Unpublished manuscript prepared for the Dahlem Workshop on
Bounded Rationality: The Adaptive Toolbox

Norms and Bounded Rationality

Robert Boyd
Department of Anthrpology
University of California
Los Angeles, CA 90095
rboyd@anthro.ucla.edu

Peter J. Richerson
Division of Environmental Studies
University of California
Davis, CA 95616
pjricherson@ucdavis.edu

ABSTRACT

Anthropologists believe that human behavior is governed by culturally transmitted norms, and that such norms contain accumulated wisdom that allows people to behave sensibly even though they do not understand why they do what they do. Economists and other rational choice theorists have been skeptical about functionalist claims because anthropologists have not provided any plausible mechanism which could explain why norms have this property. Here, we outline two such mechanisms. We show that occasional learning when coupled with cultural transmission and a tendency to conform can lead to the spread of sensible norms even though very few people understand why they are sensible. We also show that norms that help solve problems of self-control that arise from time-inconsistent preferences can spread if individuals tend to imitate successful people and are occasionally influenced by members of other groups with different norms.

DO NORMS HELP PEOPLE MAKE GOOD DECISIONS WITHOUT MUCH THOUGHT?

Many anthropologists believe that people follow the social norms of their society without much thought. On this view, human behavior is mainly the result of social norms, and rarely the result of considered decisions. In recent years, there has been increased interest within anthropology in how individuals and groups struggle to modify and reinterpret norms to further their own interests. However, we think it is fair to say that most anthropologists still believe that culture plays a powerful role in shaping how people think and what they do.

Many anthropologists also believe that social norms lead to adaptive behavior—by following norms, people can behave sensibly without having to understand why they do what they do. For example, throughout the New World people who rely on corn as a staple food process the grain by soaking it in a strong base (such as calcium hydroxide) to produce foods like hominy and masa (Katz 1974). This alkalai process is complicated, requires hard work, and substantially reduces the caloric content of corn. However, it also increases the amount of available lysine, the amino acid in which corn is most deficient. Katz (1974) argues that alkalai processing plays a crucial role in preventing protein deficiency disease in regions where the majority of calories are derived from corn. Traditional peoples had no understanding of the nutritional value of alkalai processing, rather it was a norm—We Maya eat masa because that is what we do. Nonetheless, by following the norm traditional people were able solve an important, and difficult nutritional problem. The work of cultural ecologists such as Marvin Harris (1979) provide many other examples of this kind, although few are as well worked out. Other varieties of functionalism (see Turner and Maryanski, 1979 for a discussion) also hold that social norms evolve to adapt to the local environment. While nowadays anthropologists are explicitly critical of functionalism, cryptic functionalism still pervades much thinking in anthropology (Edgerton, 1992).

Norms may also lead to sensible behavior by proscribing choices that people find tempting in the short run, but are damaging in the long run. Moral systems around the world have proscriptions against drunkeness, lazyness, gluttony, and other failures of self-control. There is evidence that such proscriptions can increase individual wellbeing. For example, Jensen and Ericson (1979) show that Mormon youths in Tucson are less likely to be involved in “victimless crimes” such as drinking, and marijuana use, than members of a nonreligious control group. Moreover these differences seem to have consequences: McEvoy and Land (1981) report that age adjusted mortalities for Mormons in Missouri are approximately 20% lower than control populations, and the differences were biggest for lung cancer, pneumonia/influenza, and violent death, sources of mortality that should be reduced if the abstentious Mormon norms are being observed. Apparently, living in a group in which there are norms against alcohol use makes it easier for young Mormons to do what is in their own long term interest.

WHAT ARE NORMS, AND WHY DO PEOPLE FOLLOW THEM?

Examples like these present a series of interesting questions to economists, psychologists, and others who start with individuals as the basic building blocks of social theory. First, what are norms? How can we incorporate the notion that there are shared social rules into models that assume that people are goal-oriented decision makers? Second, why should people follow norms? Norms will change behavior only if they prescribe behavior that differs from what people would do in the absence of norms. Finally, why should norms be sensible? If individuals cannot (or do not) determine what is sensible, why should norms prescribe sensible behavior? It seems more plausible that they will simply represent random noise or even superstitious nonsense.

A recent efflorescence of interest in norms among rational choice theorists provides one cogent answer to the first two questions. Norms are the result of shared notions of appropriate behavior, and the willingness of individuals to reward appropriate behavior and punish inappropriate behavior (See McAdams 1998 a reviews). Thus, it is a norm for men to remove their hats when they enter a Christian church because they will suffer the disapproval of others if they don't. In contrast, it is not a norm for men to remove their hats in a overheated concert hall, even if everyone does so. By this notion, people obey norms because they are rewarded by others if they do and punished if they do not. As long as the rewards and punishments are sufficiently large, norms can stabilize a vast range of different behaviors. Norms can require property to be passed to the oldest son or to the youngest; they can specify that horsemeat is a delicacy or deem it unfit for human consumption.

