

Group Beneficial Norms Can Spread Rapidly in a Structured Population

ROBERT BOYD*† AND PETER J. RICHERSON‡

**Department of Anthropology, University of California, Los Angeles, CA 90095, USA and*

‡*Department of Environmental Science and Policy, University of California, Davis, CA 95616, USA*

(Received on 20 April 2001, Accepted in revised form on 5 December 2001)

Group beneficial norms are common in human societies. The persistence of such norms is consistent with evolutionary game theory, but existing models do not provide a plausible explanation for why they are common. We show that when a model of imitation used to derive replicator dynamics in isolated populations is generalized to allow for population structure, group beneficial norms can spread rapidly under plausible conditions. We also show that this mechanism allows recombination of different group beneficial norms arising in different populations.

Introduction

Many culturally transmitted norms are group beneficial (Sober & Wilson, 1998): Property rights encourage productive effort, rules against murder and assault encourage civil order, norms governing the filling of political offices reduce the chances of civil war, and product standards, building codes, and rules of professional conduct allow more efficient commerce. For most of human history, states were weak or non-existent, and norms were not enforced by external sanctions. Nonetheless, norms were important regulators of social order, and while in modern states black-letter laws also further many of the same ends as informal norms, the evidence is that informal custom still plays a very important role in regulating behavior (Ellickson, 1991).

The *persistence* of group beneficial norms is easily explained. When people interact repeatedly, behavior can be rewarded or punished, and

such incentives can stabilize almost any behavior once there is consensus about what is normative. People conform to normative behavior in order to gain rewards or avoid punishment. The provision of rewards and punishments can be explained in several ways: first, if interactions are repeated indefinitely, punishing or rewarding also can be normative behaviors, and violators of that norm can be punished or rewarded as well (Boyd & Richerson, 1992a). Second, even if interactions do not go on indefinitely (or equivalently, people cannot remember large number of interactions), the relative disadvantage suffered by those who enforce social norms compared with those who do not rapidly becomes small as the number of interactions increases and is easily balanced by even a weak tendency to imitate the common type (Henrich & Boyd, 2001). (Of course, strong conformism can also explain the maintenance of norms without punishment, Boyd & Richerson, 1985.) Finally, punishment may be individually beneficial if it is a costly signal of an individual's qualities as a

†Author to whom correspondence should be addressed.

mate or coalition partner (Bleige Bird *et al.*, in press). Several authors suggest that the stability of such norms explains human cultural diversity—distinct groups represent alternative, stable equilibria in a complex, repeated “game of life” (Boyd & Richerson, 1992b; Binmore, 1994; Cohen, 2001).

The fact that group beneficial norms can persist does not explain why such norms are widely observed. While punishment and reward can stabilize group beneficial norms, they can also stabilize virtually any behavior (Fudenberg & Maskin, 1986; Boyd & Richerson, 1992a). We can be punished if we lie or steal, but we can also be punished if we fail to wear a tie or refuse to eat the brains of dead relatives. Thus, we need an explanation for why populations should be more likely to wind up at a group beneficial equilibrium than one of the vastly greater number of stable but non-group beneficial equilibria. Put another way, if social diversity results from many stable social equilibria, then social evolution must involve shifting among alternative stable equilibria. Group beneficial equilibria will be common only if the process of equilibrium selection tends to pick out group beneficial equilibria.

Currently, there are two different kinds of models of equilibrium selection, but neither provides a plausible explanation for the widespread existence of group beneficial norms.

Within-group models of equilibrium selection (Kandori *et al.*, 1993; Ellison, 1993; Young, 1998; Samuelson, 1997) consider the effects of random processes that act within groups to change the frequency of alternative behavioral strategies. In finite populations, sampling variation will affect patterns of interaction and replication, which in turn will lead to random fluctuations in the frequencies of types through time. As long as some mutation-like process acts to maintain variation, the probability that the population will be in any state will eventually converge to a stationary distribution. If mutation rates are low and populations are of reasonable size, most of the probability mass of the stationary distribution will pile up around the stable equilibrium of the deterministic dynamic model that has the largest basin of attraction. Since there is no necessary relationship between the size of a basin of attraction and

whether it is group beneficial, within group models do not predict that group beneficial norms will be common. Within group models also suffer from two other related problems. First, it takes a very long time for populations to shift from one equilibrium to another unless the number of interacting individuals is very small. Second, these models provide no mechanism for cumulative irreversible social change because populations are assumed to be in stochastic steady state, randomly wandering back and forth between alternative equilibria.

