working paper
department
of economics

RULES OF THUMB FOR SOCIAL LEARNING

Glenn Ellison

Drew Fudenberg

No. 92-12

June 1992

massachusetts
institute of
technology

50 memorial drive
cambridge, mass. 02139

RULES OF THUMB
FOR SOCIAL LEARNING

Glenn Ellison

Drew Fudenberg

No. 92-12

June 1992

Rules of Thumb for Social Learning [1]

Glenn Ellison[2] Drew Fudenberg[3]

June 18, 1992

[1]We thank Ahbijit Banerjee, Mathias Dewatripont, and seminar participants at University College, London, Nuffield College, Oxford, GREMAQ, Toulouse, and the Université Libre de Bruxelles for helpful comments. We are also happy to acknowledge financial support from a National Science Foundation Graduate Student Fellowship, from NSF grant SES 90-08770, the John Simon Guggenheim Foundation, and Sloan Foundation Graduate Program.

[2]Harvard University, Department of Economics

[3]MIT, Department of Economics, and IDEI, Toulouse.

ABSTRACT

This paper considers agents who use the experiences of their neighbors in deciding which of two technologies to use. We consider two learning environments, one where the same technology is optimal for all players and the other where each technology is better for some of them. In both environments, we suppose that players use exogenously specified rules of thumb that ignore all historical data but which may incorporate a tendency to use the more popular technology. These naive rules can lead to fairly efficient decisions in the long run, but adjustment can be quite slow when a superior technology is first introduced.

JEL Classifications: D8, C7, N53, O30

1. Introduction

This paper presents two simple models of how economic agents decide which of two technolgies to use when the relative profitability of the technologies is unknown. In both models, agents base their decisions, at least in part, on the experience of their neighbors; this is what we mean by "social learning." We believe that social learning is frequently an important aspect of the process of technology adoption, where "technology" should be broadly construed: Although our main example concerns the adoption of agricultural technology in the English agricultural revolution, we believe that the models may also be applicable to such choices as parents' decisions whether to send their children to a public or private school.[1]

There have been several previous models of the role of social learning in technology adoption. Perhaps the earliest is the contagion process, which models adoption as a random matching process in which players switch to the new technology the first time they meet someone who is using it; this process yields the familiar "S-shaped curve" for the time path of adoption that has been widely used in empirical work.

Recent papers by Banerjee [1991a], [1991b], Bikchandari et al [1991] and Smith [1990] study more sophisticated models of social learning, in which players must decide which of two choices is better. The primary question of interest in these models is whether the social learning is sufficient to prevent the population from locking on to the wrong decision. These papers suppose that players observe one another's choices, but that players do not observe the payoffs that these choices generate.[2] In a model where players choose sequentially, later decision makers may thus be led to make the same wrong choice as their predecessors, as the late deciders do not observe that the early ones now regret their choice. Moreover, this can occur even though the players could identify the optimal choice with certainty by pooling their information.

The learning environments we study differ from those of previous work in three ways. First, we believe that, in the context of technology adoption, it is more natural to suppose that players do observe their neighbors' payoffs, as well as their choices.[3] Second, our paper differs in supposing that individuals periodically observe other players' outcomes and reevaluate their own decision. Third, we consider the possibility that players may be sufficiently heterogeneous that under full information they would not all make the same choice.[4] In the diffusion of agricultural innovations, for example, different technologies may be appropriate for different soils and climates.

In addition to these differences in the learning environment, our paper also differs from those cited in the style of its analysis: Instead of assuming that the adoption process is described by the equilibrium of a game played by fully rational agents, we suppose that players use exogenously specified, and quite simple, "rules of thumb." We have several reasons for proceeding in this fashion. First, in some of the environments we consider, fully Bayesian learning requires calculations that may be too complicated to be realistic. A second motivation for our approach is that, to the extent that the technology choice may be substantially different than previous decisions the players have faced, we would be uncomfortable with the assumption that the technology adoption process is described by an equilibrium. A somewhat different motivation is simply technical expediency: we did not see an easy way to incorporate various considerations we feel are important into a rational-actor, equilibrium, model.

The paper is structured around two simple models of learning environments. The first one has a homogeneous population of players choosing between two competing technologies, with the payoff to each technology subject to an aggregate i.i.d. shock. Each period, only some fraction of the players has the opportunity to revise their choices; these players make their choices using simple rules of thumb.

Our analysis begins with a particularly "naive" rule of thumb in which players ignore all historical data and simply choose whichever technology worked better in the previous period. This rule will lead the popularity of the two technologies to fluctuate unless one of the technologies has a higher payoff for all values of the shock.

We subsequently consider rules which incorporate "popularity weighting," a tendency to choose a more popular technology even if it was somewhat less profitable last period. Since players observe both the choices and the payoffs of their neighbors, they would have no reason to use popularity weighting if they made full use of their information. However, we feel that real-world decision makers often do pay attention to popularity, and indeed our results may help provide some explanation for this phenomenon.

Specifically, in our model the appropriate use of popularity weighting leads players to adopt and stick with the better technology. Intuitively, a strategy which is more popular today is likely to have done well in the past, so that the relative popularity of the technologies can serve as a proxy for their historical performance. Thus is fairly clear that popularity weighting rules can lead to better decisions; we find that one particular choice of popularity weights picks out the better technology in the long run. This leads us to ask whether there are any reasons to believe that this optimal popularity rule is either particularly likely or particularly unlikely to be used. In response, we present a model of players choosing popularity rules in which the optimal rule is the unique symmetric equilibrium.

Our second model has a heterogeneous population, with each technology better for some of the players. Thus the question here is not whether the better technology will be adopted, but rather whether the new technology will be adopted by the appropriate players. We suppose that there is a continuum of players distributed uniformly over a line, and that nearby players have similar payoffs to the two technologies.[5] Moreover, we suppose that players base their decisions on the relative performance of the two technologies at locations that are within one "window width" of their own. This window width, which is exogenous in our model, can be thought of as either the result of an informational constraint- players may not observe outcomes at far-away locations- or as the result of a prior belief that players at far-away locations are not sufficiently similar for their experiences to be informative.

Once again, players revise their technology choices using simple rules of thumb. In particular, we suppose that players do not know exactly how location influences relative payoffs, and thus simply compare the average payoffs of the two technologies in their window, as opposed to using more sophisticated statistical methods.

This second model provides a number of predications about the types and magnitudes of the errors that are likely to be made. The spatial nature of the process allows some degree of social learning even without popularity weighting, and the long-run state of the system is approximately efficient when the window width is small. However, small window widths imply that the system converges more slowly, which can be costly if the initial state is far from the optimum. Roughly speaking, increasing the popularity weighting in the spatial model has about the same effect as decreasing the window width.

Before proceeding further, we should acknowledge that our belief that models of bounded rationality are a useful way to study social learning does not mean that we are completely satisfied with the particular rules we consider. In particular, in the first model use of history does not seem so complicated as to be unreasonable.[6] Our purpose is not to argue that any one of these models is particularly compelling, but rather to identify general properties that seem to occur in some of the more obvious formulations. One recurrent conclusion is that in a number of cases the long-run state of the system is fairly efficient, even though the individual decision rules are quite naive.

2. The English Agricultural Revolution

Before developing our models, we would like to use the example of the English agricultural revolution to introduce some of the issues that our models are designed to address. The revolution in question is the improvement in English agriculture between 1650 and 1850. This improvement is usually attributed to the spread of new agricultural practices, and in particular to what is called the "new husbandry," although both the size of the improvement and its causes have become more controversial in recent years. The new husbandry refers to a variety of new crops and new crop rotations which arrived in England from Flanders in the 17th century, based on the idea of growing crops such as clover or turnips instead of leaving the land fallow. These crops could rejuvenate the soil, and the hoeing they required had the by-product of eliminating weeds. The crops themselves can be fed to livestock, which allows larger herds to be carried over the winter, providing further supplies of manure, which in turn could be used for fertilizer. The net result is (ideally) increased production of both livestock and grains.[7]

The diffusion of the new husbandry seems to have involved all of the major aspects of our models that we mentioned in the introduction. First, it seems likely that farmers were able to observe, at least roughly, the output of their neighbors, as well as their neighbors' choice of crops and crop rotations. Second, the payoffs to various crops were different at different locations, depending on the soil, climate, and terrain of each farm. There is not yet a consensus on where the new husbandry should have been adopted, but it is clear that it was not a universal improvement: Turnips were most suited to the light clay soils of Norfolk, and were unprofitable in wet clay soils like those of the Midland Plain, where the crop had limited growth, and was so difficult to harvest it was often left to rot in the field[8]. Third, a crop of turnips could be ruined by excessive rain or severe frosts, leading us to incorporate the weather as an annual stochastic shock complicating the learning process.[9]

Moving away from the physical description of the situation to our (harder to verify) assumptions about the agents' behavior, it seems clear that the final adoption decisions resulted from decentralized learning, as opposed to a pronouncement from a central authority (although there were attempts made along this line, as we discuss below.) And it seems plausible that the farmers may have been less sophisticated in their use of past observations than Bayesian learning would suggest. Moreover, with capital markets poorly developed or nonexistent, and starvation a potential concern, it seems plausible that farmers' technology decisions were determined primarily by short-term considerations, and that farmers would be unlikely to experiment with a technology with a lower expected return.

The two models we discuss explore two different aspects of the adoption process. The homogeneous-population model looks at a single location in isolation, and focuses on the dynamics of the adoption process. It has been frequently noted that farmers as a group are seemingly very hesitant to try new technologies. These comments do not suggest that all farmers are equally hesitant; for example, Slicher von Bath (*op. cit.*, p. 243) notes that during the English agricultural revolution, "Land tilled in very ancient ways lay next to fields in which crop rotations were followed." This observation fits with our assumption of *inertia*, meaning that at each date only some fraction of the population considers changing technologies.

The apparent inertia in the diffusion of the new husbandry has been criticized by both contemporary and modern authors as slowing progress, and indeed it does slow the transition from a dominated technology (one that is worse in all states of the world) to a new one. However, our analysis shows that inertia may improve the long-run performance of the process if the performance of the two technologies is subject to sufficiently large random shocks. Given our assumption that players do not keep track of past outcomes, the process without inertia oscillates between the two technologies, while a combination of inertia with a tendency to use the most popular technique permits the learning process to converge to the better of the two. Our second model examines the idea that players learn from their "neighbors" when a new technology may turn out to be profitable for some but not all of the potential adopters. (In the case of farming, we take this spatial structure literally, but we also have in mind learning from "neighbors" who are believed to be similar, but who need not be geographically adjacent.) One of the most striking and frequently noted characteristics of the agricultural revolution is the slow rate at which the innovations spread. Contemporary observers in England said that the rate was only one mile per year, and indeed it did take more than a century for the new husbandry to make its way across the island[10].

Many explanations have been proposed for this slow spread, including technological factors like pests and diseases, institutional factors such as the enclosures, and the farmers' lack of education. While all of these factors may have played a role, our analysis suggests that the basic fact of a slow rate of diffusion need not be surprising, as it is can be a natural consequence of social learning, particularly when the difference in payoffs is not great, and when farmers pay attention to the relative popularity of each technology in making their decisions.