There is no consensus in this literature about why people choose to punish norm violators and reward norm followers. There have been a number of proposals: Binmore (1998) argues that social life is an infinite game and that norms are game theoretic equilibria of the kind envisioned in the folk theorem. Norm violators are punished, and so are people who fail to punish norm violators, people who fail to punish them, and so on ad infinitum. McAdams (1998) suggests that all people desire the esteem of others, and because esteem can be "produced" at very low cost, it is easy to punish norm violators by witholding esteem. Bowles and Gintis (in press) and Richerson and Boyd (1998) argue that group selection acting over the long history of human evolution created a social environment in which natural selection favored genes leading to a reciprocal psychology. Here we will simply assume that this second order free riding problem has somehow been solved.

HOW DO NORMS SOLVE PROBLEMS THAT PEOPLE CANNOT SOLVE ON THEIR OWN?

Virtually all of the recent literature on norms focuses on how norms help people solve public goods and coordination problems (e.g. Ostrom et. al 1991, Ellickson 1994). It does not explain why norms should be adaptive. If people do not understand why alkali treatment of corn is a good thing, why should they require their neighbors to eat masa and hominy, and be offended if they do not? Nor does the recent literature on norms explain why norms should commonly help people with problems of self-control. If people cannot resist the temptations of alcohol, why should they insist that their neighbors do so? Here we sketch possible answers to these questions.

Occasional learning plus conformism leads to adaptive norms

In this section we show how a small amount of individual learning when coupled with cultural transmission and a tendency to conform to the behavior of others can lead to adaptive norms, even though most people simply do what everyone else is doing

Why it may be sensible for most people to imitate

It is easy to see why people may choose to imitate others when it is costly or difficult to determine the best behavior—copying is easier than invention and plagiarism is easier than creation. However, as these examples illustrate, it is not clear that by saving such costs, imitation makes everybody better off—this is why we have patents and rules against plagiarism. We have analyzed a series of mathematical models which indicate that when decisions are difficult, everyone can be better off if most people imitate the decisions of others under the right circumstances (e.g. Boyd and Richerson, 1985, 1988, 1989, 1995, 1996). The following simple model illustrates our reasoning.

Consider a population that lives in an environment that switches between two states with a constant probability. Further assume that there are two behaviors, one best in each environmental state. All individuals attempt to discover the best behavior in the current environment. First, each individual experiments with both behaviors and then compares the results. The results of such experiments vary for many reasons, and so the behavior that is best during any particular trial may be inferior over the long run. To represent this idea mathematically assume that the observed difference in payoffs is a normally distributed random variable, *X* (Figure 1). Second, each individual can observe the behavior of an individual from the previous generation who has already made the decision.

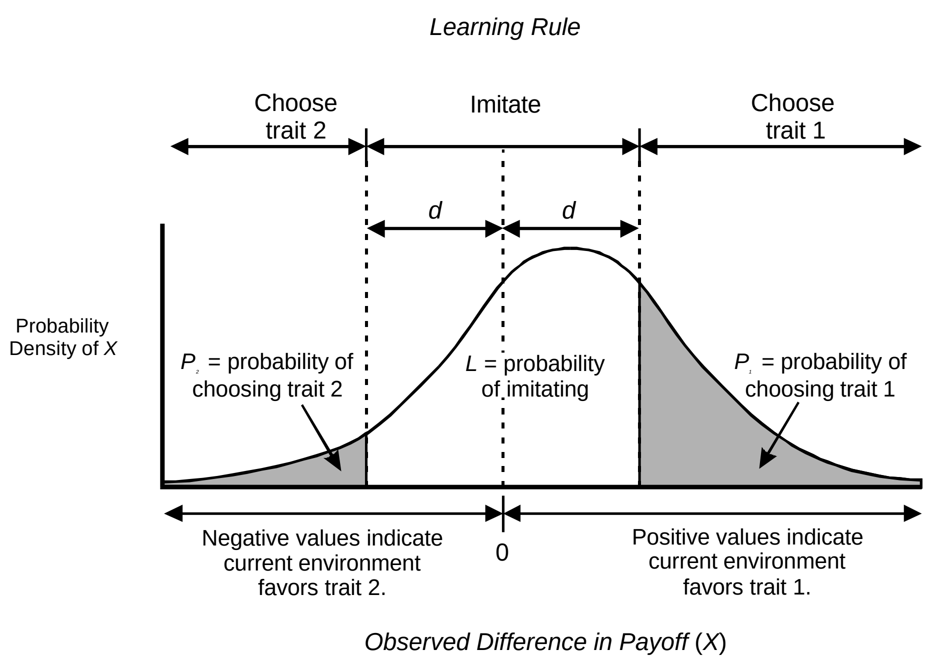

Figure 1. A graphical representation the model of individual and social learning. Each individual observes an independent, normally distributed environmental cue, *X*. A positive value of *X* indicates that the environment is in state 1; a negative value indicates that the environment is in state 2. If the value of *X* is larger than the threshold value, *d*, an individual adopts trait 1. This occurs with probability, p_1. If the value of the environmental cue is smaller than than $-d$, the individual adopts trait 2 which occurs with probability, p_2. Otherwise the individual imitates. Thus, the larger the standard deviation of the cue compared to its mean value, the is the predictive value of the cue.