Between-group models posit that equilibrium selection results from the competition between groups near alternative stable equilibria. These models assume that groups at more efficient equilibria are less likely to go extinct, or more able to compete with other groups in military or economic contests. This kind of group selection process leads to the evolution of group beneficial equilibria even when groups are large, and there is substantial migration between groups (Boyd & Richerson, 1982, 1990). However, given observed rates of group extinction, the spread of group beneficial equilibria will occur too slowly to account for much observed social evolution. Calculations based on empirical data on the social extinction of small groups in highland New Guinea suggest that even though rates of extinction are appreciable, the time-scale for the substitution of one norm by a better one is on the order of a millennium (Soltis *et al.*, 1995). Moreover, these models also lack any mechanism that allows for the efficient recombination of group beneficial innovations occurring in different groups, and thus cannot easily account for the cumulative nature of social change over the last 10 000 years.

Here, we show that when the standard replicator dynamic model of evolutionary game theory is embedded in a spatially structured population, group beneficial equilibria can spread rapidly and innovations can readily recombine to form beneficial new combinations. The basic logic of this result is simple: Evolutionary game theory is applicable to human social evolution when behavioral strategies are transmitted by imitation and people who have achieved high payoffs are most likely to be imitated. Strategies which have high average

payoffs will increase in frequency, in most cases eventually leading to a stable evolutionary equilibrium state. If the payoff structure of social interactions leads to multiple stable equilibria and a population is structured, partially isolated groups can be stabilized at different equilibria with different average payoffs. Consequently, behaviors can spread from groups at high payoff equilibria to neighboring groups at lower payoff equilibria because people imitate their more successful neighbors. Such spread can be rapid because it depends on the rate at which individuals imitate new strategies, rather than the rate at which groups become extinct.

In what follows, we first derive the dynamic equations that govern replicator dynamics in a spatially structured population. We then show that these equations can lead to the rapid spread of group beneficial traits under the plausible conditions. Finally, we show that this process readily leads to the recombination of different group beneficial traits that arise in different populations.

Replicator Dynamics in a Structured Population

In many situations, people have important social interactions shaped by social norms with one group of people, but know about the behavior, and the norms that regulate it, of a larger group of people. People interact every day with the members of their local group—they exchange food, labor, and land; aid others in need; marry and care for children—transactions that are regulated by social norms that define property rights and moral obligations. However, people also often know about the behavior of others in neighboring groups. They know that we can marry our cousins here, but over there they cannot; or anyone is free to pick fruit here, while there fruit trees are owned by individuals. With this kind of population structure, payoffs are determined by the composition of the local group, but cultural traits can diffuse among groups.

To generalize evolutionary game theory to allow for this kind of population structure, consider a population that is subdivided into n large groups in which frequent social interaction occurs. Individuals are characterized by one of the k strategies. The proportion of people in group d who have strategy i is p_{id} , and the vector

of frequencies in group d is \mathbf{p}_d . Social interaction generates a payoff, $W_i(\mathbf{p}_d)$ for individuals with behavior i in group d that depends on individuals’ own strategy and the strategies of other members of their group because frequent social interaction occurs with other group members.

To allow for the possibility of cultural diffusion between groups, we adopt the following model of cultural transmission: during each time period, each individual from group f encounters an individual, their “model”, from group d with probability m_{df} and observes that individual’s strategy and payoff from social interaction during that period. We will assume that $m_{ff} > \sum_{d \neq f} m_{df}$ so that most encounters occur within social groups. After the encounter, individuals may imitate the strategy of their model.

We assume that individuals are more likely to imitate if their model has a higher payoff than they do. More formally, if an individual with behavior i from group f encounters an individual with behavior j from group d , individual i switches to j with probability

$$Pr(j|i, j) = \frac{1}{2} \left(1 + \beta \left(W_j(\mathbf{p}_d) - W_i(\mathbf{p}_f) \right) \right), \quad (1)$$

where β is a positive parameter that scales payoffs so that $0 \leq Pr(j|i, j) \leq 1$ for all \mathbf{p}_d and \mathbf{p}_f . Equation (1) implies that individuals sometimes switch to a lower payoff strategy, unlike some recent derivations of replicator dynamics (Borgers & Sarin, 1997; Schlag, 1998; Gale *et al.*, 1995). We think this model is preferable because it captures the effect of uncertainty about the payoffs of others, and because it allows diffusion between groups even when there are no payoff differences, a conservative feature that reduces the effect of population structure.