A second and more spirited debate has surrounded the role of elite landlords and agricultural reformers on the diffusion process. The classic studies of Ernle [1912] and Mantoux [date] portrayed the agricultural revolution as the result of valiant struggle by innovators such as Jethro Tull and Lord Townshend to overcome the ignorance of the peasant farmers. Revisionist authors have attacked this view; a main component of their argument has been that the new techniques were already in use in some areas of England before many of the so-called innovators were born.[11]

Our model reinforces the pro-elitist side of this debate by emphasizing that popularizers can have a significant impact in promoting the diffusion of a technique even if they are not the first to develop it. The arrival of the new husbandry in England is often dated to the publication in 1650 of Sir Richard Weston's observations of agricultural practices in Brabant and Flanders. When one examines the spread of the new husbandry as detailed by Kerridge [*op. cit.*, pp. 272-279], it appears that the new husbandry spread out very slowly from a number of geographically distinct locations. The first adoptions were in Suffolk, but by 1680 the new husbandry had appeared in several other locations over a hundred miles away. By providing information from other counties, the agricultural popularizers may have promoted such subsequent innovations, and sped up the rate of the technology's diffusion.

In addition, our theoretical model may shed some light on another question which has received less attention in the literature: Was the new husbandry eventually adopted in all of the areas to which it was suited? As a potential guide to historical research, we discuss the conditions under which naive learning processes of the kind we consider tend to generate efficient long-run outcomes.

Beyond these historical questions, our spatial model of learning raises the basic question of how well "naive" learning rules perform that we investigated in our first, homogeneous-population, model. Our answers here focus on the tradeoff between rates of adoption and efficiency of the long-run equilibrium.

3. A Simple Model of Homogeneous Populations

Before considering social learning in systems with a heterogeneous population, it is interesting to consider the simpler case in which the same technology is optimal for all players. This model can be thought of as describing behavior at a single site in the model we consider later on, where the relative payoffs vary with location. Suppose that there is a large (continuum) population of players at a single site, each of whom must choose whether to use technology f or technology g.

In each period, all players using the same technology receive the same payoff. (Given our assumption that players observe one anothers' payoffs, nothing would be changed if we allowed each player's payoff to be subject to idiosyncratic shocks.) We suppose that the payoffs to the two technologies at date t, u_t^f and u_t^g, are related by the equation

$$(1) \quad u_t^g - u_t^f = \theta + \varepsilon_t,$$

where θ is a fixed but unknown constant parameter and the ε_t are i.i.d. shocks with zero mean and cumulative distribution function H. (In later sections the constant θ will vary with location.) We will assume that $p \equiv 1 - H(-\theta) = \text{Prob}\,[u_t^g - u_t^f \geq 0]$ is strictly between 0 and 1.

In the initial period, denoted 0, a fraction x_0 of the players are using technology g. After each period, a fraction α of the players have the opportunity to revise their choice. Very low values of α might correspond to a system in which individual player made their choices for the duration of their effective lifetimes, with the revisions corresponding to an inflow of replacement players; intermediate values might describe a system in which the choice of a technology is embodies in a costly capital good that will not be replaced until it wears out.[12]

We suppose that the players who are revising their choice can observe the average payoffs of both technologies in the previous period. However, players do not have access to the entire history of payoff observations. To justify this assumption, we suppose that individual players revise their choices too infrequently to want to keep track of each period's results, and more strongly that the market at this particular "location" is too small for a record-keeping agency to provide this service.

The simplest behavior rule we consider is the "unweighted" rule under which all players who revise their choice pick the technology which did best in the preceding period. Under this adjustment rule, the evolution of the system is described by

$$(2) \quad x_{t+1} = \begin{cases} (1-\alpha)x_t + \alpha & \text{with probability } p = \text{Prob}[u_t^g \geq u_t^f], \\ (1-\alpha)x_t & \text{with probability}(1-p) = \text{Prob}[u_t^g < u_t^f], \end{cases}$$

so that

$$(2') \qquad E(x_{t+1} | x_t) = (1-\alpha)x_t + \alpha p.$$

Note that this specification is symmetric in its treatment of the adoption and discontinuance decisions, which corresponds to the case where the costs of "transition" are small.[13] Our reading about the English agricultural revolution, as well as studies of more recent innovations cited in Rogers and Shoemaker (p.115) suggest that the amount of discontinuance is an important factor in the diffusion process.

The following result is standard; it follows from e.g. theorem 10 of Norman [1968]. (It is also a consequence of part (b) of proposition 2 below.)

<u>Proposition 1:</u> The system (2) is ergodic, i.e. the time-average of x_t converges to its expectation with respect to its unique invariant measure μ. Moreover, $E_\mu(x) = p$, and $\text{var}_\mu(x) = p(1-p)\alpha/(2-\alpha)$.

4. A Single Location with Popularity Weighting

Proposition 1 says that observing the long-run fraction of players using technology g reveals the fraction of the time that g has been the better choice. If the distribution H of ε is symmetric, the arm that is more often better is also the arm with the higher expected payoff. (The same conclusion holds so long as the amount of asymmetry in H is small compared to θ.) This suggests that if all other players in the population are choosing whichever technology has the highest current score, each player could gain by considering the relative popularity of the two technologies, as well as their recent payoffs.

Intuitively, the current popularity provides some information about the past history of the process, and thus can serve as a proxy for it. Although this information is not complete, the complete history of the process is not needed to identify the better technology. One way to interpret the results of this section is that in some cases popularity weighting is a good enough proxy for the history that the system eventually converges to the correct choice.

We now develop a simple parametric model of popularity weighting with a single location. As above, we suppose that only a fraction α of the population updates its choice each period. Now, though, instead of choosing the technology which did best last period, the choice rule is

(3) "Choose g if $u_t^g - u_t^f \geq m(1 - 2x_t)$."

Under this rule, the probability that those players who revise their choices choose g is $\text{Prob}[\ \theta+\varepsilon_t \geq m(1 - 2x_t)] = 1 - H(m(1-2x_t)-\theta)$; when all players use rule (3), the fraction using g evolves according to

$$(4)\quad x_{t+1} = \begin{cases} (1-\alpha)x_t + \alpha & \text{with probability } 1-H(m(1-2x_t)-\theta) \\ (1-\alpha)x_t & \text{with probability } H(m(1-2x_t)-\theta) \end{cases}.$$

The parameter m indexes the amount of popularity weighting; the case m = 0 corresponds to the unweighted case discussed above. When $x_t = 1/2$, both technologies are equally popular; in this case players chooses the technology with the highest current payoff for any value of m. As m grows, players become more willing to choose the currently popular technology even if its current payoff is lower[14].

We use the linear specification of popularity weighting primarily for analytic convenience. It combines nicely with a second simplifying assumption that we make in this section, that the distribution H of the per-period shocks ε_t is uniform on $[-\sigma, \sigma]$. This allows us to explicitly compute the long-run behavior of the system for any m. It also ensures that the linear class of weighting rules we consider includes one rule that leads the asymptotic distribution to concentrate on the optimal choice, namely $m = \sigma$.[15]

Beyond the presumed linearity in x_t, another point to note about decision rule (3) is that it, and any rule that compares the difference $u_t^g - u_t^f$ to a function of x_t, is invariant to additive transformations of the payoff function, but not to multiplicative ones: In order to preserve the same decision rule when the payoff functions are multiplied by a constant λ, the parameter m must be multiplied by the same constant. One way to see why this must be the case is to note that the expression (1-2x) is unitless, so the parameter m is measured in the same units as the payoff are.[16]

Our assumption that the per-period shocks have a uniform distribution makes it easy to determine the long-run behavior of the system. Since the lowest possible value of ε_t is $-\sigma$, the lowest possible observation of $u_t^g - u_t^f$ is $\theta - \sigma$. Hence, if x_t is sufficiently large that $\theta - \sigma \geq m(1-2x_t)$, or equivalently if $x_t \geq \underline{x}^g \equiv (m-\theta+\sigma)/2m$, the fraction using technology g is certain to increase. Likewise, if $x_t \leq \bar{x}^f \equiv (m-\theta-\sigma)/2m$, the fraction playing f is certain to increase. (Note that $\sigma > 0$ implies $\bar{x}^f < \underline{x}^g$.) Because the probability of a upwards step is minimized at $x_t = 0$, this probability must be at least $\text{Prob}[\theta+\varepsilon_t \geq m] = (\sigma-m+\theta)/2\sigma = -(m/\sigma)\underline{x}^f$. Thus, when $\bar{x}^f < 0$, so that the system cannot "lock on" to downwards steps, the probability of an upwards step is uniformly bounded away from zero. Similarly, if $\underline{x}^g > 1$, the probability of a downwards step is uniformly bounded away from 0.

The above shows that (ignoring knife-edge cases) there are four possibilities for the long-run behavior of the system: If $\underline{x}^g < 1$ and $\bar{x}^f < 0$, the system is certain to eventually make enough upward jumps that $x_t > \underline{x}^g$, so that from any initial position the system converges with probability 1 to $x_t = 1$. If $\underline{x}^g > 1$ and $\bar{x}^f > 0$, the system converges to $x_t = 0$ from any initial position. If $0 < \bar{x}^f$ and $\underline{x}^g < 1$, the system will converge (with probability 1) to 0 if $x_0 \leq \bar{x}^f$, will converge to 1 if $x_0 \geq \underline{x}^g$; for $x_0 \in (\bar{x}^f, \underline{x}^g)$, the system will also eventually converge to a steady state, but it has a positive probability of ending up at each of the two steady states of the system. In the remaining case, in which $\bar{x}^f < 0$ and $\underline{x}^g > 1$, the system will not converge to either steady state. Instead, the fraction x_t will continue to fluctuate, with the long-run distribution computed following the statement of proposition 2 below..

The above observations do most of the work required to establish the following claims:

Proposition 2:

(a) Popularity weighting $m = \sigma$ is "optimal" in the sense that from any x_0 the system converges with probability 1 to the state where everyone uses the better technology.

(b) $m > \sigma$ is "overweighting", in that the system converges with probability 1 to a steady state, but which steady state is selected may depend on the initial condition x_0. More precisely, the system converges to the better technology if $|\theta| \geq m-\sigma$, while for $|\theta| < m-\sigma$ the behavior of the system depends on the initial condition x_0. If $x_0 \geq (m+\sigma-\theta)/2m$, the system converges to 1 with probability 1; if $x_0 \leq (m-\sigma-\theta)/2m$, the system converges to 0 with probability 1. If $|\theta| < m-\sigma$ and $x_0 \in ((m-\sigma-\theta)/2m, (m+\sigma-\theta)/2m)$, the system will eventually converge to one of the steady states, but both steady states have positive probability.

(c) With "underweighting," i.e. $m < \sigma$, the system need not converge to a steady state. It does converge (with probability 1) to the better technology if $|\theta| \geq \sigma-m$, but for $|\theta| < \sigma-m$, the system has a non-degenerate invariant distribution μ, with

$$E_\mu x = 1/2 + \theta/2(\sigma-m), \text{ and}$$

$$\text{var}_\mu x = \alpha\sigma E_\mu x E_\mu(1-x)/[(2-\alpha)\sigma-2(1-\alpha)m].$$

Proof:

(a) If $m = \sigma$, then $\underline{x}^g = (2m-\theta)/2m$ is less than 1 iff $\theta > 0$, and $\bar{x}^f = -\theta/2m$ is greater than zero iff $\theta < 0$. The conclusion now follows from the argument in the text.

(b) It suffices to check that if $\theta > m-\sigma > 0$ then then $\bar{x}^f < 0$ and $\underline{x}^g < 1$, that $-\theta > m-\sigma > 0$ implies $\bar{x}^f > 0$ and $\underline{x}^g > 1$, while for $m-\sigma > |\theta|$, $\bar{x}^f > 0$ and $\underline{x}^g < 1$.

(c) A similar computation shows that when $|\theta| > \sigma-m$, the system must converge. Appendix B establishes that the system has a unique invariant distribution when $\sigma-m > |\theta|$, and computes the corresponding mean and variance. ∎

Proposition 2 shows that the system is certain to converge to the correct choice if the popularity weight m equals σ, and that the payoff loss from a wrong choice must be small if m is close to this level. Thus it is interesting to ask whether there is any particular reason to suppose that popularity weights equal or close to σ are likely to be used, or conversely whether there are forces in the model that would drive the players to use different weights. As a partial response, we consider a game in which players simultaneously choose their individual popularity weights, and show that the optimal weight $m = \sigma$ is its unique equilibrium outcome. This result is only a partial response, because it supposes more sophistication in the determination of the popularity weights than we find compelling. However, the result does show that popularity weighting need not conflict with individual incentives.

