



Rules of Thumb for Social Learning

Author(s): Glenn Ellison and Drew Fudenberg

Source: *The Journal of Political Economy*, Vol. 101, No. 4 (Aug., 1993), pp. 612-643

Published by: [The University of Chicago Press](#)

Stable URL: <http://www.jstor.org/stable/2138741>

Accessed: 09/12/2010 02:18

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=ucpress>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Political Economy*.

Rules of Thumb for Social Learning

Glenn Ellison and Drew Fudenberg

Harvard University

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

I. Introduction

This paper presents two simple models of how economic agents decide which of two technologies 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 examples concern the adoption of agricultural technology, we believe that the models may also be applicable to the diffusion of new man-

We thank Abhijit Banerjee, Roland Benabou, Mathias Dewatripont, Sara Fisher Ellison, and Bengt Holmström for helpful comments. We are happy to acknowledge financial support from National Science Foundation grant SES 90-08770, the John Simon Guggenheim Foundation, and the Sloan Foundation Graduate Program and the hospitality of the Institut d’Economie Industrielle (Toulouse).

[*Journal of Political Economy*, 1993, vol. 101, no. 4]
© 1993 by The University of Chicago. All rights reserved. 0022-3808/93/0104-0008\$01.50

agement practices and to parents' decisions whether to send their children to a public or a private school.¹

The learning environments we study have three main features: First, agents observe both their neighbors' choices and the payoffs that these choices generate. Second, agents periodically reevaluate their decisions, as opposed to making a once-and-for-all choice. Third, we consider the possibility that players may be sufficiently heterogeneous that under full information they would not all make the same choice.

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 from 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 model has a homogeneous population of players choosing between two competing technologies, with the payoff to each technology subject to an aggregate independent and identically distributed (i.i.d.) shock. Each period, only some fraction of the players have the opportunity to revise their choices; the other players continue using whichever technology they used in the previous period.

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 that incorporate "popularity weighting," a tendency to choose a more popular technology even if it was somewhat less profitable last period. We find that the appropriate use of popularity weighting leads players to adopt and stick with the better tech-

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 in agriculture, respectively.

nology, provided that the technology with the higher mean payoff is also the more likely to have the higher payoff *ex post*. Intuitively, a strategy that 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 it 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, regardless of the initial state of the system or the size of the payoff difference; however, this gain in long-run efficiency may have the cost of slowing the adoption of technological improvements.

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.² 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. 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 either as the result of an informational constraint—players may not observe outcomes at faraway locations—or as the result of the players’ prior belief that faraway locations are sufficiently different that experiences there are not relevant to their own decisions.

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.

The heterogeneous population model provides a number of predictions 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, improving long-run performance while slowing conver-

² 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.

gence. In contrast to the homogeneous-population model, no amount of popularity weighting will lead to an exactly efficient long-run state.

The assumptions of our models are perhaps most descriptive of the diffusion of agricultural technologies. We would expect that farmers are able to observe, at least roughly, the output of their neighbors as well as their neighbors' choices of crops and techniques. Further, farmers' payoffs are subject to aggregate stochastic shocks due to the weather. Concerning inertia, it has been frequently noted that farmers as a group seem very hesitant to try new technologies. These comments do not suggest that all farmers are equally hesitant; for example, Slicher von Bath (1963, 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." Even during the twentieth century, there is typically a substantial lag between the date farmers first learn of the existence of a technology and the date they adopt it: Ryan and Gross (1943) found that farmers in two rural communities on average adopted hybrid seed corn 7 years after they first heard of the innovation, with adoption spread over a 5–10-year period; 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.³ Finally, when capital and insurance markets are poorly developed, it seems plausible that farmers' technology decisions will be determined primarily by short-term considerations, and that farmers will be unlikely to experiment with a technology with a lower expected return.

As for the assumptions of the heterogeneous model, it seems plausible that the payoffs to various crops may differ at different locations, depending on the soil, climate, and terrain of each farm, and that when a new technology is first introduced there may be considerable uncertainty about where it should be used. Consequently, the fact that the technology did well in one area or country may not provide much reason to adopt it in another.⁴

These various features seem particularly clear in the diffusion of the agricultural practices known as the "new husbandry" during the English agricultural revolution. (The new husbandry refers to a vari-

³ 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 they do not seem to have eliminated the lag between becoming informed and deciding to adopt.

⁴ 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 since the villagers decided that the taste and consistency of the corn were inappropriate for making tortillas.

ety of new crops and new crop rotations that arrived in England from Flanders in the seventeenth century, based on the idea of growing crops such as clover or turnips instead of leaving the land fallow; see, e.g., Kerridge [1967], Timmer [1969], and Mingay [1977].) In particular, the new husbandry was attempted and then abandoned at a number of locations, which shows both that the returns to the technology varied with location and that the form of this dependence was not known.

The spatial structure of the heterogeneous-population model can be taken literally when the model is applied to agricultural innovations. We believe that the model can also be applied to other settings, such as the diffusion of management practices in large industries or the choice of private schools; the location variable should be interpreted as a point in characteristic space, and individuals with similar characteristics believe that their payoffs are similar.