Then the frequency of behavior i in group f , p'_{if} , after one time period is given by

$$p'_{if} = \sum_d m_{df} \left[p_{if} \sum_j p_{jd} \frac{1}{2} \times \left(1 + \beta \left(W_i(\mathbf{p}_f) - W_j(\mathbf{p}_d) \right) \right) + p_{id} \sum_j p_{jf} \frac{1}{2} \left(1 + \beta \left(W_i(\mathbf{p}_d) - W_j(\mathbf{p}_f) \right) \right) \right]. \quad (2)$$

The first sum inside the square brackets gives the probability that an individual with trait i in group f remains the same and the second sum gives the probability that someone who is not i initially converts to i . Some algebraic manipulation yields the following expression for the change in the frequency of behavior i in population f :

$$p'_{if} - p_{if} = \delta p_{if} \left(1 - \sum_{d \neq f} \frac{1}{2} m_{df} \right) + \sum_{d \neq f} \frac{1}{2} m_{df} [\delta p_{1d} + (p_{1d} - p_{if}) \times (1 + \beta(\bar{W}(\mathbf{p}_d) - \bar{W}(\mathbf{p}_r)))] \quad (3)$$

where $\delta p_{if} = \beta p_{if} (W_i(\mathbf{p}_f) - \bar{W}(\mathbf{p}_f))$ is the replicator dynamic equation for strategy i in group f and is the canonical description of strategy dynamics in evolutionary game theory. Thus, when individuals imitate only members of their own group ($m_{df} = 0, d \neq f$) eqn (3) says that imitation within each group causes behaviors with the highest payoff *relative* to others in the group to increase in frequency—effects on average payoff within a group are irrelevant. When there is contact between different groups, however, the effect of a behavior on average group payoff can become important. The second term in eqn (3) includes the effect of diffusion between groups that differ in trait frequency. When payoffs do not effect imitation ($\beta = 0$), this term includes only passive diffusion. However, when individuals with higher payoffs are more likely to be imitated, there is a net flow of strategies from groups with high average payoff to groups with lower average payoff.

How Group Beneficial Equilibria Spread

Next, we show how this effect can lead to the spread of group beneficial equilibria. Consider a simple model in which there are two strategies, 1 and 2. For example, strategy 1 might be a norm-forbidding cousin marriage, while strategy 2 is the norm-allowing free choice of a spouse. Within each group, individuals who deviate from the common norm suffer because they are punished by other group members. The norm

forbidding cousin marriage might lead to higher average payoff due to the formation of wider political alliances. We formalize these ideas by assuming that the payoff to an individual with behavior 1 in group d is $W_1(p_{1d}) = 1 + s(p_{1d} - \tilde{p}) + gp_{1d}$ and the payoff to an individual using behavior 2 is $W_2(p_{1d}) = 1 + gp_{1d}$. Thus, each strategy has a higher relative payoff when common. The unstable equilibrium that divides the two basins of attraction is \tilde{p} . The parameter s measures the magnitude of the difference in payoffs of the two strategies, and g measures the effect of behavior 1 on average payoff. We assume that $g > 0$, so that groups in which behavior 1 is common have higher average payoff. For example, a norm against cousin marriage might lead to more alliance formation among clans within the group. Finally, for simplicity, we assume that social groups are arranged in a ring so individuals only imitate members of their own group and the two neighboring groups. (So that $m_{df} = m$ for the two neighbors of group f and zero otherwise.)

For a novel group beneficial trait to evolve two things must occur. First, it must become common in one population, and second it must spread from that population to others. Various random processes may cause the initial shift of one population to the group beneficial equilibrium. In finite populations, sampling variation in who is imitated (Gale *et al.*, 1995) or in patterns of interaction (Kandori *et al.*, 1993; Ellison, 1993; Young, 1998) can lead to random fluctuations in trait frequencies which can tip populations into the basin of attraction of the group beneficial equilibrium. Randomly varying environments can lead to similar shifts (Price *et al.*, 1993) in populations. Finally, individual learning can be conceptualized as a process in which individuals use data from the environment to infer the best behavior. Learning experiences of individuals within a population may often be correlated, because they are utilizing the same data. Thus, random variation in such correlated learning experiences could also cause equilibrium shifts in large populations. We do not model these processes here.