To define the payoffs in this game, we suppose that players have a common prior distribution ρ over θ, and that ρ assigns positive probability to every neighborhood of $\theta = 0$. For each $x_0 \in [0,1]$, $\theta \in \text{support}(\rho)$, and m, let $\mu(\theta, x_0, m)$ be the long-run distribution on x when all players use weighting m. For each value of θ, the payoff to the profile m is defined to be the mean of the long-distribution on payoffs, and the overall payoff is the expectation of this value with respect to the prior beliefs ρ. Our use of the long-run payoff criterion here is made solely for convenience: we would prefer to consider optimal behavior for discount factors near 0, but then we would want to allow for the popularity weightings to be chosen repeatedly (popularity has no relevance in the first period) and would need to consider the distribution of the state in each period separately.

Since the profile in which all players choose $m = \sigma$ results in the full-information payoffs, this profile is clearly a Nash equilibrium of the game. Moreover, it is the only symmetric equilibrium, as shown in the following proposition.

Proposition 3: If every neighborhood of $\theta = 0$ has positive probability, the unique symmetric equilibrium of the game in which players simultaneously choose the weight m they give to popularity is for all players to choose $m = \sigma$.

Remarks: (1) We do not know whether there are asymmetric equilibria as well. The long-run behavior of the system when two or more decision rules are used by a non-negligible proportion of the population seems difficult to determine. We should also point out that no symmetric pure-strategy Nash equilibrium exists in the case of a normal distribution that we consider in appendix A. This should not be too surprising, since in that case popularity weighting permits the full-information payoff to be approximated, but not to be attained exactly. However, profiles where all players use a "large" amount of popularity weighting are ε-Nash equilibria.

(2) If players are certain that the absolute value of θ is bounded away from zero, there are equilibria in which m can

exceed θ.

(3) It is difficult to suppose that players using the kinds of naive learning rules we consider would consciously choose their popularity weights to maximize their long-run payoff. We prefer to interpret the equilibrium assumption here as the result of a long-run adaptive process, but this raises the question of the relative speeds of adjustment of the process determining m and that reflecting learning about the technologies.

<u>Proof:</u> Fix any profile where all players use some $m \neq \sigma$. The idea of the proof is simply that if $m < \sigma$, so everyone is using too little popularity weighting, then each individual player would prefer to deviate and give more weight to popularity, while if $m > \sigma$, each player would prefer to give popularity use a bit less weight.

Since there is a "large number" of players, the aggregate behavior of the system is unaffected if any single player deviates. We will show that there is always a deviation that improves the player's payoff when $|\theta|$ is sufficiently small, and has no effect on the player's payoff when $|\theta|$ is larger; the conclusion will then follow from our assumption that every neighborhood of $\theta = 0$ has positive probability.

(a) Suppose first that $m < \sigma$, and consider a player deviating to $m' = m+dm$ for some small $dm > 0$. This deviation has no effect on his long-run payoff if $|\theta| \geq \sigma-m$, for in this case the payoff difference between the two technologies is sufficiently strong that x_t converges to the optimal choice, and any $m' \geq m$ yields the first-best long-run payoff.

If $|\theta| < \sigma-m$, the system does not converge to a steady state, but as a non-degenerate ergodic distribution. In this case, deviating to m' leads the player to use g instead of f whenever

$$m'(1-2x_t) \le u_t^g - u_t^f < m(1-2x_t), \text{ or}$$

$$m'(1-2x_t) - \theta \le \varepsilon_t < m(1-2x_t) - \theta.$$

Similarly, the player will now use f instead of g whenever

$$m(1-2x_t) \le \varepsilon_t < m'(1-2x_t) - \theta.$$

Since the difference in payoff between g and f is θ, the expected change in the player's per-period payoff is

$$du_i/dm = \theta\int(2x-1)dH\Big(m(1-2x)-\theta\Big)d\mu(x),$$

where H is the uniform distribution, and μ is the invariant distribution on that was derived in proposition 2. Since $m(1-2x) - \theta \in [-\sigma,\sigma]$ for all $x \in [0,1]$, $dH\Big(m(1-2x)-\theta\Big) = 1/2\sigma$, and so

$$du_i/dm = \theta/2\sigma \int(2x-1)d\mu(x) = \theta/2\sigma\ [E_\mu 2x-1].$$

Since $E_\mu x > 1/2$ for $\theta > 0$, and $E_\mu x < 1/2$ for $\theta < 0$, $du_i/dm > 0$ for all $\theta \in (-(\sigma-m),0) \cup (0, \sigma-m)$.

(b) Now consider a profile in which players use an $m > \sigma$, and consider an individual player deviating to $m' = \sigma$. Suppose first that $\theta > 0$, so that technology g has the higher payoff. If $\theta \ge m-\sigma$, the system converges to 1 with probability 1, so that both m and m' yield the full-information long-run payoff; the same is true if $\theta \le m-\sigma$ and $x_0 \ge \underline{x}^g = (m+\sigma-\theta)/2m$.

If $0 < \theta < m-\sigma$ and $x_0 < \underline{x}^g$, or equivalently $\theta < (1-2x_0)m+\sigma$, there is positive probability that x_t converges to 0. In particular, if $\theta < \sigma$, there is positive probability of converging to 0 form any $x_0 \le 1/2$. If the system does converge to 0, using weighting m leads to f being played in almost every period, while using $m' = \sigma$, the probability of playing g converges to the probability that $\varepsilon_t \ge m-\theta = (\sigma+m+\theta)/2m$, which is greater than 0 for all positive θ.

The preceding two paragraphs show that m' does at least as well as m for all positive θ, and does strictly better if $x_0 \leq 1/2$ and $0< \theta< \min(\sigma, m-\sigma)$. A symmetric argument shows that m' does at least as well as m for all negative θ, and does strictly better if $\theta^* + z\, x_0 \geq 1/2$ and $0> \theta> \max(-\sigma, -(m-\sigma))$. Hence if the prior assigns positive probability to the neighborhood of $\theta = 0$, m' yields a strictly higher expected payoff from any initial position x_0. ■

While our formal results concern the eventual steady state of the system, the speed of convergence is of some interest as well. In particular, consider an initial position where x_0 is small, so that g corresponds to a "new" technology, and suppose that $\theta > 0$, so that the new technology is in fact an improvement. Then the share of technology g increases whenever $\theta+\varepsilon_t > m(1-2x_t)$, and since the probability of this event increases with θ, so does the expected rate of adoption.[17] Such a correlation between the extent of improvement and the speed of adoption has been noted in the empirical discussions of in Mansfield [1968] and Rogers and Shoemaker [1971], but has not, so far as we know, been addressed in the learning literature.[18]

Note also that for fixed θ, the speed of convergence decreases as σ increases, so that each period's observation becomes less informative. More generally, convergence will be slow if the new technology usually does about as well as the old one, but occasionally does much better. Furthermore, if the new technology usually does slightly <u>worse</u> than the old one, but occasionally does much better, (i.e. if the new technology has a higher mean payoff but a lower median) then naive learning rules that look only at the recent relative performance will be biased towards the wrong choice. This is consistent with the observation that seat belts, insurance, and vaccinations have been slow to diffuse.

5. Heterogeneous Population with Linear Technologies

Now we turn to the study of heterogeneous populations, in which different technologies may be optimal for different individuals. As before, we suppose that there are only two technologies, denoted f and g, with the mean difference in payoffs, $E(u_t^g - u_t^f)$, equal to θ. Now though, we think of θ as representing a location along a line, so that players at different locations have different θ's. In particular, the optimal rule (both socially and privately) is for players with positive θ to use g, and players with negative θ to use f, so that the distribution of technology choice has a cut-off or break-point at $\theta = 0$.

It will be important in the following that the relative advantage of using technology g at location θ may be correlated with the "absolute advantage" of location θ, e.g. the productivity of the "land." To capture this, we suppose that the payoffs to the technologies have the following linear form:

$$(5) \quad \begin{cases} u_t^g(\theta) = \theta + \beta\theta + \varepsilon_{1t} \\ u_t^f(\theta) = \beta\theta + \varepsilon_{2t}. \end{cases}$$

With this parameterization, $\beta > 0$ implies that technology g does better at "good" locations, while when $\beta < 0$, g does better at bad ones.

In this model, the player's location in θ-space determines his average payoff to the two technologies. We want to think of the payoff-relevant variables as being unobservable but correlated with the observed locations. The idea is that players do not know exactly which aspects of their locations are payoff-relevant, or how these aspects influence their payoffs. For this reason, we do not allow the players to regress the observed payoffs of each technology on the corresponding values of θ. Instead, we suppose that players base their decisions on the average performance of the two technologies at locations in their "observation windows," where the observation window of the player at θ is the interval $[\theta-w, \theta+w]$. We call w the "window width."

We have two interpretations in mind for this model. First, the location parameter θ may correspond to geographical location, with the performance of the technologies linked to variables such as climate or terrain that are in turn correlated with location. Second, the model may describe adoption decisions at a single village, where players are differentiated by idiosyncratic payoff-relevant characteristics such as wealth and household size.

In studying geographic diffusion, for example of an agricultural technology, the observation window might reflect the farmer only observing the outputs of his neighbors, and the window width w might be fairly small. In studying adoption at a single site, the observation window corresponds to the players' beliefs about which other players are sufficiently similar for their experiences to be relevant, and players might well observe the actions and outcome of others who are outside of their window. To the extent that the relevant characteristics are difficult to determine, the window widths in this interpretation might be fairly large. In both interpretations, players might prefer to weight observations of their immediate neighbors more heavily than those of players who are farther away, but still within the observation window; this may be particularly attractive when the observation window is large. As in the study of a homogeneous population, we begin by analyzing the simple rule where players use whichever technology did better in their window last period; later we will enrich the model to allow for popularity weighting. To define this rule formally, suppose that the distribution of players over locations has a constant density, which we normalize to equal 1, and let $\bar{u}_t^g(\theta)$ be the average score realized by those players in the interval $[\theta-w, \theta+w]$ who used g at period t, with the convention that $\bar{u}_t^g(\theta) = -\infty$ if every player in the interval used f; the average $\bar{u}_t^f(\theta)$ is defined analogously.

The (unweighted) decision rule is for the player at θ is then

(6) "Play g at period t+1 iff $\bar{u}_t^g(\theta) - \bar{u}_t^f(\theta) \geq 0$. "

In the previous sections we considered a model with a continuum of players and inertia, so that the fraction of players using each strategy can never shrink all the way to zero in finite time. In our study of spatial models, though, we will suppose that there is no inertia at individual locations, so that <u>all</u> players at each location revise their choices each period. We do so in part for reasons of convenience, and in part because in rural areas with low population density it seems plausible that a technology could be abandoned by everyone in an observation window after a few bad draws in a row.