We assume that individuals combine sources of information by adopting a particular behavior if its payoff appears *sufficiently* better than its alternative; otherwise they imitate. The larger the observed difference in the payoffs between the two behaviors, the more likely it is that the behavior with the higher payoff actually is best. By insisting on a large difference in observed payoff individuals can reduce the chance that they will mistakenly adopt the inferior behavior. Of course, being discriminating will also cause more trials to be indecisive, and then, they must imitate. Thus, there is a tradeoff. Individuals can increase the accuracy of learning but only by also increasing the probability that learning will be indecisive and having to rely on imitation

The optimal decision rule depends on what the rest of the population is doing. Assume that most individuals use a learning rule that causes them to imitate *x*% of the time— call these "common-type individuals." There are also a few rare invaders who imitate slightly more often. Compared to the common type, invaders are less likely to make learning errors. Thus, when invaders learn, they have higher payoff than the common-type individuals when they learn. When invaders imitate, they have the same payoff as the common-type individuals.. However, invaders must imitate more often and imitators always have lower fitness than learners. To see why, think of each imitator as being connected to a learner by a chain of imitation. If the learner at the end of chain learned in the current environment, then the imitator has the same chance of acquiring the favored behavior as does a learner. If the learner at the end of the chain learned in a different environment, the imitator will have a lower chance of acquiring the best behavior. Thus, the invading type will achieve a higher payoff if the advantage of making fewer learning errors is sufficient to offset the disadvantage of imitating more.

This tradeoff depends on how much the common type imitates. When the common type rarely imitates, the payoff of individuals who imitate and individuals who learn will be similar because most imitators will imitate somebody who learned, and the fact that mutants make fewer learning errors will allow them to invade. However, as the amount of imitation increases, the payoff of imitating individuals relative to those who learn declines because increased imitation lengthens the chain connecting each imitator to a learner. Eventually an equilibrium is reached at which the common type can resist invasion by mutants that change the rate of imitation. Figure 2 plots the probability that individuals imitate (denoted as *L* in figure 1) at evolutionary equilibrium as a function of the quality of the information available to individuals for three different rates of environmental change (See Boyd and Richerson 1988 for details of the calculation). Notice that when it is difficult to for individuals to determine the best behavior and when environments change infrequently more than 90% of the population simply copies the behavior of others.

As long as environments are not completely unpredictable, the average payoff at the evolutionary equilibrium is greater than the average payoff of individuals who do not imitate (Boyd and Richerson, 1995.). The reason is simple: imitation allows the population to learn when the information is good and imitate when it is bad. Figure 3 plots the average payoff of imitating and learning individuals as a function of the fraction of common type individuals who imitate. The payoff of learning individuals increases as the amount of imitation increases because learners are more selective and therefore make fewer errors. The payoff of imitating individuals *also* increases at first because they are imitating learners who make fewer errors. If imitation is common enough, payoff eventually declines because the population fails to track the changing environment. The first effect is sufficient to lead to a net increase in average payoff at evolutionary equilibrium.

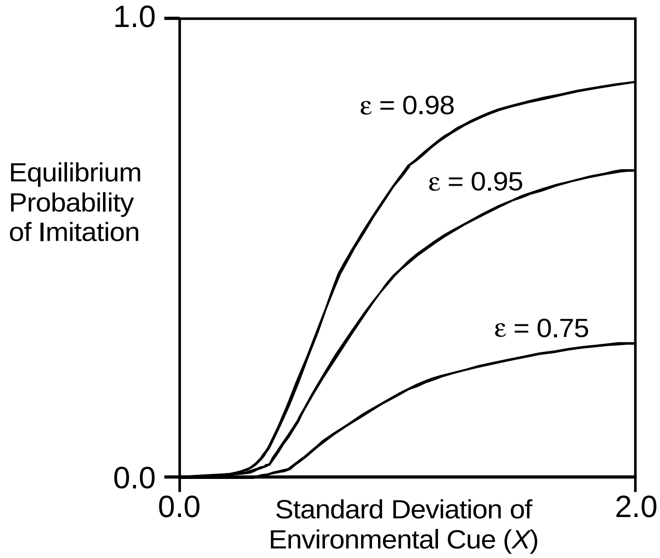

Figure 2. Probability that individuals imitate (L) at evolutionary equilibrium as a function of the quality of the environmental cue for three different rates of environmental change. The mean of the environmentat cue (X) is 1.0, so as the standard deviation of X increases, the extent to which the cue predicts the environmental state decreases. Thus, the results plotted here indicate that as the predictive quality of the cue decreases, the probability of imitation at evolutionary equilibrium increases. The parameter ***e*** is the probability that the environment remains unchanged from one time period to the next. Thus as the rate of environmental change decreases, the probability of imitation at evolutionary equilibrium increases. See Boyd and Richerson (1988) for details.