To see how imitation of the successful can lead to the spread of group beneficial strategies, assume that one of these unmodeled processes

causes the group beneficial strategy to become common in one group, while the other strategy remains common in the rest of the groups. Then, if enough individuals in the two neighboring groups imitate behavior 1, these groups will be tipped into its basin of attraction, and the group beneficial trait will increase in those two groups. This process is illustrated in Fig. 1. Trait 1 is initially common in population $i-1$. In the neighboring population i , trait 2 is common, and thus within group imitation tends to decrease the frequency of trait 1. However, individuals in population i are more likely to imitate individuals in population $i-1$ than in population $i+1$, so extra-group imitation tends to increase the frequency of trait 1 in group i . If this latter process is sufficiently strong, it can tip population i into trait 1's basin of attraction. If this occurs, the process will be repeated in group $i+1$, then group $i+2$, and so on with behavior 1 spreading throughout the population in a wave-like fashion. This process is formally similar to one recent model of the third phase of Wright's shifting balance theory (Gavrilets, 1995), but is unlike that model in two ways. First, the underlying dynamic processes arise from differential imitation, not changes in demography. Second, because the multiple equilibria arise from frequency-dependent social interaction, not underdominance, the process modeled here leads to the spread of the group beneficial trait for a wide range of parameters (Fig. 2).

It is important to see that the spread of the group beneficial trait depends crucially on the assumption that people imitate strategies that lead to success in neighboring groups, but will lower their payoff in their own group where different norms are enforced. In this simple model, a type that restricted imitation to its own group would replace the type of imitation assumed here. We think our assumption is plausible nonetheless. Empirically, the tendency to imitate the successful has been observed in a wide variety to contexts (see Henrich & Gil-White, 2000). This tendency makes sense adaptively. The world is complex and hard to understand. It is very difficult in many situations to connect behavior to outcomes with much confidence. An individual observes that in the neighboring group they never marry cousins and

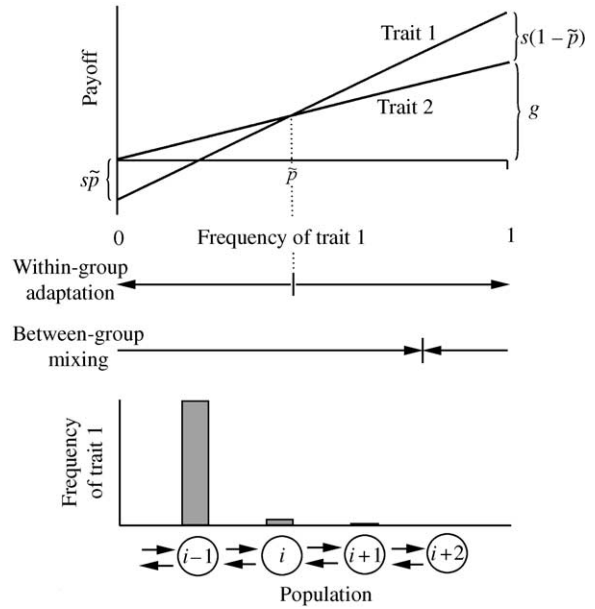


FIG. 1. This graph illustrates the assumed payoff structure and why it can lead to the spread of group beneficial traits. The top panel plots the payoffs to traits 1 and 2 as a function of the frequency of trait 1 in their local group. Each trait has a higher relative payoff when it is common, but increasing the frequency of trait 1 raises the payoff of all group members. As a result, within-group imitation increases the frequency of trait 1 above the threshold frequency \tilde{p} , and increases the frequency of trait 2 below that threshold. The lower panel shows the state of a part of population in which trait 1 is initially common in group $i-1$ and trait 2 is common in all other groups. In group i , individuals are more likely to imitate people in population $i-1$ than in population $i+1$ because the former have higher payoffs than the latter. Thus, between-group imitation tends to increase the frequency of trait 1 in population i . If this effect is strong enough, it can tip group i into the basin of attraction of trait 1, and cause the spread of this group beneficial trait.

that they are much better off. His neighbors say that the Gods punish those who marry cousins, and they have had much greater success in warfare lately. Of course, the individual knows that it will cause trouble to forbid a marriage that both his daughter and his brother want, but maybe it will be worth it. The same kinds of uncertainties beset us in the modern world despite vastly greater information gathering capacity. In the early 1990s it was commonplace to attribute Japan's economic success to their encouragement of long-term investment, their "just in time" inventory practices, or to their quality circles, and all of these practices were imitated by American firms and policy