Turning from the assumptions of the models to the conclusions, we note that the homogeneous-population model predicts that the speed with which a new technology is adopted is correlated with the extent of the payoff difference. Such a correlation between the extent of improvement and the speed of adoption has been noted in the empirical discussions of Mansfield (1968) and Rogers and Shoemaker (1971), but has not, as far as we know, been addressed in the learning literature.⁵ The homogeneous-population model also predicts that new technologies that result in a small probability of a big improvement and a large probability of a small loss will be adopted slowly, if at all; this is consistent with the slow diffusion of seat belts and vaccinations noted by Rogers and Shoemaker (1971, p. 139). Further, the fact that the combination of inertia and popularity weighting can lead to efficient long-run behavior in the homogeneous-population model may make the apparent occurrence of these phenomena less puzzling. Likewise, in the heterogeneous-population model, the parameter values that favor long-run efficiency—namely, small window widths and high popularity weights—lead to slow diffusion. This may help to explain the observed slow diffusion of some agricultural technologies.⁶

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

⁵ 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.

⁶ The most striking case of slow diffusion may be that of the new husbandry, which diffused at a rate of 1 mile per year both in England and in France.

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, for example by Griliches (1957) and Mansfield (1968).

Recent papers by Banerjee (1992, in press), Bikhchandani, Hirshleifer, and Welch (1992), and Smith (1992) study more sophisticated models of social learning in homogeneous populations, in which players must decide which of two choices is better. The primary question of interest in these models is whether social learning implies that the population eventually identifies the better choice. These papers suppose that players observe one another’s choices, but that players do not observe the payoffs that these choices generate.⁷ Manski (1990) considers estimation procedures for an individual agent in a heterogeneous population; he does not analyze the resulting social dynamics.

Although we believe that the models we develop, based on bounded rationality and players’ observance of their opponents’ payoffs, are a useful supplement to this previous work on social learning, we should say that we are not completely satisfied with the precise form of the rules we consider. In particular, in the first model, use of history does not seem so complicated as to be unreasonable.⁸ 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.

II. A Simple Model of Homogeneous Populations

Before we consider 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.⁹ We suppose that the payoffs to the two technologies at date t , u_t^f and u_t^g , are related by the equation

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

⁷ Cross (1983) develops a model of boundedly rational adaptive choice with a similar information structure.

⁸ 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 Sec. III.

⁹ Given our assumption that players observe one another’s payoffs, nothing would be changed if we allowed each player’s payoff to be subject to idiosyncratic shocks.

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 . We shall assume that $p \equiv 1 - H(-\theta) = \text{prob}[u_t^g - u_t^f \geq 0]$ is strictly between zero and one.

In the initial period, denoted zero, a fraction x_0 of the players are using technology g . After each period, a fraction α of the players are selected at random to have the opportunity to revise their choices.¹⁰ We suppose that the players who are revising their choices can observe the average payoffs of both technologies in the previous period. The simplest behavior rule we consider is the "unweighted" rule under which all players who revise their choice pick the technology that did best in the preceding period. Under this adjustment rule, the evolution of the system is described by

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

so that

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

Note that players treat the adoption and discontinuance decisions symmetrically, which corresponds to the case in which the costs of "transition" are small. This symmetry is probably extreme, but we think that it may be preferable to the standard practice in modeling technology diffusion, which supposes that once agents try the new innovation they continue using it forever: Studies of the English agricultural revolution, as well as studies of more recent innovations cited in Rogers and Shoemaker (1971, p. 115), suggest that the amount of discontinuance is an important factor in the diffusion process.

Our model supposes that 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. Also, the private gain from using history may be small in the cases, detailed below, in which the system without history converges to the efficient outcome.

The following result is standard; it follows from, for example, theo-

¹⁰ As mentioned in the Introduction, this inertia is consistent with the empirical evidence that there is often a substantial lag between the time individuals first learn of the existence of a new technology and the time they adopt it. The inertia might come from decision costs; it would also arise if the choice of a technology is embodied in a costly capital good that will not be replaced until it wears out.

rem 10 of Norman (1968). (It is also a consequence of part *b* of proposition 2 below.)

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

III. 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 in which g has been the better choice. If the distribution H of ϵ is symmetric or nearly so, the technology that is more often better has the higher expected payoff.¹¹ 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 the 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.

Of course, if all players consider popularity in guiding their choices, the level of popularity becomes less informative, and in the extreme case in which players consider only popularity, the popularity conveys no information at all. This leads us to consider the behavior of the system when all players give popularity an intermediate weight. As we shall see, there are popularity weights that, if used by all players, lead them all to adopt the better technology. One way to interpret this result is that in this case, even when popularity weighting is used by all players, the popularity remains a “sufficient statistic” for the history.

To explore the idea of popularity weighting, we develop a simple parametric model. As above, we consider a continuum of players and suppose that only a fraction α of them update their choices each period. Now, though, instead of choosing the technology that did best last period, the choice rule is

$$\text{choose } g \text{ if } u_t^g - u_t^f \geq m(1 - 2x_t). \tag{3}$$

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

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

¹¹ This conclusion holds provided that the degree of asymmetry of H is small compared to $|\theta|$.

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 choose 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.¹² Note that the expression $1 - 2x$ is unitless, so the parameter m is measured in the same units as the payoffs are. Thus 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.¹³

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 shocks per period 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 $m = \sigma$.

To analyze the dynamics of the system, we first identify situations in which it is certain to converge. 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 using f is certain to increase. (Note that $\sigma > 0$ implies $\bar{x}^f < \underline{x}^g$.) Because the probability of an upward step is minimized at $x_t = 0$, this probability must be at least $\text{prob}[\theta + \epsilon_t \geq m] = (\sigma - m + \theta)/2\sigma = -(m/\sigma)\bar{x}^f$. Thus when $\bar{x}^f < 0$, so that the system cannot "lock on" to downward steps, the probability of an upward step is uniformly bounded away from zero. Similarly, if $\underline{x}^g > 1$, the probability of a downward step is uniformly bounded away from zero.