We believe that the lessons of this model are robust. It formalizes three basic assumptions:

1. The environment varies.
2. Cues about the environment are imperfect, so individuals make errors.
3. Imitation increases the accuracy (or reduces the cost) of learning.

We have analyzed several different models which incorporate these assumptions, but differ in other features. All of these models lead to the same qualitative conclusion: when learning is difficult and environments don't change too fast, most individuals imitate at evolutionary equilibrium. At that equilibrium the population is better off, on average, than a population that does not imitate.

Adding Conformism

So far we have only shown that it may be best for most people to copy others rather than try to figure out things for themselves. Recall that for something to be a norm, there has to be a conformist element. People must agree on the appropriate behavior and disapprove of others who do not behave appropriately. We now show that individuals who respond to such disapproval by conforming to the social norm are *more* likely to acquire the best behavior. We will also show that as the tendency to conform increases, so does the equilibrium amount of imitation.

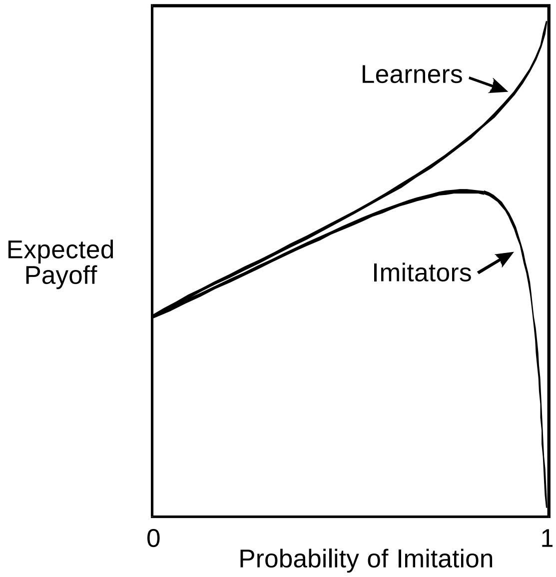

Figure 3. Individuals either learn or imitate according to the outcome of their learning trial. As individuals become more selective, the frequency of imitating individuals increases. This figure plots the expected fitness of individuals who imitate and those who learn as a function of the probability that an individual randomly chosen from the population imitates (assuming the outcome of learning experiments is normally distributed with mean 0.5 and variance 1.)

To allow the possibility for conformist pressure, we add the following assumption to the model described above: when an individual imitates, she may be disproportionately likely to acquire the more common variant. Let q be the fraction of the population using trait 1. As before individuals collect information about the best behavior in the current environment, and then if the information is not decisive, they imitate. However, now the probability that an imitating individual acquires trait 1 is:

$$\Pr(1) = q + \Delta q(1-q)(2q-1) \qquad (1)$$

Thus Δ represents the extent to which individuals respond to the blandishments of others. When $\Delta = 0$, the individuals ignore conformist pressures, and model is the same as the one described in the previous section. When $\Delta > 0$, social pressure induces individuals to adopt the more common of the two behaviors. When $\Delta \approx 1$, individuals almost always adopt the same behavior as the majority.

We now determine the equilibrium values of Δ and L (the probability of relying on imitation) in the same way that we determined the equilibrium amount of imitation above. Assume that most of the population is characterized by one pair of values of Δ and L. Then, consider whether that population can be invaded by individuals using slightly different values of Δ and L. The evolutionary equilibrium is the combination of values of Δ and L that cannot be invaded in this way.

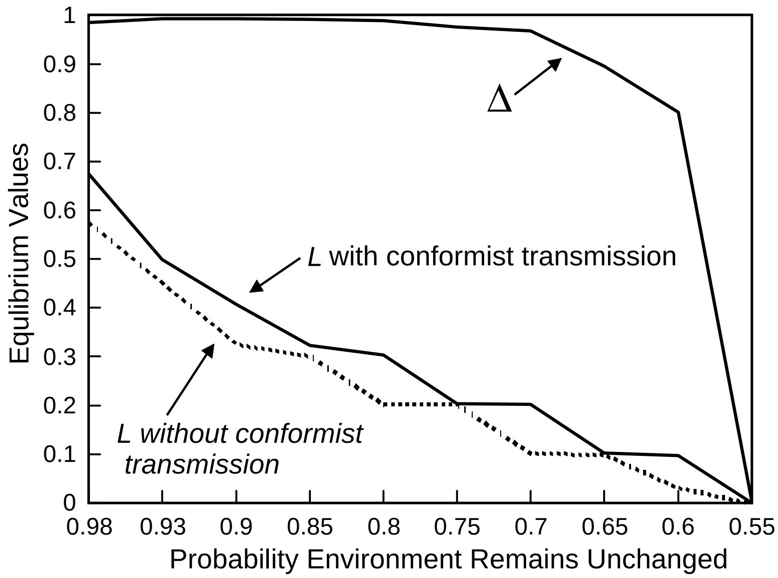

Figure 4 . Equilibrium values of *L* and Δ for different rates of environmental variation. At evolutionary equilibrium, the strength of conformist transmission is high for a wide range of rates of environmental change. However, the reliance on social learning (*L* as proportion ranging from 0 to 1.0) decreases rapidly over the same range of environmental stability. When there is no conformist effect (Δ is constrained to be zero), the evolutionary equilibrium value of *L* is lower than when Δ is free to evolve to its equilibrium value. Because the conformist effect causes the population to track the environment more effectively, it makes social learning more useful. For more details on this calculation see Henrich and Boyd (1998).