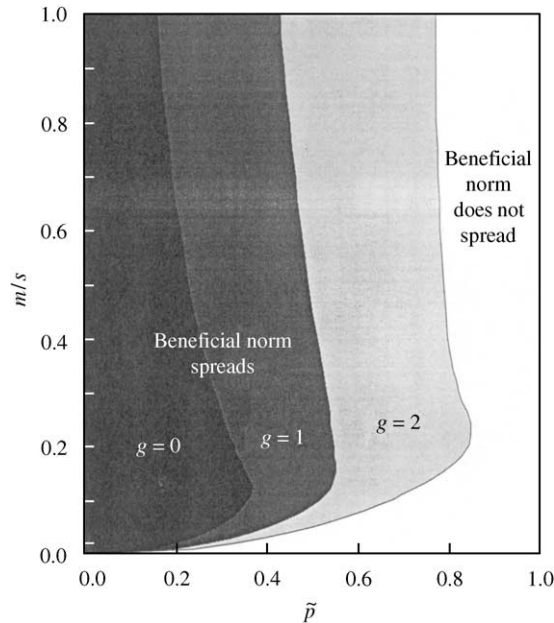


FIG. 2. This graph shows the range of parameters over which the beneficial norm spreads to all groups, eliminating the alternative norm, given that the beneficial norm is initially common in a single group. The vertical axis gives the ratio of m , the probability that individuals interact with others from one of the neighboring groups, to s , rate of change due to imitation within groups. The horizontal axis plots \tilde{p} , the unstable equilibrium that separates the basins of attraction of group beneficial and non-group beneficial equilibria in isolated groups. The shaded areas give the combinations of m/s and \tilde{p} which lead to the spread of the group beneficial strategy for three values of g . When $g=0$, neither norm is group beneficial. Larger values of g , mean that the group beneficial norm leads to a greater increase in average payoff. When m is small, the group beneficial norm cannot spread because there is not enough interaction between neighbors for the beneficial effects of the norm to cause it to spread. Very large values of m prevent the spread of the group beneficial norm because it cannot persist in the initial population. If the domain of attraction of the group beneficial strategy is too small, the flow of strategies from successful groups to less successful groups does not tip neighboring groups into its basin of attraction. Increasing the degree to which strategy 1 is group beneficial (i.e. the magnitude of g), enlarges the range of parameters which lead to the increase in that strategy. Here, the number of groups, n , was 32, but results are insensitive to n as long as it is sufficiently large. Very small values of n increase the range of parameters under which the group beneficial trait spreads. These results are from simulation—if the group beneficial trait had not spread to all groups after 10000 time periods, we assumed it would not spread. To construct the graph, we chose values of m/s and then used an interval halving algorithm to find the threshold value of \tilde{p} at which trait 1 did not spread.

makers. We have argued at length (Boyd & Richerson, 1985) that cultural transmission rules like *imitate the successful and imitate the common*

type should be seen as adaptations for dealing with this kind of uncertainty. We have a propensity to imitate the successful *because* it is often very difficult to decide what is the best behavior. These learning rules are shortcuts that on average allow us to acquire lots of useful information, but may, as in the model in this paper sometimes lead us astray.

Figure 2 plots combinations of the parameters m , s , \tilde{p} , and g that lead to the spread of the group beneficial strategy. It indicates that the group beneficial strategy fails to spread under three circumstances. If there is too much mixing between neighboring groups, the beneficial strategy cannot persist in the initial population; it is swamped by the flow of behavior 2 from the neighboring groups. If there is too little mixing, the group beneficial behavior remains common in the initial population but cannot spread because there is not enough interaction between neighbors for the beneficial effects of the norm to cause it to spread. If the domain of attraction of the group beneficial strategy is too small, the flow of ideas from successful groups to less successful groups may not be sufficient to tip neighboring groups into its basin of attraction. Increasing the degree to which strategy 1 is group beneficial (i.e. the magnitude of g) enlarges the range of parameters which lead to the increase in that strategy.

The results plotted in Fig. 3 show that the group beneficial trait spreads at a rate that is roughly comparable with the rate at which individually beneficial traits spread within a single group under the influence of the same learning process. Thus, if an individually beneficial trait can spread within a population in 10 years, a group beneficial trait will spread from one population to the next in 15–30 years, depending on the amount of mixing and the effect of the trait on average fitness. Game theorists have considered a number of mechanisms of equilibrium selection that arise because of random fluctuations in outcomes due to sampling variation and finite number of players (Kandori *et al.*, 1993; Ellison, 1993; Young, 1998; Samuelson, 1997). These processes tend to pick out the equilibrium with the largest domain of attraction. However, unless spatial structure limits interactions to a small number of