As a first step in analyzing the decision rule (6), suppose that the noise terms ε_{1t} and ε_{2t} are identically zero, so that the system is deterministic. Suppose further that the initial state of the system is described by a <u>cut-off</u> rule. That is, suppose that there is a $\hat{\theta}_t$ such that all players with $\theta \geq \hat{\theta}_t$ choose g and all those with $\theta < \hat{\theta}_t$ choose f. Then the period t+1 state will be described by a cut-off rule as well. To see this, note that all players at $\theta > \hat{\theta}_t + w$ see only g being played, and hence will play g in the next period, while all players at $\theta < \hat{\theta}_t - w$ play f. Players at every $\theta \in [\hat{\theta}_t - w, \hat{\theta}_t + w]$ see both f and g being played, with

$$(7)\quad \begin{aligned} \bar{u}_t^g(\theta) &= \int_{\hat{\theta}}^{\theta+w} (\beta+1)\,s\,ds/(\theta+w-\hat{\theta}) = (\beta+1)(\theta+w+\hat{\theta})/2, \text{ and} \\ \bar{u}_t^f(\theta) &= \int_{\theta-w}^{\hat{\theta}} \beta\,s\,ds/(\theta+w-\hat{\theta}) = \beta(\theta-w+\hat{\theta})/2. \end{aligned}$$

Thus for $\hat{\theta}_t - w < \theta' < \theta'' < \hat{\theta}_t + w$, we have

$$\bar{u}_t^g(\theta'') - \bar{u}_t^f(\theta'') = \bar{u}_t^g(\theta') - \bar{u}_t^f(\theta') + (\theta''-\theta')/2,$$

so that if the player at θ' plays g in period t+1 then so does the player at θ''. Hence the state at period t+1 is described by a cut-off rule.

A steady state cut-off rule must have the property that the player at the steady- state cut-off is indifferent between f and g given his observations. Thus the steady state is the unique solution of

$$\bar{u}_t^g(\theta^*) = \bar{u}_t^f(\theta^*), \text{ that } (\beta+1)(\theta^* + w/2) = \beta(\theta^* - w/2),$$

and so

$$(8)\qquad \theta^* = -(2\beta+1)w/2.$$

Note that although the optimal cut-off is $\hat{\theta} = 0$ for any value of β, the steady state cut-off is only at 0 if $\beta = -1/2$. When $\beta = 0$, for example, the payoff to f is identically zero, while the payoff to g is equal to θ. Hence when the cut-off is at 0, the average payoff to players using g is strictly positive, which will tempt players to the left of 0 to adopt g as well. The discrepancy between the steady state and the optimum arises from our assumption that players do not directly observe θ, and hence use only the average payoffs received by the two technologies in making their decisions.Note that the maximum steady-state payoff loss at any location is the absolute value of θ^*, which is small if β is not too large (in absolute value) and the window width w is small.

Having determined the steady- state cut-off, we next examine the behavior of the system away from the steady state. It is easy to show that, from an initial cut-off $\hat{\theta}_0$, the cut-off will move towards the steady state θ^* at a distance of w each period until it is within w/2 of θ^*. Once $\hat{\theta}_t$ is within this interval the system typically enters a stable 2-period cycle about θ^*. For ease of reference, we summarize this as a proposition.

<u>Proposition</u> <u>4:</u> From an initial cut-off $\hat{\theta}_0$, the system determined by (6) and (7) evolves according to

$$(9)\quad \hat{\theta}_t = \begin{cases} \hat{\theta}_t + w & \hat{\theta}_t < \theta^* - w/2. \\ -\hat{\theta}_t + 2\theta^* & \hat{\theta}_t \in [\theta^* - w/2, \theta^* + w/2) \\ \hat{\theta}_t - w & \hat{\theta}_t \geq \theta^* + w/2. \end{cases}$$

<u>Proof:</u> If $\bar{u}^g_t(\hat{\theta}_t - w) - \bar{u}^f_t(\hat{\theta}_t - w) > 0$, then all players who observe both technologies being played -i.e. all players in the interval

$[\hat{\theta}_t - w, \hat{\theta}_t + w]$ - use g in period t+1. Substituting $\theta = \hat{\theta} - w$ into equation (7), we see that this is the case if $(\beta+1)\hat{\theta}_t \geq \beta(\hat{\theta}_t - w)$, or $\hat{\theta}_t \geq -\beta w = \theta^* + w/2$. Similarly, if $\hat{\theta}_t < \theta^* - w/2$, all players who see both technologies being played choose f in period t+1. Finally, if $\hat{\theta}_t \in [\theta^* - w/2,\ \theta^* + w/2)$, $\hat{\theta}_{t+1}$ will satisfy $(\beta+1)(\hat{\theta}_{t+1} + \hat{\theta}_t + w) = \beta(\hat{\theta}_{t+1} + \hat{\theta}_t - w)$, so that $\hat{\theta}_{t+1} = -\hat{\theta}_t - (2\beta+1)w = -\hat{\theta}_t + 2\theta^*$.■

Next we consider the behavior of the model with noise, i.e. with ε_{1t} and ε_{2t} non-degenerate i.i.d. random variables. Let $z_t = \varepsilon_{2t} - \varepsilon_{1t}$ denote the difference in the two shocks, and let $\theta^*_t = \theta^* + z_t$; θ^*_t is the steady state of the system when $\varepsilon_{2\tau} - \varepsilon_{1\tau}$ is identically equal to z_t for all τ. Because behavior rule (6) depends only on the difference between the payoffs to f and g, and not on their levels, the evolution of the system from $\hat{\theta}_t$ when the shock is z_t is the same as that given in equation (9), with the term θ^* replaced everywhere by θ^*_t.

<u>Proposition 5:</u> If the period-t cut-off is $\hat{\theta}_t$, and the period-t shock is z_t, the period t+1 cut-off is given by

$$(10)\quad \hat{\theta}_{t+1} = \begin{cases} \hat{\theta}_t + w & \hat{\theta}_t < \theta^*_t - w/2. \\ -\hat{\theta}_t + 2\theta^*_t & \hat{\theta}_t \in [\theta^*_t - w/2, \theta^*_t + w/2) \\ \hat{\theta}_t - w & \hat{\theta}_t \geq \theta^*_t + w/2. \end{cases}$$

<u>Proof:</u> For locations $\theta \in [\hat{\theta}_t - w, \hat{\theta}_t + w]$, the difference between the average payoffs of the two technologies in θ's observation window (the interval $[\theta_t - w,\ \theta_t + w]$), i.e. $\bar{u}^g_t(\theta, \varepsilon_{1t}) - \bar{u}^f_t(\theta, \varepsilon_{2t})$, is $[\theta + \hat{\theta}_t + (2\beta+1)w]/2 - z_t = (\theta + \hat{\theta}_t - 2\theta^*_t)/2$. Since $\hat{\theta}_t > \theta^*_t + w/2$ implies $\theta + \hat{\theta}_t \geq 2\theta^*_t$ for all $\theta \geq \hat{\theta}_t - w$, $\hat{\theta}_t > \theta^*_t + w/2$ implies that all players who observe both technologies choose g.

Similarly, $\hat{\theta}_t < \theta_t^* + w/2$ implies that all players who see both technologies choose f. Finally, if $\hat{\theta}_t \in [\theta_t^* - w/2, \theta_t^* + w/2)$, the period-t cut-off is given by $\hat{\theta}_{t+1} = -\hat{\theta}_t + 2\theta_t^*$. ■

<u>Proposition</u> <u>6:</u> When the z_t are i.i.d. draws from a distribution that has a strictly positive density on a compact support, the dynamic process generated by (10) has a unique invariant distribution F, and the expected probability distribution at date t converges to F uniformly over initial probability distributions μ.

<u>Proof:</u> Appendix C shows that the system is a random contraction in the sense of Norman [1972] and satisfies uniqueness condition 2.11 of Futia [1982].

We have not been able to characterize this distribution directly. Instead, we have computed an invariant distribution of the simpler system generated by

$$(11)\quad \hat{\theta}_{t+1} = \begin{cases} \hat{\theta}_t + w & \hat{\theta}_t \le \theta_t^* \\ \hat{\theta}_t - w & \hat{\theta}_t > \theta_t^* \end{cases}.$$

Note that system (11) differs from (10) only when $\hat{\theta}_t$ falls in an interval of width w. Normally we will think of the variance of z_t as being much larger than the window width; in this case it may be reasonable to guess that the invariant distributions of (10) and (11) are close together.

We should point out that the simplified system (11), unlike (10), does not have a unique invariant distribution: Because all steps have size w, from initial position θ_0, the support of (11) is concentrated on the grid $\theta_0 + kw$, and so different initial conditions lead to different invariant distributions. Moreover, the supports for of the date-t distribution are different for t even and for t odd. Despite these qualitative differences between systems (10) and (11), the absolute magnitude of the effect of the initial condition is small when w is small, which supports the conjecture that the two systems are similar. Table 1 below provides further support for this belief by comparing Monte Carlo estimates of the steady-state variance of (10) with the variance of the particular invariant distribution of (11) that is computed in proposition 7. As conjectured, the two variances are close when w is small.

To examine the invariant distributions of (11), suppose that the noise terms z_t are are i.i.d with mean 0 and c.d.f. H. Then $\hat{\theta}_t$ follows a Markov process with the transition from $\hat{\theta}_t$ to $\hat{\theta}_t + w$ having probability $\text{Prob}[\theta^* + z_t \geq \hat{\theta}_t] = 1 - H(\hat{\theta}_t - \theta^*)$. The invariant distribution has a particularly simple form when the z_t are uniform on $[-\sigma, \sigma]$ and the parameters are such that there is an invariant distribution whose support is a symmetric grid containing the points $\theta^* - \sigma$ and $\theta^* + \sigma$.

<u>Proposition 7</u>: Suppose the z_t are uniform on $[-\sigma, \sigma]$, and that $M = \sigma/w$ is an integer. Then one invariant distribution of (11) is the binomial $\text{Prob}(\theta = \theta^* + kw) = [((2M!)/(M-k)!(M+k)!]2^{-2M}$; this is the limit of the time-average distribution when the initial condition belongs to the grid $\theta^* \pm kw$, $k \leq M$.

<u>Remark:</u> Recall that the mean of this distribution is θ^*, its variance is $\sigma w/2$, and that the distribution is asymptotically normal as w tends to zero.

<u>Proof</u>: To show that f is an invariant distribution, it is

sufficient to verify that it meets the "detailed balance condition" that for all θ and θ', the (unconditional) probability flow from θ to θ' equals the probability flow in the reverse direction. Thus, we will verify that

$$f(\theta)\ \mathrm{Prob}(\theta_{t+1} = \theta' | \theta_t = \theta) = f(\theta')\ \mathrm{Prob}(\theta_{t+1} = \theta | \theta_t = \theta'),$$

or equivalently that

$$f(\theta)/f(\theta') = \mathrm{Prob}(\theta_{t+1} = \theta | \theta_t = \theta') / \mathrm{Prob}(\theta_{t+1} = \theta' | \theta_t = \theta).$$

Since the probability of a jump of more than w is zero, it suffices to check this conditions between adjacent states, so take $\theta = \theta^{*}+kw$ and $\theta' = \theta^{*}+(k+1)w$ for some integer k between $-M/w$ and $(M-1)/w$. For such states, we have

$$f(\theta)/f(\theta') = \frac{2^{-2M}\left[(2M!)/(M+k!)(M-k)!\right]}{2^{-2M}\left[(2M!)/(M+k+1)!(M-k-1)!\right]} = (M+k+1)/(M-k),$$

and

$$\mathrm{Prob}(\theta_{t+1} = \theta | \theta_t = \theta') / \mathrm{Prob}(\theta_{t+1} = \theta' | \theta_t = \theta) =$$

$$\left[(\sigma+(k+1)w)/2\sigma\right]/(\sigma-kw)/2\sigma = (M+k+1)w/(M-k)w,$$

so detailed balance holds. ■

w/σ	(10)	(11)
.5	.21	.25
.1	.048	.05
.05	.024	.025
.001	.00496	.005

TABLE 1:
STEADY STATE VARIANCE FOR UNIFORM NOISE

As one would expect, the variance of the steady state is decreasing in w, because small w corresponds to small steps in each period. Note that the social optimum is the constant $\hat{\theta} = 0$, and that the expected welfare loss (compared to $\hat{\theta} = 0$) when the cut-off is $\hat{\theta}_t$ is

$$\int_0^{\hat{\theta}_t} \theta \, d\theta = \hat{\theta}_t^2/2.$$

Hence, in the long run the average per-period welfare loss (using the invariant distribution computed in proposition 7) is $1/2 \ E(\theta^2) = 1/2 \ (E(\theta))^2 + 1/2 \ \text{var}(\theta) = [(2\beta+1)^2/8 + \sigma/2]w$, so that steady state welfare is decreasing in w. For small w, despite the lack of either memory or popularity weighting, the spatial nature of the process allows the long-run outcome to be approximately efficient.[19]

While small w's are thus desirable from the viewpoint of the time-average payoff, they entail a significant short-run welfare loss when the initial state is far from the optimum, because in this case the system will take a long time to approach the neighborhood of the optimum. This is true for two reasons: First, $\hat{\theta}_t$ is limited to move at most w per period. Second, in the presence of noise a typical path is likely to take far more than θ_0/w periods to reach a neighborhood of θ^*, because many steps will be in the wrong direction.