The discussion above shows that (with knife-edge cases ignored) there are four possibilities for the long-run behavior of the system: (1) 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

¹² The empirical literature suggests that popularity weighting is a factor, but reliable estimates of m are hard to come by. Rogers and Shoemaker (1971, p. 142) say that "many students of peasant life feel" that innovations must be 20–30 percent better to be adopted; they also cite a President's Science Advisory Committee figure of 50–100 percent. From our reading, it is not clear whether these premia reflect popularity weighting or switching costs.

¹³ An alternative explanation of the need for rescaling is to use the fact that the rule $m = \sigma$ yields the optimal long-run decision. Since σ is the standard deviation of the payoff differences per period, rescaling the utility function rescales σ in the same way.

the system converges with probability one to $x_t = 1$. (2) If $\underline{x}^g > 1$ and $\bar{x}^f > 0$, the system converges to $x_t = 0$ from any initial position. (3) If $0 < \bar{x}^f$ and $\underline{x}^g < 1$, the system will converge (with probability one) to zero if $x_0 \leq \bar{x}^f$ and will converge to one 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. (4) If $\bar{x}^f < 0$ and $\underline{x}^g > 1$, the system will not converge to either steady state. Instead, the fraction x_t will continue to fluctuate.

These 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 one to the state in which everyone uses the better technology.
- b) $m > \sigma$ is “overweighting” in that the system converges with probability one to a steady state; which steady state is selected may depend on the initial condition x_0 . More precisely, the system converges to the better technology if $|\theta| \geq m - \sigma$, whereas 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 one with probability one; if $x_0 \leq (m - \sigma - \theta)/2m$, the system converges to zero with probability one. 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 a positive probability.
- c) With “underweighting,” that is, $m < \sigma$, the system need not converge to a steady state. It does converge (with probability one) to the better technology if $|\theta| \geq \sigma - m$, but for $|\theta| < \sigma - m$, the system has a nondegenerate invariant distribution μ , with

$$E_\mu x = \frac{1}{2} + \frac{\theta}{2(\sigma - m)}$$

and

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

Proof. (a) If $m = \sigma$, then $\underline{x}^g = (2m - \theta)/2m$ is less than one 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 for $\theta > m - \sigma > 0$, $\bar{x}^f < 0$ and $\underline{x}^g < 1$; for $-\theta > m - \sigma > 0$, $\bar{x}^f > 0$ and $\underline{x}^g > 1$; and 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 distri-

bution when $\sigma - m > |\theta|$ and computes the corresponding mean and variance. Q.E.D.

COROLLARY. For any prior distribution over θ with a finite expected value whose support contains $[-\sigma, \sigma]$, the expected long-run payoff is continuous and increasing in m on the interval $[0, \sigma]$.

Proof. Increasing m increases the set of θ 's for which the system converges to the better choice and increases the "tilt" of the mean toward the better choice for smaller θ 's. Q.E.D.

Proposition 2 shows that the system is certain to converge to the correct choice if the popularity weight $m = \sigma$ 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, our working paper (Ellison and Fudenberg 1992) considers a game in which players simultaneously choose their individual popularity weights, and it shows 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. We conjecture that optimal popularity weighting might emerge from an adaptive process because individuals have a private incentive to increase m whenever $m < \sigma$.

The reader may be concerned that the results in proposition 2 seem to rely on the fact that the uniform distribution has compact support: an observation that $u_i^g - u_i^f > \sigma$ implies that $\theta > 0$. However, similar conclusions can be obtained without compact support. Appendix A shows that the nonlinear rule "switch only 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." This nonlinear rule has the additional advantage that it does not depend on the exact form of the distribution of the noise, whereas the optimal linear rule for the uniform distribution must be tailored to the distribution's support. Appendix A also reports simulations of a more complex rule that seems to work well even when the tails are not infinitely revealing.

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 in which 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 $\theta + \epsilon_t > m(1 - 2x_t)$; since the probability of this event increases with θ , so does the expected rate of adoption.¹⁴ As we noted in the Introduction, such a correlation between the extent of improvement and the speed of adoption has been noted in empirical work.

Note also that for fixed θ , the speed of convergence decreases as σ increases, so that each period's observation becomes less informative. Furthermore, if the new technology usually does slightly *worse* 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 toward the wrong choice. This is consistent with the observation that seat belts, insurance, and vaccinations have been slow to diffuse.

Finally, before leaving the homogeneous-population framework, 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. We now suppose that players consider "trends" in the relative popularity of the two technologies as well as 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 in which the system converges to a steady state, the system still converges to the better technology with probability one when $m = \sigma$. However, if the initial state is far from the optimum, as is 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 three simulations, each 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 $.5\sigma$; in the second, $\alpha = .5$ and $\theta = .1\sigma$; in the third, $\alpha = .1$ and $\theta = .02\sigma$. In all 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 1, show that trends can improve the speed of convergence.

¹⁴ 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 $m = 0$, provided that θ is smaller than σ .

TABLE 1
TREND WEIGHTING AND THE SPEED OF CONVERGENCE

	$\alpha = .5, \theta = .5\sigma$	$\alpha = .5, \theta = .1\sigma$	$\alpha = .1, \theta = .02\sigma$
$c = 0$	11	39	940
$c = 5$	10	26	710
$c = 10$	10	28	470

NOTE.—On the basis of estimated standard errors, the first two digits are correct at the .95 level.

IV. Heterogeneous Populations 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, $E(u_i^g - u_i^f)$, 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 cutoff or break point at $\theta = 0$.