FIG. 3. This figure plots a measure of the length of time necessary for the spread of the group beneficial trait relative to the length of time necessary for the spread of an individually advantageous trait. In the simulations reported, the group beneficial trait spreads from one group to the next at a constant rate after an initial transient period. Here, we plot the ratio of the time necessary to increase from a frequency of 0.1 to 0.9 in a single group at the boundary of the wave spreading at the constant rate divided by the length of time necessary for a purely advantageous trait with dynamics $\Delta p = sp(1 - p)$ to spread from 0.1 to 0.9 in a single isolated population for two different values of the ratio m/s . As in Fig. 1, m is the probability of interacting with, and potentially imitating, an individual in each of the two neighboring groups. In both graphs, $g=1.0$, and the parameter \tilde{p} is the unstable equilibrium that divides the basins of attraction of the group beneficial trait and the other trait. These results indicate that spatial structure causes an initially individually disadvantageous but group beneficial trait to spread on roughly the same time scale as a simple individually advantageous trait whose within-group dynamics are governed by the same rate parameter s .

individuals, the rate at which this occurs in a large population is very slow. Similarly, group selection models appear to require unrealistically high group extinction rates to explain many examples of the spread of group beneficial cultural traits (Boyd & Richerson, 1990; Soltis *et al.*, 1995). In contrast, the process we describe here leads to the deterministic spread of the group beneficial trait on roughly the same time-scale as the same social learning processes cause individually beneficial traits to spread within groups.

Of course, we have not accounted for the processes that influence the rate at which the beneficial behavior initially becomes common in

a particular group. However, if the conditions for spread are satisfied, the group-beneficial trait needs to become common only in a single group. If we imagine that group-beneficial traits mainly arise as a result of random processes in small populations, only the initial group, not the whole population, needs to be small, and the group must remain small only for long enough for random processes to give rise to an initial “group mutation” which can then spread relatively rapidly to the population as a whole. If we imagine that rare events, such as the emergence of uniquely charismatic reformers or alignment of the particular constellations of political forces, are required to affect a group-favoring innovation, the same considerations apply. Only one group need make the original innovation; any others with substantial cultural contact can rapidly acquire the trait by the mechanism we model here.

Recombination at the Group Level

The process described here readily leads to the recombination of group beneficial strategies that initially arise in different groups. The exact combination of strategies necessary to support complex, adaptive social institutions would seem unlikely to arise through a single chance event. It is much more plausible that complex institutions are assembled in numerous small steps. Previous group selection models of equilibrium selection are analogous to the evolution of an asexual population in that they lack any mechanism that allows the recombination of beneficial strategies that arise in different populations, and thus require innovations to occur sequentially in the same lineage. Within group models in which equilibrium selection occurs through random sampling processes assume that the population has reached a stationary distribution, and thus while recombination is possible there is no cumulative, irreversible change. By contrast, the present model allows recombination of different strategies and irreversible, cumulative change. To see this, consider a model in which strategies consist of two components (x, y) each with two values $(0,1)$. Let p_d and q_d be the frequencies of $x = 1$ and $y = 1$ in group d , respectively. Let the payoff of an individual in

group d be

$$\begin{aligned} W_d(x, y) = & 1 + sx(p_d - \tilde{p}) \\ & + sy(q_d - \tilde{p}) \\ & + g(q_d + p_d). \end{aligned} \quad (4)$$

Thus, both $x = 1$ and $y = 1$ have an independent group beneficial effect, and all four combinations of x and y can be stable equilibria in isolated groups. Finally, suppose that individuals occasionally learn the x component of their strategy from one individual and the y component from another, leading to recombination of behavioral strategies at the individual level. Once again suppose that the population is initially all strategy $(0,0)$, and that random shocks cause $(1,0)$ to become common in one population and $(0,1)$ common in a second population. Then, if conditions are right, both strategies will begin to spread [(Fig. 4(a)]. When the two waves meet, the frequency of $x = 1$ is equal to one half and the frequency of $y = 1$ is equal to one half at the boundary between the two expanding fronts. The outcome depends on the value of \tilde{p} . If $\tilde{p} < \frac{1}{2}$, the strategy $(1,1)$ has the highest payoff in the group on the boundary, increases deterministically in that group, and eventually spreads throughout the population as a whole [(Fig. 4(b)]. If $\tilde{p} > \frac{1}{2}$, the strategy $(1,1)$ has a lower payoff than $(1,0)$ or $(0,1)$, and the two waves form a stable boundary. However, in the boundary group, the most beneficial combination $(1,1)$, has a relatively small payoff disadvantage compared with $(0,1)$, and $(0,1)$, is present at substantial frequency. In this situation, a shift to the most beneficial combination due to random shocks is much more likely than the shifts that were necessary to cause $(0,1)$ and $(1,0)$ to become common in the first place. Thus, existing group beneficial traits will recombine more rapidly than new ones arise.