For a fixed initial condition and social discount factor, the socially optimal window width will trade off the speed of convergence and the steady state variance, with larger w's being optimal the farther the initial condition is from 0. If the social planner does not know the initial condition and/or the location of the social optimum, the size of the optimal w will depend on the planner's prior beliefs. This tradeoff between speed of adjustment and the variance of the steady state seems a natural feature of the sorts of model we consider[20].

At this point we would like to make a few observations about how the conclusions might change if the players did keep records of their past observations. Since players at locations within σ of θ^* will play both technologies infinitely often, they could eventually learn which technology is better for themselves by keeping such records. However, a few calculations suggest that this learning process will be fairly slow if the random shock to the payoffs has a sizable common component and w is small.

To see this, suppose that the payoffs to each technology are subject to a common shock η_t as well as the idiosyncratic shocks we assumed before, so that system (5) is replaced by

$$(5') \quad \begin{cases} u_t^g(\theta)=\theta + \beta\theta + \varepsilon_{1t} + \eta_t. \\ u_t^f(\theta)=\beta\theta + \varepsilon_{2t} + \eta_t. \end{cases}$$

If the variance of η_t is relatively large, then observations of only one technology at date t are not very informative, and only observations of both technologies in the same period will be of much help. Players at locations more than three standard deviations from θ^* -that is, outside of the interval $\theta^* \pm 3(\sigma w/2)^{1/2}$- rarely see both technologies played, and hence would need a very long memory to learn. Players at locations θ closer to θ^* do see both technologies played more often, but for these players the systematic payoff difference between the technologies is smaller, and hence it may require many observations to be fairly confident one is better. Our informal approximations, reported in appendix E, suggest that this is indeed the case, and in particular that the number of periods required to be fairly confident which technology is better is of the order of $\sigma_\varepsilon^2/ w^2$, so that when w is small a very long history would be required for players to do much better than with our simple rule. Of course, players could use history even when the advantage to doing so is slight or slow to develop, but in these cases it seems less obvious that players would be led to abandon simple rules.

6. Examples of non-linear Technologies

Before considering the implications of popularity weighting in a heterogeneous population, we would like to discuss some examples of what can happen without popularity weighting when the payoffs as a function of location do not take the linear form presumed in equation (5). Suppose for example that the "old" technology f has returns that are identically zero, while $g(\theta) = \cos(\theta)$, so that regions where g is optimal alternate with regions where f is. (See Figure 1) If there is no noise in the system, and the window width is relatively small, then even if all players in locations $\theta \in [-\pi/2, \pi/2]$ adopt the new technology g, the new technology will not spread to the other regions where it is optimal. In this example there are substantial social gains from having the new technology "tested" at a number of diverse locations. (It is for this reason that we would argue that the gentleman farmers may have played an important role in the spread of the new husbandry, even though they were not among the first to adopt it.)

It may also be interesting to note that when the local

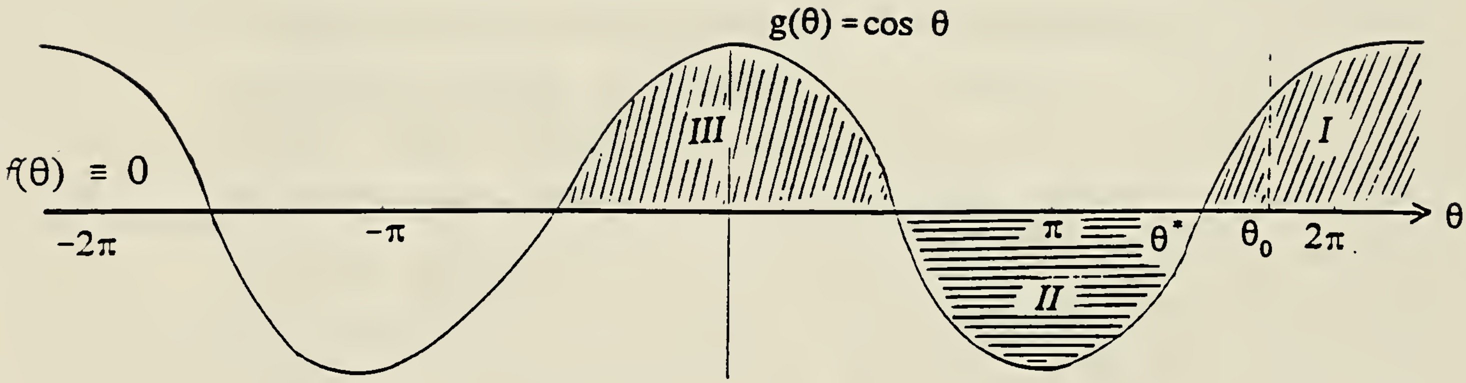

Figure 1

process may fail to spread as widely as it should, random shocks to payoffs can increase social welfare, that is, welfare can increase as the variance of the noise term z_t increases from zero. Suppose that the technologies are $f(\theta) = 0$ and $g(\theta) = \cos(\theta)$, and that the initial state has all players to the right of θ_0 using g and players to the left using f. Without noise, the cutoff will move to $\theta^* \cong 3\pi/2$ and stay there. (See figure 1.) When the support of z_t is sufficiently large, there will eventually be enough consecutive draws of very negative z_t that the cutoff reaches $\pi/2$. From this point, the system may no longer have a single cutoff, as players to the left of $\pi/2$ will tend to switch to g, while those to the right switch back to f. Essentially, the noise leads the players in region II to use the new technology long enough that it can spread from region I to region III.

The next example shows that in certain extreme cases the specification error involved in ignoring how payoffs vary with "location" can allow a technology that is everywhere inferior to completely drive out a better one. This is the case depicted in figure 2 below, in which $f(\theta) = \theta$ and $g(\theta) = \theta-\varepsilon$. If the current cut-off is at $\hat{\theta}$, then the player at $\theta \in [\hat{\theta} - w, \hat{\theta}+w]$ computes $\bar{u}^g(\theta) = \hat{\theta}-\varepsilon + (\theta-(\hat{\theta}-w))/2$, and $\bar{u}^f(\theta) = \hat{\theta} - w + (\theta-(\hat{\theta}-w))/2$. Since $\bar{u}^g(\theta) - \bar{u}^f(\theta) = w - \varepsilon$, if $w > \varepsilon$ all players who observe both technologies choose technology g. Hence $\hat{\theta}_{t+1} = \hat{\theta}_t - w$, and eventually g will take over the entire population.

We should point out that these technologies are quite special: an inferior technology can only drive out a better one if the difference in payoffs $|f-g|$ is small compared to the

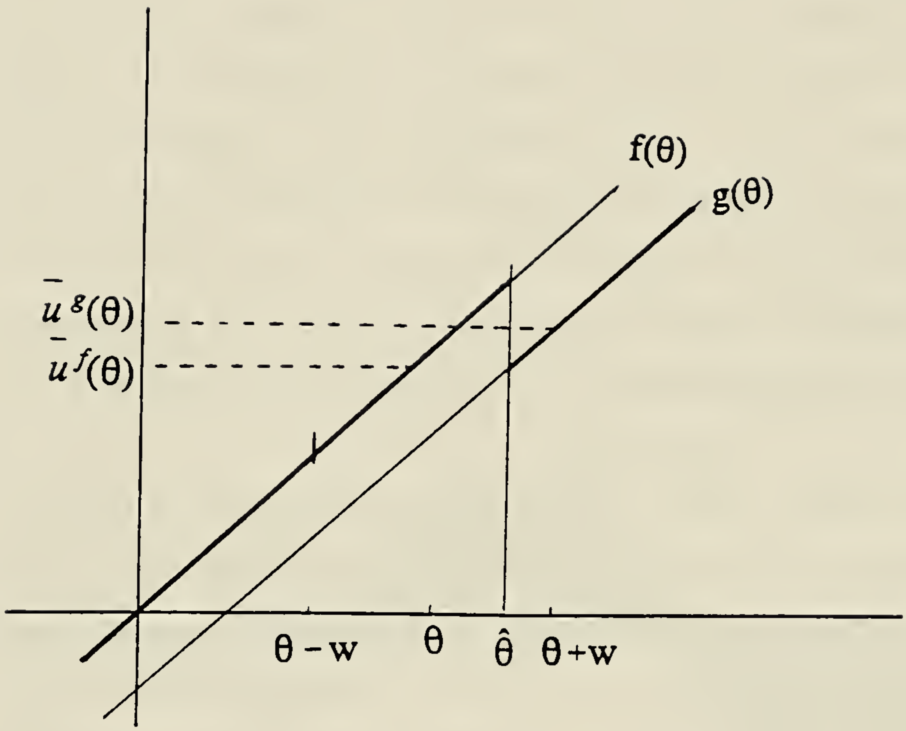

Figure 2

errors caused by estimating the payoffs by their average values in the window. These errors are of the magnitude of $w\, df/d\theta$ and $w\, dg/d\theta$, which bound the difference between the payoffs at $(\theta-w)$ and $\theta+w$. Thus, if w is small, the difference in payoffs $|f-g|$ must be small as well in order for the inferior technology to dominate, and hence even though the wrong technology is adopted everywhere, the payoff loss at each location is not substantial. (In the example above, the payoff loss at each location is ε, and ε must be less than w in order for g to dominate.)

For small window widths, a more substantial payoff loss arises when the new technology is not adopted in a region where it is a substantial improvement. This was the case in the example where $g = \cos(\theta)$ and $f = 0$, so that the regions where g should be adopted are disconnected. We can also modify the example of figure 2 so that g is better than f at every location (and so in particular is better on a connected set) and yet a substantial payoff loss results from g failing to spread. In figure 3, the payoffs to f and g are such that g is much better than f in the neighborhood of $\theta = 0$, but is only slightly better than f for extreme θ values. Hence, if technology g is first introduced at these extreme values, it will be driven out of the population before it can be tried in the center region.