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

$$\begin{aligned} u_i^g(\theta) &= \theta + \beta\theta + \epsilon_{1i}, \\ u_i^f(\theta) &= \beta\theta + \epsilon_{2i}. \end{aligned} \tag{5}$$

With this parameterization, $\beta > 0$ implies that technology g does better at "good" locations, and $\beta < 0$ implies that g does better at bad ones. The player's location in θ -space determines his average payoff to the two technologies. We suppose that players base their decisions on the average performance of the two technologies at locations in their "observation windows"; the observation window of the player at θ is the interval $[\theta - w, \theta + w]$. We call w the "window width."

As in the study of a homogeneous population, we begin by analyzing the simple rule in which players use whichever technology did better in their window last period; later we shall 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 one, and let $\bar{u}_i^g(\theta)$ 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 $\bar{u}_i^g(\theta) = -\infty$ if every player in the interval used f ; the average $\bar{u}_i^f(\theta)$ is defined analogously.

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

$$\text{play } g \text{ at period } t + 1 \text{ iff } \bar{u}_i^g(\theta) - \bar{u}_i^f(\theta) \geq 0. \quad (6)$$

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 shall suppose that there is no inertia at individual locations, so that *all* players at each location revise their choices each period. We do so in part for reasons of convenience and in part because learning seems too easy when all players observe the payoffs to both technologies in every period. Moreover, it seems plausible that in rural areas with low population density a technology could be abandoned by everyone in an observation window after a few bad draws in succession.

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. 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. This is why we do not allow the players to regress the observed payoffs of each technology on the corresponding values of θ . When one is studying geographic diffusion, the observation window might reflect the player only observing the outputs of his neighbors, and the window width w might be fairly small. When one is 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.¹⁵

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 deter-

¹⁵ 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.

ministic. Suppose further that the current state of the system is described by a *cutoff* 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 cutoff 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; similarly, 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

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

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

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

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 cutoff rule.

Our analysis restricts attention to the evolution of these cutoffs. Given our assumption that the payoff difference between the two technologies is monotone in location, cutoff rules seem natural; one might suspect that even if the initial state is not a cutoff rule, the system will converge to one. However, we have not attempted to verify this result because in the absence of a cutoff rule the system is much harder to analyze.

A steady-state cutoff rule must have the property that the player at the steady-state cutoff is indifferent between f and g given his observations. Thus the steady state is the unique solution of $\bar{u}_t^g(\theta^*) = \bar{u}_t^f(\theta^*)$. This gives $(\beta + 1)[\theta^* + (w/2)] = \beta[\theta^* - (w/2)]$ and thus

$$\theta^* = \frac{-(2\beta + 1)w}{2}.\tag{8}$$

Note that although the optimal cutoff is $\hat{\theta} = 0$ for any value of β , the steady-state cutoff equals zero only if $\beta = -1/2$. When $\beta = 0$, for example, so that the payoff to f is identically zero and the payoff to g is equal to θ , the steady state occurs at $-w/2$ (see fig. 1). 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




FIG. 1

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 cutoff, we next examine the behavior of the system away from the steady state. It is easy to show that, from an initial cutoff $\hat{\theta}_0$, the cutoff will move toward 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 two-period cycle about θ^* . For ease of reference, we summarize this as a proposition.

PROPOSITION 3. From an initial cutoff $\hat{\theta}_0$, the system determined by (6) and (7) evolves according to

$$\hat{\theta}_{t+1} = \begin{cases} \hat{\theta}_t + w & \hat{\theta}_t < \theta^* - \frac{w}{2}, \\ -\hat{\theta}_t + 2\theta^* & \hat{\theta}_t \in \left[\theta^* - \frac{w}{2}, \theta^* + \frac{w}{2} \right), \\ \hat{\theta}_t - w & \hat{\theta}_t \geq \theta^* + \frac{w}{2}. \end{cases} \quad (9)$$

Proof. If $\bar{u}_t^g(\hat{\theta}_t - w) - \bar{u}_t^f(\hat{\theta}_t - w) > 0$, 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 $\theta_t = \hat{\theta}_t - w$ 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 $\hat{\theta}_{t+1} = -\hat{\theta}_t - (2\beta + 1)w = -\hat{\theta}_t + 2\theta^*$. Q.E.D.

Next we consider the behavior of the model with noise, that is, with ϵ_{1t} and ϵ_{2t} nondegenerate i.i.d. random variables. Let $z_t = \epsilon_{2t} - \epsilon_{1t}$ denote the difference in the two shocks, and let $\theta_t^* = \theta^* + z_t$; θ_t^* is the steady state of the system when $\epsilon_{2\tau} - \epsilon_{1\tau}$ 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^* .

PROPOSITION 4. If the period t cutoff is $\hat{\theta}_t$ and the period t shock is z_t , the period $t + 1$ cutoff is given by

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

Proof. 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 $[\theta - w, \theta + w]$), that is, $\bar{u}_t^g(\theta, \epsilon_{1t}) - \bar{u}_t^f(\theta, \epsilon_{2t})$, is

$$\frac{\theta + \hat{\theta}_t + (2\beta + 1)w}{2} - z_t = \frac{\theta + \hat{\theta}_t - 2\theta_t^*}{2}.$$

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 cutoff is given by $\hat{\theta}_{t+1} = -\hat{\theta}_t + 2\theta_t^*$. Q.E.D.

PROPOSITION 5. 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 .

Proof. 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).