Conclusions

Many anthropologists and sociologists have long believed that human behavior is regulated by culturally transmitted norms in ways that promote the survival and growth of human societies. Economists and other rational choice

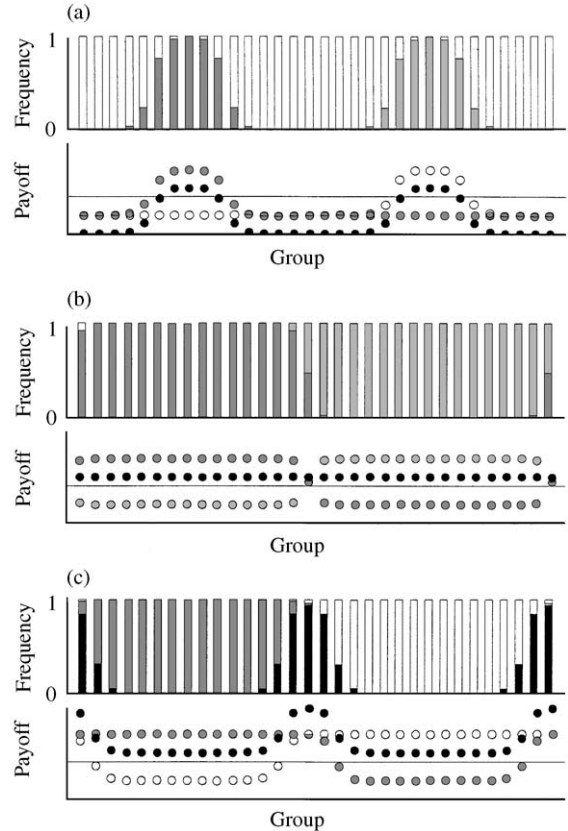


FIG. 4. In (a), (b), and (c) the upper graph plots the frequencies of the four possible strategies as a stacked bar graphs for each of 32 groups: $(0,0)$ white, $(1,0)$ light gray, $(0,1)$ dark gray, and $(1,1)$ black. The lower graph plots the payoff to each strategy net of the group effects in each group. The (—) line gives the payoff of $(0,0)$ and the (•••) circles give the payoffs of the other three strategies. The parameters are $m=0.02$, $s=0.1$, $\tilde{p}=0.4$, and $g=2$. (a) Initially $(0,1)$ is common in group 8 and $(1,0)$ is common in group 24, and the two group beneficial traits begin to spread. (b) When the two spreading fronts meet, the frequencies of $x=1$ and $y=1$ are one-half, which means that the strategy $(1,1)$ has a highest payoff. (c) Recombination at the individual level introduces strategy $(1,1)$ into the boundary group and strategy $(1,1)$ then spreads deterministically, first in that group, and then to adjacent groups.

theorists have been skeptical about such functionalist claims because there was no plausible mechanism to explain why such norms should be common. Social scientists influenced by evolutionary biology tend to share this skepticism based upon theoretical models and empirical findings suggesting that group selection is generally a weak force in nature. We believe that humans are an exception to this rule because cultural variation is much more susceptible to

group selection than genetic variation. The cultural group selection hypothesis explains both why humans cooperate on such a large scale and why the pattern of this cooperation is so different from that of other ultrasocial animals (Richerson & Boyd, 1999). Human societies are based upon cooperation between non-relatives, while kinship underlies cooperation and complex sociality in other taxa like the social insects.

Despite a general fit between the existing models of cultural group selection and the facts of human sociality, much uncertainty remains. Earlier work suggests that the differential survival of culturally distinctive groups can lead to the evolution of group beneficial behavior under plausible circumstances, but that this process is quite slow and likely to produce historically contingent group-level adaptations (Boyd & Richerson, 1982, 1990; Soltis *et al.*, 1995). Since the evolution of human social institutions does have a time-scale of millennia and the resulting institutions are highly variable, such group selection processes may have had a role in shaping these institutions. On the other hand, some social institutions do diffuse from one society to another and on time-scales shorter than a millennium. The spread of the joint stock company on time-scales of a century is a recent example. Such events accord better with a mechanism like the one we model here.

We suspect that both differential survival and differential diffusion may affect the evolution of human social institutions. The operation of many social institutions is opaque even to the people who enact them (Nelson & Winter, 1982, Chap. 5) and such institutions are even harder for outsiders to understand. In such cases, diffusion may be ineffective because actors cannot connect the attributes of particular institutions to their success, and this fact may explain why the path from the origins of agriculture to our complex modern industrial nations took some ten millennia to traverse. Other institutions spread much more readily because their costs and benefits are more readily understood. Proselytizing religions, for example, take pains to be transparent to potential converts and thus may readily spread. The rate of diffusion of institutions may also be affected by how much people know about other societies.