7. Heterogeneous populations and popularity weighting

Our analysis of social learning in homogeneous populations showed that popularity weighting could improve the aggregate performance of the learning process, and that the optimal level of popularity weighting is consistent with individual

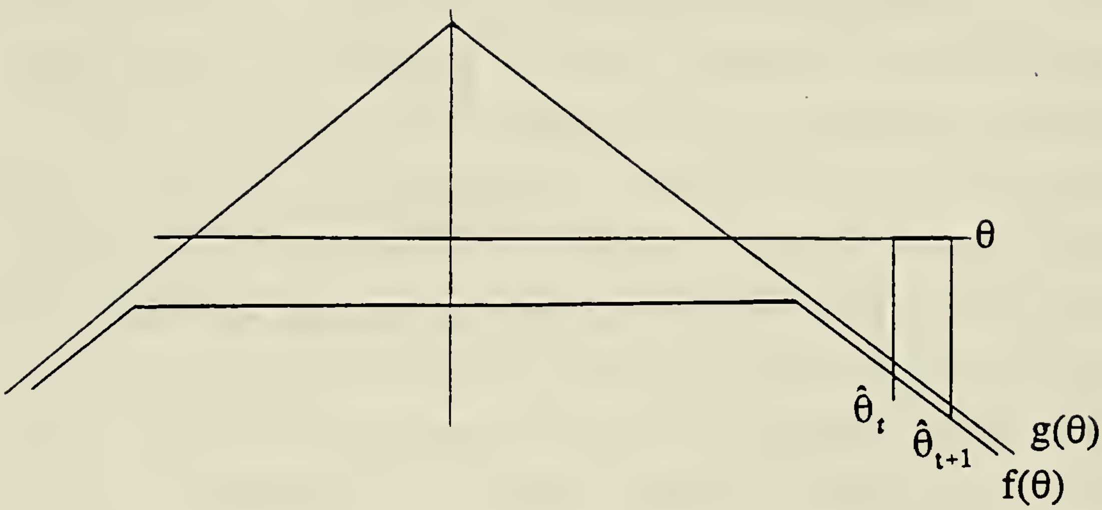

Figure 3

incentives. We will now investigate the implications of popularity weighting in our model of a heterogeneous population with linear technologies.

To model popularity weighting, let $x_t(\theta)$ be the fraction of players in the interval $[\theta-w, \theta+w]$ who use technology g. In the spirit of the popularity weighting rule (3), we now modify the the decision rule (6) used in sections 5 and 6 and suppose that players use the decision rule

(12) "Play g at period t+1 iff $\bar{u}_t^g(\theta) - \bar{u}_t^f(\theta) \geq m(1-2x_t(\theta))$,"

where, as before, the parameter m indexes the importance of popularity in the players' decisions.

Since the analysis of this system is quite close to that of the system without popularity weighting, we will give the results without proof. As in section 5, if the state in period t corresponds to a cut-off rule, so will the state in period t+1. In addition, without noise terms the system has the same, unique, steady-state cut-off $\theta^* = -(2\beta+1)w/2$. However, the introduction of popularity weighting does change the dynamics in two ways. First, in the absence of noise terms, the system converges to the steady-state cut-off from any initial cut-off; the oscillations described in proposition 4 do not arise. Second, (and relatedly) movements of less than one window width become more common, as players are more hesitant to a less popular technology.

The following proposition gives a more precise description of the dynamics.

<u>Proposition</u> <u>8:</u> From an initial cut-off $\hat{\theta}_0$, the system described by decision rule (12) and payoffs (5) evolves according to

(13):

$$\hat{\theta}_{t+1}=\begin{cases}\hat{\theta}_t+w & \text{if } \hat{\theta}_t<\theta^*_t-(m+w/2)\\ \theta^*_t+(2m-w)/(2m+w)(\hat{\theta}_t-\theta^*_t) & \text{if } \hat{\theta}_t\in[\theta^*_t-(m+w/2),\ \theta^*_t+(m+w/2)]\\ \hat{\theta}_t-w & \text{if } \hat{\theta}_t>\theta^*_t+(m+w/2).\end{cases}$$

Proof: Omitted. The calculations involved are straightforward, and quite similar to those of proposition 4. Note that the dynamics above reduce to those of proposition 4 when $m = 0$, as they should do.

To see that, in the absence of noise, the system converges to θ^* from any initial cutoff, note that the cut-off moves a full window width so long as $|\hat{\theta}_t-\theta^*| > m + w/2$. Eventually then $|\hat{\theta}_t-\theta^*| \leq m + w/2$, and from then on $\hat{\theta}_{t+1}-\theta^* = [(2m-w)/(2m+w)](\hat{\theta}_t-\theta^*)$, so that the system converges to θ^* at a geometric rate.

Note also that for a given $\hat{\theta}_t$, the system will move less than a full window width whenever the realization of θ^*_t is in an interval of $2m + w$. This show that popularity weighting makes the system more "sluggish," and suggests that it will reduce the variance of the long-run distribution. To verify this intuition, and determine the extent to which popularity weighting reduces the variance, we characterize the long-run distribution in one special case.

Proposition 9: (a) If the z_t are i.i.d. draws from a distribution that has a strictly positive density on a compact support, the dynamic process defined by (5) and (12) has a unique invariant distribution.

(b) If the z_t are i.i.d. draws from the uniform distribution on

$[-\sigma,\sigma]$ and $m \geq 2\sigma$, the invariant distribution f is concentrated on the interval $[\theta^*-\sigma- w/2, \theta^*+\sigma +w/2]$ and satisfies

$$E_f(\hat{\theta}) = \theta^*,$$

and

$$\mathrm{var}_f(\hat{\theta}) = \sigma^2 w/6m.$$

<u>Proof:</u> (a) Omitted; the argument is very close to that for proposition 6.

(b) Appendix D shows that there is a deterministic, finite time T for which the cut-off $\hat{\theta}_T$ is in the interval $[\theta^*- \sigma - w/2, \theta^*+ \sigma +w/2]$, and that once this interval is reached, $\hat{\theta}_{T+s}$ remains in the interval for all subsequent periods T+s.

Given a T satisfying these claims, we have $|\hat{\theta}_{T+s} -\theta^*_{T+s}| < |\hat{\theta}_{T+s}-\theta^*| + |\theta^*_{T+s} -\theta^*| < (\sigma + w/2) + \sigma$, which is less than $m + w/2$ from our assumption that $m > 2\sigma$. Hence we the evolution of $\hat{\theta}_{T+s}$ from T on is determined by the second case in proposition 8.

Writing $c = (2m-w)/(2m+w)$, and applying this rule repeatedly, we find that

$$\hat{\theta}_{T+s} = (1-c)\sum_{\tau=0}^{s} c^{\tau} \theta^*_{T+s-\tau}.$$

Hence,

$$E(\hat{\theta}_{T+s}|\hat{\theta}_T) = (1-c)\sum_{\tau=0}^{s} c^{\tau}E(\theta^*) + c^s(1-c)\hat{\theta}_T \rightarrow E(\theta^*),$$

and

$$\mathrm{var}(\hat{\theta}_{T+s}|\hat{\theta}_T) = (1-c)^2 \sum_{\tau=0}^{s} c^{2\tau}\mathrm{var}(\theta^*) =$$

$$(1-c)\sigma^2/3(1+c)= 2w\sigma^2/12m = \sigma^2 w/6m. \blacksquare$$

Comparing the steady state distributions for $m > 2\sigma$ with that for $m = 0$, we see that popularity weighting reduces the long-run variance by a factor of $1/3m$.

The welfare consequences of increasing m for fixed w are similar to those of decreasing w for fixed m: in both cases, the steady state distribution becomes more efficient, while the speed at which the system converges decreases. It may be interesting to note, however, that in this simple model there is one way to change the parameters to speed up the rate of convergence (when the initial cutoff is far from the optimum) without altering the steady-state variance, namely increasing the window width w while holding the ratio of w/m fixed.[21]

8. Concluding Remarks

The various models we have presented suggest that even very naive learning rules can lead to quite efficient long-run social states, at least if the environment is not too highly nonlinear. Moreover, popularity weighting can contribute to this long-run efficiency, and the use of popularity-weighting passes a crude first-cut test of consistency with individual incentives. Of course, there are many other plausible specifications of behavior rules for social learning, so it is interesting to speculate about the robustness of our conclusions.

In this light, we would like to report simulation results for one simple modification of popularity weighting that seems to improve the short-run performance of the system without changing its long-run behavior. For this purpose, we return to the homogeneous-population model of section 4, and now suppose that players give weight to trends in the relative popularity of the two technologies as well as to the popularity itself.

More precisely, suppose that players now choose technology g iff the realized difference in payoffs $u_t^g - u_t^f$ exceeds the expression $m(1-2x_t)-c(x_t-x_{t-1})$, where x_t-x_{t-1} is the trend in popularity. Since the trend variable converges to zero along any path where the system converges to a steady state, the system still converges to the better technology with probability 1 when when $m = \sigma$. However, if the initial state is far from the optimum, as in the case when a superior technology is first introduced, one would expect that responsiveness to trends would help to increase the speed with which the new technology is adopted.

To test this intuition, we ran two simulations, both with the noise term ε uniformly distributed on $[-\sigma, \sigma]$ and popularity weighting $m = \sigma$. In the first, the fraction α who adjust each period was .5, and the mean payoff difference θ was $.1\sigma$; in the second, $\alpha = .1$ and $\theta = .02\sigma$. In both cases, we counted the number of periods required for the system to move from initial state $x_0 = .05\sigma$ to $x = .99\sigma$. The results, reported in Table 2, show that at least one obvious modification of our assumptions reinforces the impression that naive rules can perform fairly well.

	α=.5, θ=.1	α=.1, θ=.02
c = 0	39	940
c = 5	26	710
c = 10	28	470

Table 2: Trend Weighting and the Speed of Convergence[22]

There are a number of other extensions that we have not considered but seem important. Players might use rules of thumb which make some use of historical data. Also, players might be arranged in more complex networks than the simple linear structure we have considered. Finally, our results suppose either that rules of thumb are exogenous, or, in Proposition 3, are equilibrium choices of a static game. It would be interesting to complement these results with an analysis of a dynamic process by which players adjust their rules of thumb along with their choice of technology.

Finally, we should point out that popularity weighting is not always as beneficial as our results might suggest. Consider the problem of children is a poor neighborhood choosing whether to pursue higher education. If students who have done so in the past tend to move out of the neighborhood, and past residents are underrepresented in the observation windows, then the choice of higher education will appear less popular than it really is, and decisions based on popularity may be biased against this choice.[23]

Appendix A: Optimal Popularity Weighting with Other Distributions

To better understand the forces generating proposition 2a- that a single choice of popularity weight yields the optimal long-run distribution uniformly over all values of θ- we show that analogous results obtain when the per-period noise term ε_t has distribution F with unbounded support.

Suppose first that $\alpha = 1$, so that the entire population adjusts every period, and hence the state x_t takes on only the values 0 and 1. If we let s_t denote the vector $[Prob(x_t)= 0, Prob(x_t) = 1]$, we have $s_{t+1} = s_t A$, where the transition matrix is A =

$$\begin{bmatrix} F(m-\theta) & 1-F(m-\theta) \\ F(-m-\theta) & 1-F(-m-\theta) \end{bmatrix}.$$

Since this matrix is strictly positive, the system is ergodic; the unique invariant distribution μ^* is given by

$$\mu^*(x=0) = F(-m-\theta)/[F(-m-\theta)+ 1-F(m-\theta)].$$

If F is the standard normal distribution, then as m increases, the ratio $F(-m-\theta)/ 1- F(m-\theta)$ converges to 0 if $\theta > 0$, and converges to ∞ if $\theta < 0$. Hence for large m, the ergodic distribution of the system places probability near 1 on the correct choice. Moreover, the same is true for any distribution for which the the ratio $F(-m-\theta)/ 1- F(m-\theta)$ converges to 0 if $\theta > 0$, and to ∞ if $\theta < 0$. (This is what is meant by saying that the tails of the distribution are "infinitely revealing.")