This analysis leads to two robust results. First, all conditions that lead a substantial fraction of the population to rely on imitation also lead to very strong conformism. Consider, for example, Figure 4 which plots the equilibrium values of Δ and *L* as a function of the rate of environmental change. Notice that as long as the environment is not completely unpredictable, the equilibrium value of Δ is near its maximum value—when people imitate, they virtually always do what the majority of the population is doing. As detailed in Henrich and Boyd (1998), the equilibrium values of Δ and *L* are equally insensitive to other parameters in the model. Second, as conformism increases, so does the fraction of the population that relies on imitation. Figure 4 shows that the equilibrium value of *L* when both *L* and Δ are allowed to evolve is larger than the equilibrium value of *L* in a model in which Δ is constrained to be zero. Thus, a tendency to conform increases the number of people who follow social norms and decreases the number who think for themselves.

These results are easy to understand. Just after the environment switches, most people acquire the wrong behavior. Then, the combination of occasional learning and imitation causes the best behavior to gradually become more common in the population until an equilibrium is reached at which most of the people are characterized by the better behavior. For rates of environmental change that favor substantial reliance on imitation, the best behavior is more common than the alternative averaged over this entire cycle. Thus, individuals with a conformist tendency to adopt the most common behavior when in doubt are more likely to acquire the best behavior. Conformism continues to increase until it becomes so strong that it prevents the population from responding adaptively after an environmental shift. Increasing conformism leads to increased imitation because on average conformism causes imitators to be more likely to acquire the best behavior in the current environment.

Imitation of successful neighbors leads the spread of beneficial norms.

There is by now a large literature that indicates that people often have time-inconsistent preferences and as a result, they often make choices in the short run that they know are not in their long run interest. It is plausible that social norms help people solve these problems by creating short run incentives to do the right thing. I may not be able to resist a drink when the costs are all in the distant future, but make a different decision if I suffer immediate social disapproval. It is also easy to see why such norms persist once they are established. If everyone agrees that self-control is proper behavior and punish people who disagree, then the norm will persist. The problem is that the same mechanism can stabilize any norm. People could just as easily agree that excessive drinking is proper behavior and punish people teetotallers. If it *is* true that norms often promote self-control, then we need an explanation of why such norms are likely to arise and spread. In this section, we sketch one such mechanism. Namely, if people modify their beliefs by imitating the successful and if they sometime imitate people from neighboring groups with different norms, then under the right circumstances norms that solve self-control problems will spread from one group to another because their enforcement makes people more successful and therefore more likely to be imitated.

Consider a model in which a population is subdivided into *n* social groups (number $d = 1, \ldots n$). There are two alternative behaviors: Individuals can be self-indulgent or abstentious. Self-indulgent individuals succumb to the temptations of strong drink, while abstentious individuals restrain themselves. Abstentious individuals are better off in the long run. They make more money, live longer, are healthier, and so on, and everyone agrees that the short run pleasures of the bottle are not sufficient to compensate for the long run costs that result. Nonetheless, because individuals do not have time-consistent preferences, everyone succumbs to the temptations of the table and drinks to excess.

Next, assume that there are two social norms governing consumption behavior. People can be *Puritanical* or *Tolerant*. Puritans believe that alcohol consumption is wrong and disapprove of those who drink. Tolerant people believe everyone should make their own consumption decisions. Each type disapproves of the other—puritans do not believe that one should tolerate excess, and the tolerant think that others should be tolerant as well. These norms affect the costs and benefits of the two behaviors. When puritans are present, the people who drink suffer social disapproval, and because this cost is incurred *immediately*, it can cause people to choose not to drink when they otherwise might. Thus as the proportion of the population who hold puritanical beliefs increases, the proportion of people who drink decreases, and people are better off in the long run. To formalize these ideas, let p_d be the frequency of the puritanical norm in group *d*. Then W_1, the average payoff of puritanical individuals in group *d*, is given by:

$$W_1(p_d) = W_0 - s(1 - p_d) + gp_d \qquad (2)$$

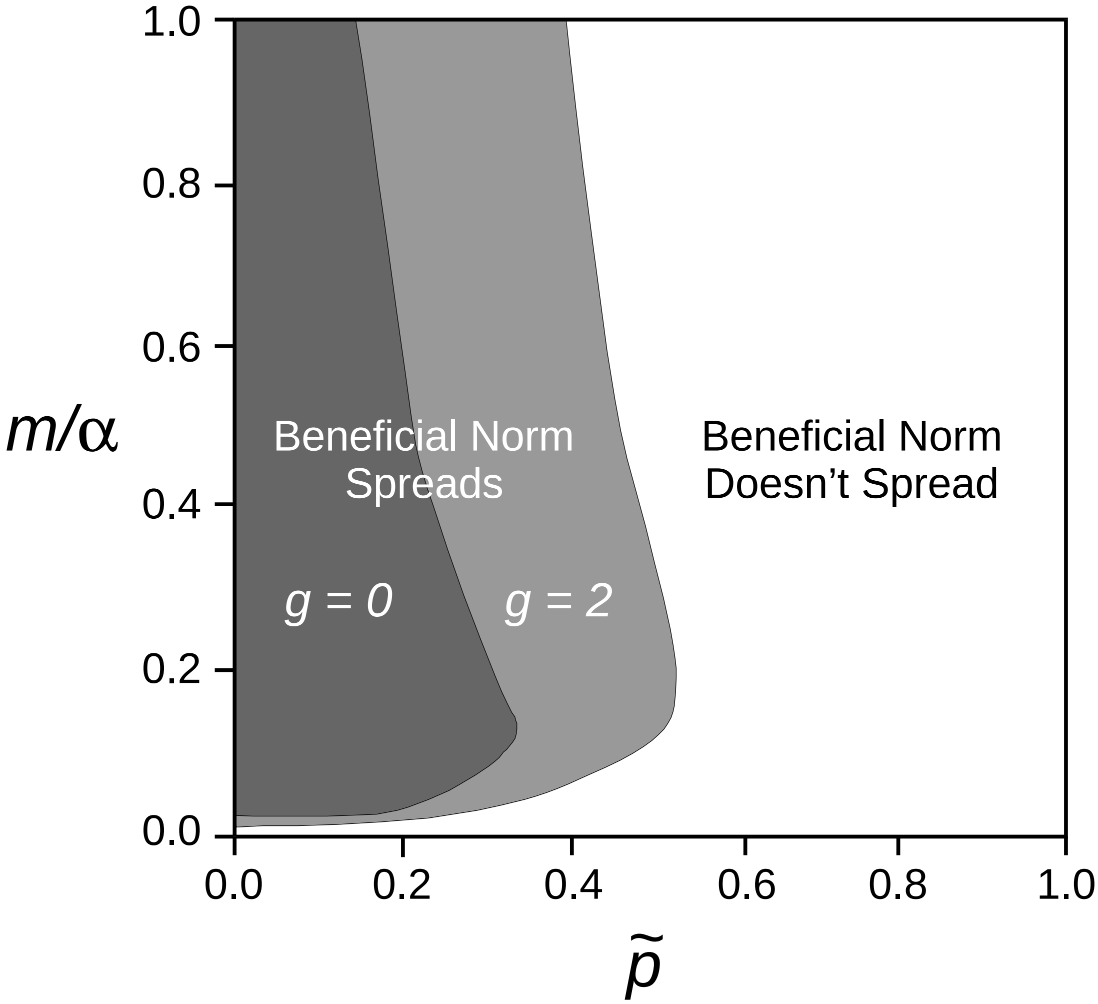

Figure 5. Plots parameter combinations which lead to the spread of the group beneficial norm between groups. The vertical axis give the ratio of *m*, the probability that individuals interact with others outside of their group to *a* rate of change due to imitation within groups, and horizontal axis plots $\tilde{p}$, the unstable equilibrium that separates the domains of attraction of puritanical and tolerant equilibria in isolated groups. Each axis was divided into 51 increments, and the model was simulated until a stationary state was reached for all 2601 combinations. There were 20 subpopulations in all simulations. Boundaries were smoothed by interpolation. The shaded areas give the combinations of *m/a* and $\tilde{p}$ which lead to the spread of the puritanical norm given that it has become common in a single group for two values of *g*, the extent to which individual behavior is affected by norms. Notice that the beneficial norm spreads when levels of interaction between groups is intermediate. If there is too much mixing the puritanical norm cannot persist in the initial population. If there is too little mixing, it can persist in the intitial population but cannot spread.

W_2, the average payoff of tolerant individuals in group *d*, is given by:

$$W_2(p_d) = W_0 - sp_d + (g+d)p_d \tag{3}$$

Individuals of each type suffer disapproval and a reduction in welfare when the other type is present in their social group. These social effects on welfare are represented by the terms proportional to s in two expressions above. However, the welfare of all individuals is increased by the fraction of puritanical individuals because everybody is less likely to drink when puritans are present to shame them. These effects are represented by the terms proportional g and $g+d$. The parameter d captures the idea that puritans may have a different effect on each other than they do on the tolerant, perhaps bigger because they are more sensitive to the opinions of their own kind, perhaps smaller because they are already avoiding strong drink.