TABLE 2
STEADY-STATE VARIANCE FOR UNIFORM NOISE

w/σ	VARIANCE OF	
	System (10)	System (11)
.5	.2134	.25
.1	.0483	.05
.05	.0246	.025
.01	.0050	.005

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

$$\hat{\theta}_{t+1} = \begin{cases} \hat{\theta}_t + w & \hat{\theta}_t \leq \theta_t^*, \\ \hat{\theta}_t - w & \hat{\theta}_t > \theta_t^*. \end{cases} \tag{11}$$

Note that system (11) differs from (10) only when $\hat{\theta}_t$ falls in an interval of width w . Normally we 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 $\theta_0 \pm kw$, and so different initial conditions lead to different invariant distributions. Moreover, the supports 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 2 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 6 below. 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 i.i.d. with mean zero and cumulative distribution function 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 grid $\{\theta_0 \pm kw\}$ contains the points θ^* , $\theta^* - \sigma$, and $\theta^* + \sigma$.

PROPOSITION 6. Suppose that 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 $\text{prob}(\theta = \theta^* \pm kw) = [2M!/(M-k)!(M+k)!]2^{-2M}$; this is the limit of the time average distribution when the initial condition belongs to the grid $\theta^* \pm kw$.

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

Proof. 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 shall verify that

$$f(\theta)\text{prob}(\theta_{t+1} = \theta' | \theta_t = \theta) = f(\theta')\text{prob}(\theta_{t+1} = \theta | \theta_t = \theta')$$

or, equivalently, that

$$\frac{f(\theta)}{f(\theta')} = \frac{\text{prob}(\theta_{t+1} = \theta | \theta_t = \theta')}{\text{prob}(\theta_{t+1} = \theta' | \theta_t = \theta)}$$

Since the probability of a jump of more than w is zero, it suffices to check that this condition holds 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

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

and

$$\frac{\text{prob}(\theta_{t+1} = \theta | \theta_t = \theta')}{\text{prob}(\theta_{t+1} = \theta' | \theta_t = \theta)} = \frac{[\sigma + (k+1)w]/2\sigma}{(\sigma - kw)/2\sigma} = \frac{(M+k+1)w}{(M-k)w},$$

so detailed balance holds. Q.E.D.

As one would expect, the variance of the steady state is increasing 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 cutoff is $\hat{\theta}_t$ is

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

Hence, in the long run the average welfare loss per period (from the invariant distribution computed in proposition 6) is

$$\frac{1}{2}E(\theta^2) = \frac{1}{2}[E(\theta)]^2 + \frac{1}{2}\text{var}(\theta) = \left[\frac{(2\beta + 1)^2 w}{8} + \frac{\sigma}{4} \right] w,$$

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.¹⁶

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 taken 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 zero. If the social planner does not know the initial condition or the location of the social optimum, the size of the optimal w will depend on the planner's prior beliefs. This trade-off between speed of adjustment and the variance of the steady state seems a natural feature of the sorts of model we consider.¹⁷

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

$$\begin{aligned} u_t^g(\theta) &= \theta + \beta\theta + \epsilon_{1t} + \eta_t, \\ u_t^f(\theta) &= \beta\theta + \epsilon_{2t} + \eta_t. \end{aligned} \tag{5'}$$

¹⁶ 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. As noted earlier, Ryan and Gross (1943) found that the experiences of neighbors were an important factor in the adoption of hybrid seed corn by twentieth-century Iowa farmers.

¹⁷ 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 faster convergence with a small long-run variance.

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 helpful. Players at locations far from θ^* 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. For these players the systematic payoff difference between the technologies is smaller, and hence it may require many observations to be fairly confident that 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 on the order of $(\sigma/w)^{3/2}$; 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.

V. Examples of Nonlinear 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, and $g(\theta) = \cos(\theta)$, so that regions in which g is optimal alternate with regions in which f is. 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 in which it is optimal. In this example there are substantial social gains from having the new technology “tested” at a number of diverse locations. It may also be interesting to note that when the local 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 fig. 2). 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, since players to the left of $\pi/2$ will tend to switch to g and those to the right switch back to f . Essentially, the noise leads the players in region II




FIG. 2

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 3, in which $f(\theta) = \theta$ and $g(\theta) = \theta - \epsilon$. If the current cutoff occurs at $\hat{\theta}$, then the player at $\theta \in [\hat{\theta} - w, \hat{\theta} + w]$ computes $\bar{u}^g(\theta) = \hat{\theta} - \epsilon + \{[\theta - (\hat{\theta} - w)]/2\}$, and $\bar{u}^f(\theta) = \hat{\theta} - w + \{[\theta - (\hat{\theta} - w)]/2\}$. Since $\bar{u}^g(\theta) - \bar{u}^f(\theta)$




FIG. 3

= $w - \epsilon$, if $w > \epsilon$ 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 drive out a better one only if the difference in payoffs $|f - g|$ is small compared to the errors caused by estimating the payoffs by their average values in the window. These errors have magnitudes 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; 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 in which it is a substantial improvement. This was the case in the example in which $g = \cos(\theta)$ and $f = 0$, so that the regions in which g should be adopted are disconnected. We can also modify the example of figure 3 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 4, 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




FIG. 4

at these extreme values, it will be driven out of the population before it can be tried in the center region.

VI. 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. We shall 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 decision rule (6) used in Sections IV and V and suppose that players use the decision rule

$$\text{play } g \text{ at period } t + 1 \text{ iff } \bar{u}_t^g(\theta) - \bar{u}_t^f(\theta) \geq m[1 - 2x_t(\theta)], \quad (12)$$

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 shall give the results without proof. As in Section IV, if the state in period t corresponds to a cutoff rule, so will the state in period $t + 1$. In addition, without noise terms the system has the same, unique, steady-state cutoff $\theta^* = -(2\beta + 1)w/2$. 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 cutoff from any initial cutoff; the oscillations described in proposition 3 do not arise. Second (and relatedly), movements of less than one window width become more common since players are more hesitant to use a less popular technology.