It is plausible that the spread of literacy and the development of ever better means of transportation have gradually increased the importance of the rapid processes based on borrowing relative to the slower ones based on group extinction. In the 20th century, social institutions like central banks, soccer, and government bureaucracies have become all but universal in about a century. Nevertheless, globalization is incomplete; dramatic differences exist even between modern societies (Nisbett *et al.*, in press). Some elements of culture likely still have time-scales of change measured in millennia.

We thank Sam Bowles, Ernst Fehr, Daniel Friedman, Francisco Gil-White, Herb Gintis, Joe Henrich, Richard McElreath, Rajiv Sethi, David Sloan Wilson, and several anonymous reviewers for useful comments on the research reported here.

REFERENCES

- BLEIGE BIRD, R., SMITH, E. A. & BIRD, D. (in press). The hunting handicap: costly signaling in human foraging strategies. *Behav. Ecol. Sociobiol.*
- BINMORE, K. (1994). *Game Theory and the Social Contract*. Cambridge: MIT Press.
- BORGERS, T. & SARIN, R. (1997). Learning through reinforcement and replicator dynamics. *J. Econ. Theory* **77**, 1–14.
- BOYD, R. & RICHERSON, P. J. (1992a). Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethol. Sociobiol.* **13**, 171–195.
- BOYD, R. & RICHERSON, P. J. (1992b). How microevolutionary processes give rise to history. In: *Evolution and History* (Nitecki, M. H. & Nitecki, D. V., eds), pp. 179–209. Albany: State University of New York Press.
- BOYD, R. & RICHERSON, P. J. (1990). Group selection among alternative evolutionarily stable strategies. *J. theor. Biol.* **145**, 331–342.
- BOYD, R. & RICHERSON, P. J. (1985). *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.
- BOYD, R. & RICHERSON, P. J. (1982). Cultural transmission and the evolution of cooperative behavior. *Hum. Ecol.* **10**, 325–351.
- COHEN, D. (2001). Cultural variation: considerations and implications. *Psychol. Bull.* **127**, 451–471.
- ELICKSON, R. C. (1991). *Order Without Law: How Neighbors Settle Disputes*. Cambridge: Harvard University Press.
- ELLISON, D. (1993). Learning, social interaction, and coordination. *Econometrica* **61**, 1047–1071.
- FUNDENBURG, D. S. & MASKIN, E. (1998). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* **51**, 533–554.
- GALE, J., BINMORE, K. G. & SAMUELSON, L. (1995). Learning to be imperfect: the ultimatum game. *Games Econ. Behav.* **8**, 56–90.

- GAVRILETS, S. (1995). On phase three of the shifting balance theory. *Evolution* **50**, 1034–1041.
- HENRICH, J. & BOYD, R. (2001). Why people punish defectors: weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. *J. theor. Biol.* **208**, 79–89.
- HENRICH, J. & GIL-WHITE, F. (2001). The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evol. Hum. Behav.* **22**, 165–196.
- KANDORI, M., MAILATH, G. & ROB, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica* **61**, 29–56.
- NELSON, R. R. & WINTER, S. G. (1982). *An Evolutionary Theory of Economic Change*. Cambridge: Harvard University Press.
- NISBETT, R. E. PENG, K., CHOI, I. & NORENZAYAN, A. (2001). Culture and systems of thought: holistic vs. analytic cognition. *Psychol. Rev.* **108**, 291–310.
- PRICE, T., TURELLI, M. & SLATKIN, M. (1993). Peak shifts produced by correlated response to selection. *Evolution* **4**, 280–290.
- RICHERSON, P. J. & BOYD, R. (1999). The evolutionary dynamics of a crude super organism. *Hum. Nat.* **10**, 253–289.
- SAMUELSON, L. (1997). *Evolutionary Games and Equilibrium Selection*. Cambridge: MIT Press.
- SCHLAG, K. H. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *J. Econ. Theory* **78**, 130–156.
- SOBER, E. & WILSON, D. S. (1998). *Unto Others*. Cambridge: Harvard University Press.
- SOLTIS, J., BOYD, R. & RICHERSON, P. J. (1995). Can group functional behaviors evolve by cultural group selection? An empirical test. *Curr. Anthropol.* **36**, 473–494.
- YOUNG, P. (1998). *Individual Strategy and Social Structure*. Princeton: Princeton University Press.