With a more involved argument, we have shown that the same conclusion holds for any $\alpha \in (0,1)$ when players use the (discontinuous) popularity weighting "if $x_t \geq 1/2$, choose g iff $u_t^g-u_t^f \geq -m$; if $x_t < 1/2$, choose g iff $u_t^g-u_t^f \geq -m$." The details are available on request, here is the intuition for the result. Note first that when $m = \infty$ the system is deterministic with stable steady states at 0 and 1. If m is finite but very large compared to α and to the distribution, then steps the "wrong way" (i.e. decreasing steps when $x_t > 1/2$) are rare "innovations", and when the distribution is symmetric, transits from 0 to 1/2 and from 1 to 1/2 both take same number of innovations. If the tails of the distribution are infinitely revealing, then as $m \to \infty$ innovations towards the better technology become infinitely more likely than innovations towards the inferior one, and the analysis of Freidlin and Wentzell [1984] suggests that the limit of the ergodic distributions will be oncentrated on the better technology. To establish this formally, we partitionthe interval into a large number of (appropriately chosen) small subintervals, and approximate the original system by two finite-state Markov processes, whose ergodic distributions will serve as bounds on the ergodic distribution of the original system. We then use the discrete-time, finite-state translation of Freidlin and Wentzell's results (Kandori, Mailath and Rob [1992], Young [1992]) to confirm the intuition above, i.e. the limits of the ergodic distributions of thefinite-state process are concentrated on the subinterval corresponding to the better choice. The above suggests that infinitely-revealing tails are sufficient for there to be a single popularity rule that is approximately optimal for all θ. Moreover, this rule has the nice feature that it need not be tailored to the exact form of the distribution. Even when the tails are not infinitely revealing, however, there is another popularity rule that seems to perform very well, namely "choose g iff $u_t^g - u_t^f \geq F^{-1}(1-x_t)$."

With this rule,

$$E(x_{t+1}|x_t) = (1-\alpha)x_t + \alpha \, \mathrm{Prob}[\theta + \varepsilon_t \geq F^{-1}(x_t)]$$

$$= x_t + \alpha\left(F(\theta+F^{-1}(x_t))-x_t\right),$$

so that $E(x_{t+1}|x_t) > x_t$ if and only if $\theta > 0$; the system tends to drift towards the correct choice. Although the system may converge to the wrong technology with positive probability, the simulation results reported in table 2 for the logistic and Laplace distributions (which both have non-revealing tails) suggest that when α is small the system is very likely to converge to the right choice. Intuitively, when α is small, the system evolves through a series of small steps that allow the drift to outweigh the random forces. We conjecture that there may be a general result these lines.

Appendix B :Proof of Proposition 2c

If σ-m > $|\theta|$, then neither 0 nor 1 is an absorbing state. Our first step is to show that there is a unique invariant distribution. To do so, we first note that the stochastic system (4) is a random contraction in the sense of Norman [1972].[24] A random contraction is a stochastic system in which the realization of an i.i.d. auxiliary variable (call it ω) is used to determine which of a family of mappings $\varphi \in V$ is used to send x_t to x_{t+1}, and each φ_ω is a contraction "on average." In our context, ω corresponds to the realized difference in payoffs, and there are only two maps φ_ω: $\varphi_+(x_t) = (1-\alpha)x_t+\alpha$, and $\varphi_-(x_t)=(1-\alpha)x_t$, both of which are contractions, so that (4) is indeed a random contraction. Norman's results then imply that the Markov operator associated with system (4) is quasi-compact. We next note that when $|\theta| < \sigma-m$, the system (4) satisfies the uniqueness criterion 2.11 of Futia [1982]: for any neighborhood U of the point x = 1/2, and any point x' in [0,1], there is an n such that the probability the system starting at x' in is in U exactly n periods later is strictly positive. (If $m \geq \sigma$, the uniqueness condition fails, as both x = 0 and x = 1 are absorbing.)

The last step is to compute the mean and variance of the invariant distribution μ. Using $E_\mu(x_t) = E_\mu(x_{t+1})$, we have

$$E_\mu(x) = (1-\alpha)E_\mu(x) + \alpha\int p(x)d\mu(x),$$

where $p(x) = (\sigma-m+\theta)/2\sigma + (m/\sigma)x_t$ is the probability that $\theta+\varepsilon_t \geq m(1-2x_t)$, which is the probability that $x_{t+1} = (1-\alpha)x_t+\alpha$. Simple algebra then shows that $E_\mu x = 1/2 + \theta/2(\sigma-m)$.

To compute the variance, we first write the identity

$$E_\mu(x^2) = \int\left[(1-p(x))[(1-\alpha)x]^2 + p(x)[(1-\alpha)x+\alpha)]^2\right]d\mu(x) =$$

$$E_\mu(x^2)\left[(1-\alpha)^2 + 2\alpha(1-\alpha)m/\sigma\right] + E_\mu(x)\left[2\alpha(1-\alpha)(\sigma-m+\theta)/2\sigma + \alpha^2 m/\sigma\right] +$$

$$\alpha^2(\sigma-m+\theta)/2\sigma;$$

solving for $E_\mu(x^2)$ and computing $var(x) = E_\mu(x^2) - (E_\mu(x))^2$ gives the desired result.∎

APPENDIX C: PROOF OF PROPOSITION 6

To begin we rewrite (10) in the equivalent form (10')

$$(10') \qquad \hat{\theta}_{t+1} = \begin{cases} \min[\hat{\theta}_t+w, 2(\theta^*+z_t)-\hat{\theta}_t)]\theta^*+z_t \geq \hat{\theta}_t \\ \max[\hat{\theta}_t-w, 2(\theta^*+z_t)-\hat{\theta}_t)]\theta^*+z_t < \hat{\theta}_t \end{cases}$$

To show that the system (10) is a "random dynamical system" as described by Futia [1982], we note that the auxiliary events are the z_t. The probability distribution Q on the z's does not depend on the current state, and so in particular is continuous in the state, and the map $\varphi(\theta,z)$ defined by $\hat{\theta}_{t+1} = \varphi(\hat{\theta}_t, z_t)$ is easily seen to be continuous in θ for fixed z, so that (10) is indeed a random dynamical system.

Next we check that it is a random contraction, as in Futia's definition 6.2. Because the map Q is constant in θ, the constant M in part (a) of the definition can be taken to equal 0. Next we must show that for all z, and all $\theta \neq \theta'$, $d(\varphi(\theta,z), \varphi(\theta',z)) \leq d(\theta,\theta')$, and that for all θ and θ' there is a positive probability of z such that $d(\varphi(\theta,z), \varphi(\theta',z)) < d(\theta,\theta')$.

To show that $d(\varphi(\theta,z), \varphi(\theta',z)) \leq d(\theta,\theta')$, we note that for all θ and θ' and all z, either (a) both θ and θ' move in the same direction (e.g. $(\varphi(\theta,z)-\theta)(\varphi(\theta',z)-\theta') > 0$) or (b) $\varphi(\theta,z)-\theta \geq 0 \geq \varphi(\theta',z)-\theta'$. Case a has three subcases: either (1) φ moves both locations by w, so that $d(\varphi(\theta,z), \varphi(\theta',z)) = d(\theta,\theta')$; or (2) the location closer to θ^*+z_t moves less than w, and the state farther away moves w, so that $d(\varphi(\theta,z), \varphi(\theta',z)) < d(\theta,\theta')$, or (3) both locations move by less than w, in which case the two locations are reflected about the point θ^*+z_t, and $d(\varphi(\theta,z), \varphi(\theta',z)) = d(\theta,\theta')$.

In case (b), suppose w.l.o.g. that $\theta < \theta'$; then case b implies that $\theta \leq \theta^*+z \leq \theta'$, and so $d(\theta, \theta') = d(\theta, \theta^*+z) + d(\theta^*+z, \theta')$. Using the triangle inequality, this implies that

$$(C1) \qquad d(\varphi(\theta,z), \varphi(\theta',z)) - d(\theta,\theta') \leq$$

$$d(\varphi(\theta,z), \theta^*+z) + d(\theta^*+z, \varphi(\theta',z)) - d(\theta, \theta^* +z) - d(\theta^*+z, \theta')$$

$$=[d(\varphi(\theta,z), \theta^*+z)- d(\theta, \theta^*+z)] +[d(\theta^*+z, \varphi(\theta',z))- d(\theta^*+z,\theta')],$$

and inspection of (10') shows that each of the terms in square brackets is non-positive. Thus $d(\varphi(\theta,z), \varphi(\theta',z)) \leq d(\theta,\theta')$ for all z, θ, and θ'.

To show that for all θ and θ' there is positive probability that $d(\varphi(\theta,z), \varphi(\theta',z)) < d(\theta,\theta')$, let $\theta < \theta'$, and suppose first that $\theta - \theta^* > -\sigma + w/2$. Then for sufficiently small $\varepsilon > 0$ there is a positive probability that z lies in any sufficiently small neighborhood of $\theta - \theta^* + \varepsilon - w/2$, and for z's in this neighborhood, moves less than w to the left, while θ' moves w, so that $d(\varphi(\theta,z), \varphi(\theta',z)) < d(\theta,\theta')$. If $\theta - \theta^* \le -\sigma + w/2$ but $\theta' - \theta^* < \sigma - w/2$, a similar argument establishes the existence of a range of z's such that both θ and θ' move to the right, with θ' moving less than θ. Finally, if $\theta - \theta^* \le -\sigma + w/2$ and $\theta' - \theta^* \ge \sigma - w/2$, then $\theta' - \theta > w$, and $d(\varphi(\theta,z), \varphi(\theta',z)) < d(\theta,\theta')$ for z's in a neighborhood of $\theta + w/2$. Thus (10') is a random contraction.

The last step in the proof is to verify that (10') satisfies Futia's uniqueness condition 2.11, which requires that there be a point θ such that for any neighborhood U of θ and any θ, there is an n such that when the system begins at θ, it has a positive probability of being in U in period n. It is easy to see that e.g. $\theta^{\frac{1}{2}} = \theta^*$ satisfies this condition.∎

<u>APPENDIX D Proof of Proposition 9, part b.</u>

To complete the proof, we must show that there exists a deterministic, finite time T such that (i) $|\hat{\theta}_T - \theta^*| < \sigma + w/2$, and (ii) that $|\hat{\theta}_{T+s} - \theta^*| < \sigma + w/2$ for all subsequent dates T+s. Define $d_t = |\hat{\theta}_t - \theta^*|$. Note that since $(\theta^*_t - \hat{\theta}_t)$ and $(\hat{\theta}_{t+1} - \hat{\theta}_t)$ have the same sign, and $|\theta^*_t - \theta^*| \le \sigma$, $(\theta^* - \hat{\theta}_t)$ and $(\hat{\theta}_{t+1} - \hat{\theta}_t)$ have the same sign whenever $d_t > \sigma + w/2$. Hence,

$$\text{(D1)} \quad d_{t+1} = |d_t - (|\hat{\theta}_{t+1} - \hat{\theta}_t|)| \text{ whenever } d_t > \sigma + w/2.$$

As a first step towards proving (i), we show that for any initial condition there exists a finite T' such that regardless of the sample path, <u>either</u> $d_{T'} < \sigma + w/2$ <u>or</u> $d_{T'} \le |\hat{\theta}_{T'+1} - \hat{\theta}_{T'}|$. To see this, note that (D1) implies that until such a T' is reached, $d_t - d_{t+1} = |\hat{\theta}_{t+1} - \hat{\theta}_t|$, and from proposition 8,

$$|\hat{\theta}_{t+1} - \hat{\theta}_t| = \min\{ w, [2w/(2m+w)](\hat{\theta}_t - \theta^*_t)\} \ge$$
$$\min\{w, [2w/(2m+w)]w/2\}.$$

Thus, until the conditions defining T' are satisfied, the decrease in d_t is bounded below by a positive constant which is independent of the sample path.

If $d_{T'} \le \sigma + w/2$, setting $T = T'$ completes the proof of (i). The remaining case is $\sigma + w/2 < d_{T'} \le |\hat{\theta}_{T'+1} - \hat{\theta}_{T'}|$. In this case, (D1) implies that $d_{T'+1} = |\hat{\theta}_{T'+1} - \hat{\theta}_{T'}| - d_t$, which is less than $w - (\sigma + w/2) = w/2 - \sigma < w/2 + \sigma$. Hence we can set $T = T'+1$ to complete the proof of (i).