Next, the evolution of these norms within a population is governed by the following process. During each time period, each individual encounters another person, compares her welfare to the person she encounters, and then with probability proportional to the difference between their payoffs during the last time period, adopts that person's norm. In particular, suppose that an individual with norm *i* from group *f* encounters an individual with norm *j* from group *d*. After the encounter the probability that an individual switches to *j* is:

$$\Pr(j \mid i, j) = \tfrac{1}{2}\left(1 + b\left(W_j(p_f) - W_i(p_d)\right)\right) \tag{4}$$

Notice that since an individual's payoff depends on the composition of his group, there will be a tendency for ideas to spread from groups in which beneficial norms are common to groups in which less beneficial norms are common.

Let m_{df} be the probability that an individual from group *d* encounters an individual from group *f*. Then, the change in the frequency of a belief is governed by the payoff of each belief relative to the average payoff in the population and, Δp_f, the change in p_f during one time period is given by:

$$\begin{aligned}\Delta p_f &= bp_f\left(W_1(p_f) - \overline{W}(p_f)\right)\\ &\quad + \sum_{d \neq f} m_{df}\, b\left[p_d\left(W_1(p_d) - \overline{W}(p_d)\right) - p_f\left(W_1(p_f) - \overline{W}(p_f)\right)\right]\\ &\quad + \sum_{d \neq f} m_{df}\left(p_d - p_f\right)\left(1 + b\left(\overline{W}(p_d) - \overline{W}(p_f)\right)\right)\end{aligned} \tag{5}$$

To make sense of this expression, first assume that people only encounter individuals from their own social group. Then

$$\Delta p_f = bp_f\left(W_1(p_f) - \overline{W}(p_f)\right) \tag{6}$$

This is the ordinary replicator dynamic equation. This equation simplifies to have the form:

$$\Delta p_f = ap_f(1 - p_f)(p - \tilde{p}) \tag{7}$$

where $a = b(2s - d)$ and $\tilde{p} = \frac{s}{2s - d}$. Thus, when each social group is isolated and the effects of social sanctions are large compared to the differences in their effects on drinking ($s > d$) there are two stable evolutionary equilibria, groups consisting of all puritans or all tolerant individuals. If the presence of puritans benefits other puritans more than it benefits the tolerant ($d < 0$), then the all puritan equilibrium has a larger basin of attraction. If puritans benefit the tolerant more, then the all tolerant equilibrium has a larger basin of attraction.

When there is contact between different groups, the second two terms in (5) affect the change in frequency of norms within social groups. The third term is of most interest here. If $b = 0$, this term is proportional to the difference in the frequency between the groups and simply represents passive diffusion. If, however, $b > 0$ there is a greater flow of norms from groups with high average payoff to groups with lower average payoff. This differential flow arises because people imitate the successful and norms affect the average welfare of group members. Can this effect lead to the systematic spread of beneficial norms?

For the beneficial puritanical norm to spread two things must occur. First, such a norm must increase to substantial frequency in one group. Second, it must spread to other groups. Here we address only the second question. To keep things simple, we further assume that social groups are arranged in a ring, and that individuals only have contact with members of two neighboring groups. Now, suppose that a random shock causes the puritan norm to become common in a single group. Will this norm spread? To answer this question we have simulated this model for a range of parameter values. Representative results are shown in Figure 5 which plots the ranges of parameters over which the beneficial norm spreads. The vertical axis gives the ratio of m, the probability that individuals interact with others outside of their group to a, rate of change due to imitation within groups, and the horizontal axis plots $\tilde{p}$, the unstable equilibrium that separates the domains of attraction of puritanical and tolerant equilibria in isolated groups. The shaded areas give the combinations of m/a and $\tilde{p}$ which lead to the spread of the puritanical norm to all groups given that it was initially common in a single group for two values of g.

First, notice that the beneficial norm spreads when the level of interaction between groups is intermediate. If there is too much mixing the puritanical norm cannot persist in the initial population. It is swamped by the flow of norms from its two tolerant neighbors. If there is too little mixing, the puritanical norm 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.

Second, to understand the affect of g, consider the case in which $g = 0$. Even when the norm produces no benefit to individuals when it becomes common, it still can spread if the puritanical norm has a larger basin of attraction in an isolated population ($d < 0$). To understand why, consider a focal group at the boundary between the spreading front of puritan groups and the existing population of tolerant groups. The focal group, in which both norms are present, is bounded on one side by a group in which puritan norms are common, and on the other side a group in which tolerant norms are common. Since groups on both sides of the boundary have the same average payoff, flow of norms will tend to move the focal group towards an even balance of the two norms. If the domain of attraction of the puritanical norm includes 0.5 and if there is enough mixing, then mixing with neighboring groups can be enough to tip the focal group into the basin of attraction of the puritanical norm. To see why increasing g increases the range of values of $\tilde{p}$ which allow the beneficial norm to spread, consider again a focal group on the boundary between the regions in which the puritanical norm is common and uncommon. When $g > 0$, individuals in the focal group are more likely to imitate someone from the neighboring group where the puritanical norm is common than the other neighboring group where tolerant individuals are common because individuals from the former group are more successful. Therefore, the flow of norms will tend to move the focal group toward a frequency greater than 0.5.

