

WORD-OF-MOUTH COMMUNICATION AND SOCIAL LEARNING*

GLENN ELLISON AND DREW FUDENBERG

This paper studies the way that word-of-mouth communication aggregates the information of individual agents. We find that the structure of the communication process determines whether all agents end up making identical choices, with less communication making this conformity more likely. Despite the players' naive decision rules and the stochastic decision environment, word-of-mouth communication may lead all players to adopt the action that is on average superior. These socially efficient outcomes tend to occur when each agent samples only a few others.

I. INTRODUCTION

Economic agents must often make decisions without knowing the costs and benefits of the possible choices. Given the frequency with which such situations arise, it is understandable that agents often choose not to perform studies or experiments, but instead rely on whatever information they have obtained via casual word-of-mouth communication. Reliance on this sort of easily obtained information appears to be common in circumstances ranging from consumers choosing restaurants or auto mechanics to business managers evaluating alternative organizational structures.

This paper studies two related environments in arguing that individuals' reliance on word-of-mouth communication has interesting implications for their aggregate behavior. First, motivated by the diffusion of new technologies, we consider a choice between two competing products with unequal qualities or payoffs, and show that the structure of communication is important in determining whether the population as a whole is likely to learn to use the superior product. Second, we consider a choice between two products or practices that are equally good, and ask whether consumers are likely to "herd" onto a single choice, or whether "diversity" will obtain even in the long run.

We explore the implications of word-of-mouth communication in a simple nonstrategic environment. There is a large population of identical players, each of whom repeatedly chooses between two possible actions. Each player's payoff is determined by his own

*We thank Edward Glaeser for asking questions that helped prompt us to study this topic. National Science Foundation Grants SBR 93-10009 and SBR 92-23320 provided financial support.

actions and some stochastic shocks, but is not influenced by the actions chosen by others, so there are no direct or “strategic” reasons that the player’s choice should depend on the choices of others. The only reason the players’ actions are linked is because of the way that information is transmitted. To model the difficulties that might arise when choosing between technologies or methods of organization, we assume that payoffs can only be learned from experience. Further, payoffs in each period are subject to both common shocks (e.g., weather, business conditions, time-varying quality) and idiosyncratic ones. Because of the idiosyncratic component it seems sensible that players would want to collect information from others.

In this environment we examine the implications of exogenously specified rules for behavior. This involves two separate departures from rationality. First, motivated by the idea that the decision problem at hand is but one of many informed by the same set of casual conversations, we do not model the acquisition of information as an optimal search problem. Instead, we simply assume that in each period each player hears of the current experiences of a random sample of N other players and treat N as an exogenous parameter describing possible communication structures. Second, we specify a simple model for boundedly rational choice given the information players acquire. The population displays inertia, with a fraction of the players ignoring the information they receive in a given period (or not inquiring at all) and continuing with their previous choice. The players who do consider switching pick the action that appears to give the highest average payoff given their most recent experience and the reports they have heard in the current period. Moreover, we impose the “must-see” restriction that players who only obtain reports about the choice they are currently using do not contemplate switching, even if all observed payoffs to their current choice are low.¹

These naive rules are intended to capture the essence of word-of-mouth communication. Some economic agents are probably more sophisticated than this, and it would be interesting to explore the implications of more sophisticated, but still boundedly rational behavior rules. However, given the complexity of the

1 We feel that this assumption is plausible on behavioral grounds. It can also be derived as the outcome of rational Bayesian learning in the neighborhood of the efficient state, as in the closely related model of Banerjee and Fudenberg [1994]. Banerjee [1993] offers a different explanation for the must-see property by appealing to unmodeled but unlikely states of the world.

decision problems, we find the extreme simplicity of the behavior we study at least as plausible as the complex behavior required by the standard “fully rational” approach, which requires players to update optimally from a consistent and common prior distribution over the possible structures of payoffs and noise processes.

Our most basic conclusion is that the structure of the word-of-mouth process affects the tendency of a population to display conformity or diversity, with less communication (that is, smaller but nonzero samples for each individual) making conformity more likely. Despite the players’ naive behavior and the complex stochastic environment of our model, we find that word-of-mouth communication may allow efficient social learning in the sense that all players eventually adopt the action which is on average superior. Surprisingly, these socially efficient outcomes tend to occur only when each individual receives very little information.