To prove claim (ii), note that when $|\hat{\theta}_T - \theta^*| < \sigma + w/2$, we have $|\hat{\theta}_T - \theta^*_T| \le (\sigma + w/2) + \sigma$, which is less than $m + w/2$ from the assumption that $m > 2\sigma$. From proposition 8 we then have $\hat{\theta}_{T+1} = \theta^*_T + [(2m-w)/(2m+w)](\hat{\theta}_T - \theta^*_T)$, and since both θ^*_T and $\hat{\theta}_T$ lie in the interval $[\theta^* - \sigma - w/2, \theta^* + \sigma + w/2]$, so does $\hat{\theta}_{T+1}$. The claim now follows from induction on s. ∎

Appendix E:

This appendix gives a rough computation of how many periods would be needed on average for a player using history to be reasonably sure he knew which technology is better. If we let $\sigma_\varepsilon = \mathrm{var}(\varepsilon_{1t}) + \mathrm{var}(\varepsilon_{2t})$ denote the variance of $u^g_t - u^f_t$, then the player at θ will need about $(\sigma_\varepsilon/\theta)^2$ observations of both technologies to be fairly confident (about 85%- i.e.15% chance of a false rejection) that he knew which technology was better. For $\theta \le \theta^* + (\sigma w/2)^{1/2}$, this requires on the order of $2\sigma^2_\varepsilon/\sigma w$ observations of both technologies. For $\theta^* + (\sigma w/2)^{1/2} < \theta \le \theta^* + 3(\sigma w/2)^{1/2}$, at least $2\sigma^2_\varepsilon/9\sigma w$ observations of both technologies are required.

Our next step is to approximate the frequency with which these observations arrive. To do so, we approximate the steady-state distribution by a normal distribution, and then note that the density of the normal is less than $.25/\sigma$ at points more than 1 standard deviation away from the mean, and that the density of the normal at the mean is less than $.5/\sigma$. Finally, we recall that only players within w of the current period's cutoff see both technologies being used. Thus we conclude that for players within one standard deviation of θ^* the required observations arrive on average every $[2w\cdot.5/\sigma]^{-1} = \sigma/w$ periods; since these players need about $2\sigma_\varepsilon^2/\sigma w$ observations, $2\sigma^2{}_\varepsilon/w^2$ periods are required. For players between 1 and 3 standard deviations away, the required observations arrive about every $2\sigma/w$ periods; these players need $2\sigma_\varepsilon{}^2/9\sigma w$ observation, so that about $4\sigma^2{}_\varepsilon/9w^2$ periods are required.

REFERENCES

Apodaca, A. [1952] "Corn and Custom: Introduction of Hybrid Corn to Spanish American Farmers in New Mexico," in E.H. Spicer, ed., Human Problems in Technological Change," Russel Sage Foundation, New York.

Banerjee, A. [1991a] "The Economics of Rumours," forthcoming in the Review of Economic Studies

Banerjee, A. [1991b] "A Simple Model of Herd Behavior," , forthcoming in the Quarterly Journal of Economics.

Bikchandari, S., D. Hirshleiffer, and I. Welch [1991] "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," forthcoming in the Journal of Political Economy.

Bush, R.R. and F. Mosteller [1955] Stochastic Models for Learning, Wiley, New York.

Chambers, J.D. and G.D. Mingay [1966] The Agricultural Revolution 1750-1880, Schocker, New York.

Conlisk, J. [1980] "Costly Optimizers versus Cheap Imitators," Journal of Economic Behavior and Organization, 1:275-293.

Cross, J.G. [1983] A Theory of Adaptive Economic Behavior, Cambridge University Press

Ellison, G. [1991] "Learning, Local Interaction, and Coordination," mimeo.

Ernle, R.E.P. [1912] English Farming Past and Present, Longmans Green and Co., London.

Futia, C. [1982] "Invariant Distributions and the Limiting Behavior of Markovian Economic Models, Econometrica, 50, 377-408

Kandori, M., G. Mailath, and R. Rob [1992] "Learning, Mutation, and Long Run Equilibrium in Games," mimeo.

Kerridge, E. [1967] The Agricultural Revolution, George Allan & Unwin, London.

Manski, C. [1990] "Dynamic Choice in a Social Setting," University of Wisconsin SSRI discussion paper 9003.

Mantoux, P. [1905] The Industrial Revolution in the 18th Century, Harper and Row, N.Y.; reprinted in 1961.

Mingay, G.E. [1977] The Agricultural Revolution, Adam & Charles Black, London.

Norman, M.F. [1968] "Some Convergence Theorems for Stochastic Learning Systems with Distance-Diminishing Operators," J. Mathematical Psychology, 5, 61-101.

Norman, M.F. [1972] Markov Processes and Learning Models, Academic Press.

Rogers, E. with F. Shoemaker [1971] Communication of Innovations: A Cross-Cultural approach, 2nd edition, New York, the Free Press.

Rothschild, M. [1974] "A Two-Armed Bandit Theory of Market Pricing," J. Econ. Theory, 9, 185-202.

Ryan, B. and N. Gross [1943] "The Diffusion of Hybrid Seed Corn in Two Iowa Communities," Rural Sociology

Schmalansee, R. [1975] "Alternative Models of Bandit Selection," J. Econ Theory, 10, 333-342.

Slicher van Bath, H.S. [1965] The Agrarian History of Europe A.D. xxxx-1850, Edward Arnold, London.

Smith, L. [1990] "Error Persistence and Experimental versus Observational Learning," mimeo.

Timmer, C.P. [1969] "The Turnip, the New Husbandry, and the English Agricultural Revolution," Quarterly Journal of Economics, 83, 375-395.

REFERENCES

[1]See Rogers and Shoemaker [1971] for an extensive discussion of empirical research on adoption processes, especially in development. Mansfield [1968] and Ryan and Gross [1943] are classic studies of technology adoption in basic industries and agriculture, respectively.

[2] Cross [1983] develops a model of boundedly rational, adaptive choice with a similar information structure.

[3] Smith's paper models the pricing decisions of monopolistically competitive firms, for which the assumption of unobserved payoffs seems more plausible. Banerjee [1991a] is a model of investment decisions; Banerjee [1991b] is less explicit about its intended interpretations. Bikchandari et al suggest that their model is appropriate for a wide range of decision problems, including that of technology adoption.

[4] Manski [1990] considers heterogeneous populations from an econometric viewpoint. His paper provides a sufficient condition for an individual to be able to obtain a consistent estimate of his own optimal choice by observing the choices and outcomes of others.

[5] Note that when the players are heterogeneous, a central planner would need to know the relative payoffs of the competing technologies for every player in order to implement the optimum by fiat. Centrally-based agricultural reformers are often hampered by their lack of understanding of the variation in farmers' tastes and production costs. For example, Apodaca [1952] describes how a planner tried to induce a New Mexico community to adopt a hybrid corn. The innovation was adopted and then discontinued despite doubling yields, as the villagers decided the taste and consistency of the corn were inappropriate for making tortillas.

[6](In the second model the environment is complicated enough that a great many periods would be required to obtain good estimates, as we discuss in section 5.)

[7] See Mingay [1977] and Slicher von Bath [1965] for more thorough descriptions of the technology. Timmer [1965] discusses the extent of the resulting gains in productivity.

[8]Kerridge [1967] pp. 28-34, 339-341.

[9]Chambers and Mingay [1966] p. 55.

[10]See Timmer [1969] and Kerridge [1967]. Slicher von Bath [1965] p.243 gives a similar figure for the rate of diffusion in France.

[11]See Timmer [1969] for an excellent summary of this debate.

[12] The consideration of inertia is further motivated by the empirical evidence that there is typically a substantial lag between the time individuals first learn of the existence of a technology and the time they adopt it. Ryan and Gross [1943] found that farmers in two rural communities on average adopted hybrid seed corn 9 years after they first heard of the innovation. Other studies cited in Rogers and Shoemaker [1971], p. 129, report lags of 2-4 years for the adoption of weed spray in Iowa and fertilizer in Pakistan. Note that the spread of literacy and modern communication media will speed up the rate at which farmers become aware of a new technology's existence, but do not seem to have eliminated the lag between becoming informed and deciding to adopt.

[13] If we interpret x_t as the probability that a singleindividual chooses g, as opposed to the population fraction, (2) is an example of the linear stochastic learning theory (LSLT) of Bush and Mosteller [1952]. This theory describes a traditional one-player bandit problem, in which only 1 arm is observed in each period; it is also assumed that the only possible outcomes are "reward" and "failure," with the probability of arm i being rewarded being π. Our model corresponds to the special case in which $\pi_1 = 1-\pi_2$, and the two arm's outcomes in a given period are perfectly negatively correlated. (In the bandit problem, the joint distribution is irrelevant.) See Schmalansee [1975] for a brief survey of these models, and a discussion of their applicability to market pricing. In our system (2), Schmalansee's constants L, L, G, and G all equal α, while Schmalansee's α_1 and α_2 both equal 1, and his β_1 and β_2 equal 0. Schmalansee argues that in some cases the behavior these models prescribe, with both actions taken infinitely often, is more realistic than that of the optimal solution to the discounted bandit problem.

[14] The empirical literature suggests that popularity weighting is a factor, but reliable estimates of m are hard to come by. Rogers and Shoemaker (*op cit.*, p. 142) say that "many students of peasant life feel" that innovations must be 20% to 30% better to be adopted; they also cite a President's Science Advisory Committee figure of 50% to 100% . From our reading, it is not clear whether these premia reflect popularity weighting or inertia.

[15]We should warn the reader that the results we obtain for the uniform case will seem to rely on the fact that this distribution has compact support: an observation that $u_t^g - u_t^f > \sigma$ implies that θ is greater than zero. However, compact support is not what underlies our conclusions. Appendix A shows that the non-linear rule "only switch if the observed payoff difference is large compared to the popularity" leads to a long-run distribution that places most of its weight on the better choice whenever the distribution of errors is "infinitely revealing in the tails." The appendix also reports simulations of a more complex rule that seems to work well even when the tails are not infinitely revealing.

[16] An alternative explanation is to use the fact that the rule $m = \sigma$ yields the optimal long-run decision, and that since σ is the standard deviation of the per-period payoff differences, rescaling the utility function rescales σ in the same way.

[17] Unless the payoff difference is so extreme that $\theta - \sigma > m$, in which case the rate of adoption is independent of θ. Note that the rate is also an increasing function of θ when $m = 0$, s provided that θ is smaller than σ.

[18]However, the correlation is easy to explain as the result of an optimal investment policy under complete information if adopting the innovation requires investing in a capital good.

[19] Although our leading example of very small window widths is the English agricultural revolution, small window widths should not be seen as requiring illiterate agents. Anecdotal evidence suggests that farmers often distrust the information of central authorities and experts, and prefer to see how innovations work out in their neighborhood. Ryan and Gross [1943] found that the experiences of neighbors was an important factor in the adoption of hybrid seed corn by 20th century Iowa farmers.

[20]Although we have not checked the details, it seems that a combination of large window widths with a rule of proximity-weighted averages could combine quick convergence with a small long-run variance.

[21]However, as w increases the specification bias grows. When w is large, it may be more natural to suppose that players weight the experience of those nearby more than that of those who are farther away but still within their window.

[22]Based on estimated standard errors, the first two digits are correct at the .95 level.

[23]We thank Roland Benabou for this observation.

[24](See Futia's [1982 survey for a summary of Norman's results, and other techniques for establishing that the invariant distribution is unique.)