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

PROPOSITION 7. From an initial cutoff $\hat{\theta}_0$, the system described by decision rule (12) and payoffs (5) evolves according to

$$\hat{\theta}_{t+1} = \begin{cases} \hat{\theta}_t + w & \text{if } \hat{\theta}_t < \theta_t^* - \left(m + \frac{w}{2}\right), \\ \theta_t^* + \left[\frac{2m - w}{2m + w}(\hat{\theta}_t - \theta_t^*)\right] & \text{if } \hat{\theta}_t \in \left[\theta_t^* - \left(m + \frac{w}{2}\right), \theta_t^* + \left(m + \frac{w}{2}\right)\right], \\ \hat{\theta}_t - w & \text{if } \hat{\theta}_t > \theta_t^* + \left(m + \frac{w}{2}\right). \end{cases} \quad (13)$$

Proof. The proof is omitted. The calculations involved are straightforward and quite similar to those of proposition 3. Note that the dynamics above reduce to those of proposition 3 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 cutoff moves a full window width as 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 width $2m + w$. This shows 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 8.

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

Proof. The proof of part *a* is omitted; the argument is very close to that for proposition 5. For part *b*, Appendix D shows that there is a deterministic, finite time T for which the cutoff $\hat{\theta}_T$ 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

$$|\hat{\theta}_{T+s} - \theta_{T+s}^*| < |\hat{\theta}_{T+s} - \theta^*| + |\theta_{T+s}^* - \theta^*| < \left(\sigma + \frac{w}{2}\right) + \sigma,$$

which is less than $m + (w/2)$ from our assumption that $m > 2\sigma$. Hence, the evolution of $\hat{\theta}_{T+s}$ from T on is determined by the second case in proposition 7. Writing $c = (2m - w)/(2m + w)$ and applying this rule repeatedly, we find that

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

Hence,

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

and

$$\text{var}(\hat{\theta}_{T+s} | \hat{\theta}_T) = (1 - c)^2 \sum_{\tau=0}^{s-1} c^{2\tau} \text{var}(\theta^*) \rightarrow \frac{(1 - c)\sigma^2}{3(1 + c)} = \frac{2w\sigma^2}{12m} = \frac{\sigma^2 w}{6m}.$$

Q.E.D.

Comparing the steady-state distributions for $m > 2\sigma$ with that for no popularity weighting, we see that popularity weighting reduces the long-run variance by a factor of $\sigma/3m$.

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, whereas 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.¹⁸

VII. 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.

We discussed one extension, the use of trends, in Section III. There are a number of other extensions that we have not considered but that seem important. Players might use rules of thumb that make some use of historical data. Also, players might be arranged in more complex networks than the simple linear structure we have considered. In addition, our results suppose that rules of thumb are exogenous. It would be interesting to complement these results with an

¹⁸ 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.

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 in 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 window, 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.¹⁹

Appendix A

Optimal Popularity Weighting with Other Distributions

To better understand the forces generating part *a* of proposition 2—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 zero and one. If we let \mathbf{s}_t denote the vector [$\text{prob}(x_t) = 0$, $\text{prob}(x_t) = 1$], we have $\mathbf{s}_{t+1} = \mathbf{s}_t \mathbf{A}$, where the transition matrix is

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

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

$$\mu^*(x = 0) = \frac{F(-m - \theta)}{F(-m - \theta) + 1 - F(m - \theta)}.$$

If F is the standard normal distribution, then as m increases, the ratio $F(-m - \theta)/[1 - F(m - \theta)]$ converges to zero if $\theta > 0$ and converges to infinity if $\theta < 0$. Hence for large m , the ergodic distribution of the system places probability near one on the correct choice. Moreover, the same is true for any distribution for which the ratio $F(-m - \theta)/[1 - F(m - \theta)]$ converges to zero if $\theta > 0$ and to infinity 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; the intuition for the result follows.

Note first that when $m = \infty$ the system is deterministic with stable steady states at zero and one. If m is finite but very large compared to α and to the standard deviation of the distribution, then steps the “wrong way” (i.e., decreasing steps when $x_t > 1/2$) are rare “innovations,” and when the distribu-

¹⁹ We thank Roland Benabou for this observation.

TABLE A1
 $\text{prob}[x_t \rightarrow 1 | \theta = .2\sigma, x_0 = .2\sigma]$

Distribution	$\alpha = .2$	$\alpha = .1$	$\alpha = .05$
Normal	.53	.77	.95
Laplace	.55	.79	.94
Logistic	.53	.78	.96
Standard error	.0005	.0004	.0002

tion is symmetric, transits from zero to one-half and from one to one-half both take the same number of innovations. If the tails of the distribution are infinitely revealing, then as $m \rightarrow \infty$ innovations toward the better technology become infinitely more likely than innovations toward the inferior one. The analysis of Freidlin and Wentzell (1984) suggests that the limit of the ergodic distributions will be concentrated on the better technology. To establish this formally, we partition the 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 1993; Young 1993) to confirm the intuition above; that is, the limits of the ergodic distributions of the finite-state process are concentrated on the subinterval corresponding to the better choice.