At first sight it may be surprising that, in contrast to our previous paper on boundedly rational learning [Ellison and Fudenberg 1993], social learning in this model can lead to long-run efficiency in the absence of explicit popularity weighting. The explanation is that the must-see restriction incorporates a form of popularity weighting, as it implies that players tend not to hear about unpopular technologies, and hence allows these to die out. Previous models of word-of-mouth communication include the contagion model, Smallwood and Conlisk’s [1979] model of replacement on breakdown, Satterthwaite’s [1979] model of reputation goods, and Kirman’s [1993] model of search by ants. Our model differs from these papers both in the structure of the environment in which people learn and the form of the players’ behavioral rules. Typically, previous work has assumed that the aggregate performance of each choice is fixed over time and that players always adopt whichever choice they hear about regardless of how its performance compares with their own experience. In these respects, our model is closer to that of our [1993] paper, although the environment is more complex due to the addition of idiosyncratic noise. Indeed, the simplest unweighted single population model of Ellison and Fudenberg may be thought of as the limiting case of the model presented here as the sample size N becomes infinite: word-of-mouth sampling with an infinite sample allows each player to average out the idiosyncratic shocks, but not the common one, and so in the limit the players all choose the technology that had the higher average payoff in the previous period.

In addition to the differences in the decision environment and decision rules, our analysis is also somewhat different from most of the previous literature in its focus on a categorization of the dynamic pattern of the market shares of the two choices. We begin with the case where the average payoffs are identical. Here we distinguish two possibilities. If all consumers eventually make the same choice (formally, if the market share of one choice converges to one), we say that the population exhibits “conformity.” If instead the market shares perpetually fluctuate or settle down to a nontrivial split, we say that the population exhibits diversity. Our main conclusion, presented in Section IV, is that sufficiently limited word-of-mouth communication results in conformity. When players receive more information, i.e., when N is large, the system exhibits diversity.

In the case where the choices have unequal payoffs, we distinguish three possible patterns of behavior. First, as before, the population may exhibit diversity. Second, we may have efficient social learning with everyone adopting the superior choice. Finally, there is “inefficient herding” if everyone eventually adopts the same choice but the common choice is not necessarily optimal.

To understand how these regions arise, it is instructive to think of the superior average payoff as a force pushing all players toward using the better choice. The long-run behavior of the system is then determined by how this force combines with the word-of-mouth process’s “inherent” favoring of conformity or diversity. When the sample size N is large, we show in Section V that the diversity forces are overwhelming, and the first case obtains. When N is smaller, there is efficient social learning, as a not too strong tendency toward diversity and the superior payoff combine to prevent convergence to the inferior extreme while allowing convergence to the superior extreme. For some parameter values, very small N produces inefficient herding. We find, however, that this behavior is rare when significant inertia is present, as it may be that even $N = 1$ is large enough to produce efficient learning.

To assess the robustness of our conclusions, we examine two related models in Section VI. In the first variant, players practice popularity weighting, combining information about payoffs that they obtain from their samples with information about overall market shares. We find that the same basic patterns of behavior occur, although efficient social learning may now require that the players see larger samples. Our second variant supposes that the

common shock to the system changes more frequently than any individual player observes the system, so that when information is obtained about someone's most recent experience it may reflect the payoffs of a previous date. In this way, we think of the players' sampling process as taking advantage of a social memory of past payoffs. This might be expected to give a clearer advantage to the product which is better on average, and possibly make larger samples more efficient. In examining the extreme cases of two period and infinite memories, we find that sampling only one other player is still efficient, but note that learning may be nearly efficient with large sample sizes as well.

II. THE MODEL

We examine word-of-mouth communication and social learning in the context of a population of agents who are faced with a choice between two competing products, technologies, or practices, which we label f and g . There is a continuum of identical agents, and each chooses between the two products at a discrete set of points in time, $t = 1, 2, 3, \dots$. The aggregate behavior of the population at each point in time can be summarized by a "state" variable x_t , giving the fraction of the population who are using technology g . Our model does not attempt to explain the very first adoptions of a new technology. Rather, we take the initial state x_0 to be exogenous, and focus on the social learning reflected in the subsequent patterns of the market share movements.

The specification of the payoffs provided by the products incorporates a number of elements which complicate decision-making in technology adoption problems. We assume that the payoff to consumer i using brand f at date t is $\bar{f}_t + \epsilon_{ift}$, and that i 's payoff to using brand g at date t is $\bar{g}_t + \epsilon_{igt}$. Here the $\epsilon_{i,t}$ represent player-specific idiosyncratic shocks, which could correspond to variations both in the quality of individual units of the product, and its suitability for the consumer in question. (For example, they may reflect both whether a restaurant meal is well cooked, and whether the customer ordered the right entree given the way he is feeling on a particular day.) We suppose that the $\epsilon_{i,t}$ are i.i.d. both over time and across consumers, which implies that there are no persistent differences between individuals.² For computational