It is interesting that the rate at which this process of equilibrium selection goes on seems to be roughly comparable to the rate at which traits spread within a single group under the influence of the same learning process. 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 numbers of players (e.g. Samuelson, 1997). These processes also tend to pick out the equilibrium with the highest average payoff and the largest domain of attraction. However, unless the number of individuals in the model is very small, the rate at which this occurs is very slow. In contrast, in the simulations we performed the group beneficial trait spread from one population to the next at a rate roughly half the rate at which the same imitate-the-successful mechanism led to the spread of a trait within an isolated group. Of course, we have not accounted for the rate at which the beneficial norms becomes common in an initial group. This requires random processes. However, only the group, not the whole population, needs be small, and the group must be small only for a short period of time for random processes to give rise to an initial "group mutation" which can then spread relatively rapidly to the population as a whole.

CONCLUSION: ARE NORMS USUALLY SENSIBLE?

We have shown that it is possible for norms to guide people towards sensible behavior that they would not choose if left to their own devices. Norms could be sensible, just as functionalists in anthropology have claimed. However, the fact that they could be sensible doesn't mean that they *are* sensible. There are some well studied examples like the alkali treatment of corn which are well established, and there are many other plausible examples of culturally transmitted norms that seem to embody adaptive wisdom. However, as documented in Robert Edgerton's book, *Sick Societies* (1992), there are also many examples of norms that are not obviously adaptive and in fact some seem spectacularly maladaptive. More careful quantitative research on the costs and benefits of alternative norms would clearly be useful.

We believe that it is also important to focus more attention on the processes by which norms are shaped and transmitted. What happens to a Maya who doesn't utilize the normative form of alkali treatment of corn in her traditional society? What are the nutritional effects? The social effects? From whom do people learn how to process corn? How does this affect which variants of the process are transmitted and which are not? Only by answering such questions will we learn why societies have the norms they have and when norms are adaptive.

REFERENCES

Binmore, K. G. 1998. Just Playing: Game Theory and the Social Contract. Cambridge, MA: MIT Press.

Bowles, S. and Gintis, H. in press. The evolution of Strong Reciprocity. Manuscript.

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

Boyd, R. and P.J. Richerson. 1988. An evolutionary model of social learning: the effects of spatial and temporal variation. In: Social Learning: A Psychological and Biological Approaches, eds T. Zentall and B. G. Galef, pp 29–48. Hillsdale, NJ: Lawrence Erlbaum Associates.

Boyd, R. and P.J. Richerson, 1989. Social Learning as an Adaptation. *Lectures on Mathematics in the Life Sciences*, **20**:1–26,

Boyd, R., and P.J. Richerson. 1995.Why does culture increase human adaptability? *Ethnology and Sociobiology*. **16**: 125–143.

Boyd R., and P. J. Richerson. 1996.Why culture is common, but cultural evolution is rare, *Proceedings of the British Academy*, **88**: 77–93.

Edgerton, R. 1992. Sick Societies: Challenging the Myth of Primitive Harmony. NY: Free Press.

Ellickson, R. C. 1994. Order Without Law: How Neighbors Settle Disputes, Cambridge, MA: Harvard University Press.

Harris, M. 1979. Cultural Materialism: The Struggle for a Science of Culture. New York: Random House.

Henrich, J., and R. Boyd. 1998. The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior* **19**: 215–242.

Jensen, G. F. and M. L. Erickson. 1979. The religious factor and deliniquency: Another look at the Hellfire hypothesis. In: The Religious Dimension: New Directions in Quantitative Research. Ed. R Wenthow, NY: Academic Press. Pp. 157–177.

Katz, S. H., M. L. Hediger, and L. A. Valleroy. 1974. Traditional maize processing techniques in the New World. *Science*. **184**: 765–476.

McAdams, R. H. in press. The origin, development, and regulation of norms. Michigan Law Review, **96:** 338**.**

McEvoy, L. and G. Land. 1981. Life-style and death patterns of Missouri RLDS Church members. *American Journal of Public Health.* **71:** 1350–1357.

Ostrom, E. 1991. Governing the Commons: The Evolution of Institutions for Collective Action. Cambridge: Cambridge University Press.

Richerson, P. J. and R. Boyd. 1998. The evolution of human ultra-sociality. In: Indoctrinability, Ideology, and Warfare: Evolutionary Perspectives, eds. I. Eibl-Eibisfeldt and F. Salter,. NY: Berghahn Books, Pp. 71–96.

Samuelson, L. 1997. Evolutionary Games and Equilibrium Selection. Cambridge, MA: MIT Press.

Turner, J. H. and A. Maryanski. 1979. Functionalism. Menlo Park CA: Benjamin/Cummings.