The discussion 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

$$\text{choose } g \text{ iff } u_t^g - u_t^f \geq F^{-1}(1 - x_t). \tag{A1}$$

With this rule,

$$\begin{aligned} E(x_{t+1} | x_t) &= (1 - \alpha)x_t + \alpha \text{prob}[\theta + \epsilon_t \geq F^{-1}(1 - x_t)] \\ &= x_t + \alpha[1 - F(-\theta + F^{-1}(1 - x_t)) - x_t], \end{aligned}$$

so that $E(x_{t+1} | x_t) > x_t$ if and only if $\theta > 0$; the system drifts toward the correct choice. Although the system may converge to the wrong technology with positive probability, simulations for the logistic and Laplace distributions (which both have nonrevealing tails) suggest that when α is small the system is very likely to converge to the right choice. Table A1 displays one set of simulations, for the case $\theta = .2\sigma, x_0 = .2\sigma$. The table suggests that the behavior for all three distributions is similar, even though the latter two do not have infinitely revealing tails. 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 along these lines.

Appendix B

Proof of Part c of Proposition 2

If $\sigma - m > |\theta|$, then neither zero nor one 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).²⁰ 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 that the system starting at x' is in U exactly n periods later is strictly positive. (If $m \geq \sigma$, the uniqueness condition fails since 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 $\theta + \epsilon_t \geq m(1 - 2x_t)$, 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

$$\begin{aligned} E_\mu(x^2) &= \int \{[1 - p(x)][(1 - \alpha)x]^2 + p(x)[(1 - \alpha)x + \alpha]^2\} d\mu(x) \\ &= E_\mu(x^2) \left[(1 - \alpha)^2 + \frac{2\alpha(1 - \alpha)m}{\sigma} \right] \\ &\quad + E_\mu(x) \left[\frac{2\alpha(1 - \alpha)(\sigma - m + \theta)}{2\sigma} + \frac{\alpha^2 m}{\sigma} \right] + \frac{\alpha^2(\sigma - m + \theta)}{2\sigma}; \end{aligned}$$

solving for $E_\mu(x^2)$ and computing $\text{var}(x) = E_\mu(x^2) - [E_\mu(x)]^2$ give the desired result. Q.E.D.

Appendix C

Proof of Proposition 5

To begin we rewrite (10) in the following equivalent form:

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

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

²⁰ See Futia's (1982) survey for a summary of Norman's results and other techniques for establishing that the invariant distribution is unique.

distribution Q on the z 's does not depend on the current state, and so in particular it 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 zero. 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 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: (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 $\theta^* + z$ 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 $\theta^* + z$, and $d(\varphi(\theta, z), \varphi(\theta', z)) = d(\theta, \theta')$.

In case *b*, suppose without loss of generality that $\theta < \theta'$; then case *b* implies that $\theta \leq \theta^* + z \leq \theta'$, and so $d(\theta, \theta') = d(\theta, \theta^* + z) + d(\theta^* + z, \theta')$. With the triangle inequality, this implies that

$$\begin{aligned} d(\varphi(\theta, z), \varphi(\theta', z)) - d(\theta, \theta') &\leq d(\varphi(\theta, z), \theta^* + z) + d(\theta^* + z, \varphi(\theta', z)) \\ &\quad - d(\theta, \theta^* + z) - d(\theta^* + z, \theta') \\ &= [d(\varphi(\theta, z), \theta^* + z) - d(\theta, \theta^* + z)] \\ &\quad + [d(\theta^* + z, \varphi(\theta', z)) - d(\theta^* + z, \theta')], \end{aligned} \tag{C1}$$

and inspection of (10') shows that each of the terms in brackets is nonpositive. 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 a 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 $\epsilon > 0$, there is a positive probability that z lies in any sufficiently small neighborhood of $\theta - \theta^* + \epsilon - (w/2)$, and for z 's in this neighborhood, θ moves less than w to the left and θ' moves w , so that $d(\varphi(\theta, z), \varphi(\theta', z)) < d(\theta, \theta')$. If $\theta - \theta^* \leq -\sigma + (w/2)$ 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 $\theta - \theta^* \leq -\sigma + (w/2)$ and $\theta' - \theta^* \geq \sigma - (w/2)$, 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^0 such that, for any neighborhood U of θ^0 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, for example, $\theta^0 = \theta^*$ satisfies this condition. Q.E.D.

Appendix D

Proof of Part *b* of Proposition 8

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) $|\hat{\theta}_{T+s} - \theta^*| < \sigma + (w/2)$ for all subsequent dates $T + s$. Define $d_t = |\hat{\theta}_t - \theta^*|$. Note that since $(\theta_t^* -$

$\hat{\theta}_i$) and $(\hat{\theta}_{i+1} - \hat{\theta}_i)$ have the same sign and $|\theta_i^* - \theta^*| \leq \sigma$, $(\theta^* - \hat{\theta}_i)$ and $(\hat{\theta}_{i+1} - \hat{\theta}_i)$ have the same sign whenever $d_i > \sigma + (w/2)$. Hence,

$$d_{i+1} = |d_i - (|\hat{\theta}_{i+1} - \hat{\theta}_i|)| \quad (D1)$$

whenever $d_i > \sigma + (w/2)$.

As a first step toward proving claim i, we show that for any initial condition there exists a finite T' such that, regardless of the sample path, *either* $d_{T'} < \sigma + (w/2)$ or $d_{T'} \leq |\hat{\theta}_{T'+1} - \hat{\theta}_{T'}|$. To see this, note that (D1) implies that until such a T' is reached, $d_i - d_{i+1} = |\hat{\theta}_{i+1} - \hat{\theta}_i|$, and from proposition 7,

$$|\hat{\theta}_{i+1} - \hat{\theta}_i| = \min \left\{ w, \left(\frac{2w}{2m+w} \right) (\hat{\theta}_i - \theta_i^*) \right\} \geq \min \left\{ w, \left(\frac{2w}{2m+w} \right) \left(\frac{w}{2} \right) \right\}.$$

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

If $d_{T'} \leq \sigma + (w/2)$, setting $T = T'$ completes the proof of claim i. The remaining case is $\sigma + (w/2) < d_{T'} \leq |\hat{\theta}_{T'+1} - \hat{\theta}_{T'}|$. 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 claim i.