2 This involves a standard abuse of the law of large numbers.

convenience we further assume that each ϵ has a normal distribution with mean 0 and standard deviation σ

Common shocks, i.e., factors affecting the relative payoffs of the two products for all consumers, are represented by allowing $\theta_t \equiv \bar{g}_t - \bar{f}_t$ to be random. We assume that the θ_t are also i.i.d. and independent of the individual specific shocks, and that the marginal distribution over θ_t is binomial, with probability p , $0 < p < 1$, that $\theta_t = \theta > 0$, and probability $1 - p$ that $\theta_t = -\theta < 0$.³ The common shocks might reflect the impact of weather conditions on two agricultural technologies, the different effect of changing economic conditions on firms who have chosen one of two business practices, or variation in the quality of the designs that two firms produce in a given model year.

Given the distribution of the shocks, the probability that g is better than f for player i at date t is $p \text{ Prob}(\epsilon_{igt} - \epsilon_{ift} > -\theta) + (1 - p) \text{ Prob}(\epsilon_{igt} - \epsilon_{ift} > \theta)$.

Note that even if $p = 1$, so that at each time t brand g has a higher expected payoff, the probability that g does better than f for a given individual may be less than one due to the idiosyncratic shocks. Note also that this formulation places no restriction on the distributions of the mean payoffs \bar{f}_t and \bar{g}_t , but only on their difference which allows for the possibility that the quality of both products is increasing over time. The individual means will not matter until we consider an extension of the model in subsection VI 2.

In choosing between the two alternatives, the agents in our model are assumed to follow exogenously specified, boundedly rational behavior rules that incorporate the notions that there is inertia in consumer choices, and that consumers use the information gathered by word-of-mouth communication in a simple way, without the explicit considerations of the information they would see in each possible state of the world that is required in a Bayesian Nash equilibrium

Inertia is modeled with the assumption that, at each date $t > 0$, some fraction α , $0 < \alpha < 1$, of the agents decide to reevaluate their choice.⁴ When agents do reevaluate their choice, they ask N other individuals chosen at random from the population distribu-

3 The assumed symmetry between payoffs in the good and bad state reduces the number of cases we must discuss but does not significantly change our qualitative conclusions

4 A more realistic model would specify that consumers are more likely to switch if they are dissatisfied with their current choice, that is, if their current payoff is low

tion about their current choice and current payoff. The sampling agents then compute the average payoff of the two alternatives in their sample, and choose whichever one had the higher average payoff. This extremely simple decision rule ignores both the precise stochastic structure of the payoffs, and the implicit information that may be conveyed by the relative popularity of the two choices. Consequently, it may be more plausible in complex and novel environments.

Several aspects of this rule deserve emphasis. First, a player who is currently using one choice cannot switch to the other one unless he talks to someone who is using it. As remarked above, this implies that the system incorporates a form of “popularity weighting,” where the weight is decreasing in N . Second, players weigh observations of other players’ experience as heavily as they weight their own. Since the individual-specific shocks are assumed to be i.i.d., this behavior is consistent with the model, and the more plausible case where one’s experience is more relevant corresponds to a model with persistent individual shocks. Third, customers use only current information in making decisions. The lack of information on other agents’ past payoffs may be justified by saying that agents simply do not pay attention to others’ casual comments about products they are not considering buying at the time. The agents’ failure to incorporate their own past experience is harder to justify fully, although it is perhaps not unreasonable, for example, for consumers facing a choice between computers or other products for which mean quality levels change so rapidly that it is hard to compare the quality of products purchased at different times. To a degree, we use this extreme assumption simply to capture the idea that boundedly rational consumers will not fully incorporate all historical information.

Our analysis of the model focuses on the dynamic pattern of the market shares of the two products over time. Since the aggregate shock θ_t takes on only two values, and we have supposed that the aggregate distribution of idiosyncratic shocks over the continuum of consumers is deterministic, there are only two possible values of x_{t+1} for each value of x_t . For example, when $N = 1$, the evolution of the system is given by

$$(1) \quad x_{t+1} = \begin{cases} (1 - \alpha)x_t + \alpha[x_t(x_t + (1 - x_t)q^+) + (1 - x_t)x_tq^+] & \text{with prob. } p \\ (1 - \alpha)x_t + \alpha[x_t(x_t + (1 - x_t)q^-) + (1 - x_t)x_tq^-] & \text{with prob. } 1 - p, \end{cases}$$

where q^+ is the probability that a randomly drawn g -user had a higher payoff than a randomly drawn f -user when $\theta_t = \theta$, and q^- is the analogous term when $\theta_t = -\theta$.