To prove claim ii, note that when $|\hat{\theta}_T - \theta^*| < \sigma + (w/2)$, we have $|\hat{\theta}_T - \theta_T^*| \leq [\sigma + (w/2)] + \sigma$, which is less than $m + (w/2)$ from the assumption that $m > 2\sigma$. From proposition 7 we then have

$$\hat{\theta}_{T+1} = \theta_T^* + \left(\frac{2m-w}{2m+w} \right) (\hat{\theta}_T - \theta_T^*),$$

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 . Q.E.D.

Appendix E

This appendix gives a rough approximation of how many periods a player using the entire history of observations would need to identify the better technology with a confidence level of 85 percent. Suppose that $\epsilon_t = \epsilon_{1t} - \epsilon_{2t}$ is uniform on $[-\sigma, \sigma]$, so that $\text{var}(u_t^1 - u_t^2) = \sigma^2/3$. Then the player at location θ will need about $\sigma^2/3\theta^2$ observations of the payoff difference to have an 85 percent confidence level. (Recall that we have assumed that only observations of the payoff difference are used since the levels are subject to a common shock.)

Now the player at location θ will not observe both technologies in every period; he observes both at t only if $\hat{\theta}_t \in [\theta - w, \theta + w]$. We approximate the distribution of $\hat{\theta}_t$ by an $N(0, \sigma w/2)$ random variable and approximate the probability of the event $\hat{\theta}_t \in [\theta - w, \theta + w]$ by $2w$ times the density of this variable at θ . The expected wait for $\sigma^2/3\theta^2$ observations of both technologies simultaneously is then approximately

$$\left(\frac{\sigma^2}{3\theta^2} \right) \left(\frac{1}{2w} \right) \left(\frac{2\pi\sigma w}{2} \right)^{1/2} \exp \left(\frac{\theta^2}{\sigma w} \right).$$

This expression is minimized at $\theta = (\sigma w)^{1/2}$, where its value is $(e\pi^{1/2}/6) \times (\sigma/w)^{3/2}$.

References

- Apodaca, Anacleto. "Corn and Custom: Introduction of Hybrid Corn to Spanish American Farmers in New Mexico." In *Human Problems in Technological Change: A Casebook*, edited by Edward H. Spicer. New York: Sage Found., 1952.
- Banerjee, Abhijit. "A Simple Model of Herd Behavior." *Q.J.E.* 107 (August 1992): 797-817.
- . "The Economics of Rumours." *Rev. Econ. Studies* (in press).
- Bikhchandani, Sushil; Hirshleifer, David; and Welch, Ivo. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *J.P.E.* 100 (October 1992): 992-1026.
- Cross, John G. *A Theory of Adaptive Economic Behavior*. Cambridge: Cambridge Univ. Press, 1983.
- Ellison, Glenn, and Fudenberg, Drew. "Rules of Thumb for Social Learning." Working paper. Cambridge: Massachusetts Inst. Tech., 1992.
- Freidlin, Mark, and Wentzell, Alexander. *Random Perturbations of Dynamical Systems*. New York: Springer-Verlag, 1984.
- Futia, Carl A. "Invariant Distributions and the Limiting Behavior of Markovian Economic Models." *Econometrica* 50 (March 1982): 377-408.
- Griliches, Zvi. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica* 25 (October 1957): 501-22.
- Kandori, Michihiro; Mailath, George; and Rob, Rafael. "Learning, Mutation, and Long Run Equilibrium in Games." *Econometrica* 61 (January 1993): 29-56.
- Kerridge, Eric. *The Agricultural Revolution*. London: Allan and Unwin, 1967.
- Mansfield, Edwin. *The Economics of Technological Change*. New York: Norton, 1968.
- Manski, Charles. "Dynamic Choice in a Social Setting." Discussion Paper no. 9003. Madison: Univ. Wisconsin, Soc. Sci. Res. Inst., 1990.
- Mingay, G. E., ed. *The Agricultural Revolution: Changes in Agriculture, 1650-1880*. London: Adam and Charles Black, 1977.
- Norman, M. Frank. "Some Convergence Theorems for Stochastic Learning Models with Distance Diminishing Operators." *J. Math. Psychology* 5 (February 1968): 61-101.
- . *Markov Processes and Learning Models*. New York: Academic Press, 1972.
- Rogers, Everett, and Shoemaker, Floyd. *Communication of Innovations: A Cross-cultural Approach*. 2d ed. New York: Free Press, 1971.
- Ryan, Bryce, and Gross, Neal C. "The Diffusion of Hybrid Seed Corn in Two Iowa Communities." *Rural Sociology* 8 (March 1943): 15-24.
- Slicher van Bath, B. H. *The Agrarian History of Europe, A.D. 500-1850*. London: Arnold, 1963.
- Smith, L. "Error Persistence and Experimental versus Observational Learning." Manuscript. Cambridge: Massachusetts Inst. Tech., 1992.
- Timmer, C. Peter. "The Turnip, the New Husbandry, and the English Agricultural Revolution." *Q.J.E.* 83 (August 1969): 375-95.
- Young, H. Peyton. "The Evolution of Conventions." *Econometrica* 61 (January 1993): 57-84.