To understand the top line of this equation, suppose that $\theta_t = \theta$. Recall first that each period a fraction $(1 - \alpha)$ of the consumers do not consider switching. Since a fraction x_t of these are currently using g , there will be at least $(1 - \alpha)x_t$ consumers using g next period. Of the α consumers who do consider switching, three subsets choose g with positive probability. Those who use g and meet another g -user choose g with probability one, while those who use g and meet an f -user and those who use f and meet a g -user both choose g with probability q^+ . Multiplying these probabilities by the corresponding masses, which are αx_t^2 , $\alpha x_t(1 - x_t)$, and $\alpha x_t(1 - x_t)$, respectively, yields the top line of equation (1); the bottom line is the same with q^- replacing q^+ .

For larger values of N , the system has the same general form, with $x_{t+1} = H^+(x_t)$ when $\theta_t = \theta^+$, and $x_{t+1} = H^-(x_t)$ when $\theta_t = \theta^-$, but the form of the equations H becomes more complicated in two related ways. First, the binomial distribution for the consumer's word-of-mouth sampling is replaced by a multinomial (N, x_t) distribution; e.g., the probability that k of N consumers sampled use g is $\binom{N}{k} x_t^k (1 - x_t)^{(N-k)}$. Second, for each value of k , there are different numbers corresponding to q^+ and q^- , namely the probability that brand g has a higher average payoff than f in a sample of k g 's and $N - k$ f 's.

Consequently, the general form of the dynamics is

$$x_{t+1} = \begin{cases} (1 - \alpha)x_t + \alpha \left(x_t^{N+1} + \sum_{k=1}^N \frac{(N+1)!}{k!(N+1-k)!} x_t^k (1 - x_t)^{N-k} \right. \\ \left. \times \Phi \left(\frac{\theta}{\sigma \sqrt{1/k + 1/(N+1-k)}} \right) \right) & \text{with probability } p \\ (1 - \alpha)x_t + \alpha \left(x_t^{N+1} + \sum_{k=1}^N \frac{(N+1)!}{k!(N+1-k)!} x_t^k (1 - x_t)^{N-k} \right. \\ \left. \times \Phi \left(\frac{-\theta}{\sigma \sqrt{1/k + 1/(N+1-k)}} \right) \right) & \text{with probability } 1 - p, \end{cases}$$

where Φ is the standard normal cumulative distribution function.

III. LOCK-ON AND HERDING: A LOCAL ANALYSIS

We have not been able to completely determine the long-run dynamics of our model. Rather than simplify the model further, we have chosen to provide a partial characterization. In particular, our results concern only the question of whether the system will converge to one of its endpoints, in which case we will say that it exhibits “lock-on” or “herding,” or whether to the contrary there will be long-run diversity in the sense that both choices continue to be used by nonnegligible fractions of the population.

A great practical advantage of this sort of characterization is that it only requires a local analysis of the system in the neighborhood of its endpoints. In this section we temporarily defer any discussion of our word-of-mouth learning model to present a mathematical result which greatly simplifies such analyses. Essentially, the lemma shows that when one is trying to determine whether a member of a class of Markov processes can converge to an extreme point, one can work instead with a linearized version of the system. Given the complexity of our model, the ability to work with a linear approximation will be an advantage.

LEMMA 1. Let x_t be a Markov process on $(0,1)$ with

$$x_{t+1} = \begin{cases} H_1(x_t) & \text{with probability } p \\ H_2(x_t) & \text{with probability } (1-p) \end{cases}.$$

Suppose that in a neighborhood of 0, $H_i(x_t) = \gamma_i x_t + o(x_t)$, with $\gamma_2 < 1 < \gamma_1$.

(a) If

$$\frac{p}{1-p} > -\frac{\log(\gamma_2)}{\log(\gamma_1)},$$

then x_t cannot converge to 0 with positive probability.

(b) If

$$\frac{p}{1-p} < -\frac{\log(\gamma_2)}{\log(\gamma_1)},$$

then there are strictly positive δ and ϵ such that $\text{Prob}[x_t \rightarrow 0 | x_0 \leq \delta] \geq \epsilon$.

(c) If

$$\frac{p}{1-p} > -\frac{\log(\gamma_2)}{\log(\gamma_1)},$$

there is an $x^* > 0$ such that for all $x_0 > 0$, $\text{Prob}[x_t < x^* \forall t | x_0] = 0$.

Remark. The proof is deferred to Appendix A; the following discussion tries to provide an intuition. Since this lemma concerns approximately linear systems, the best way to understand it is to first understand the case of systems that are exactly linear. Such systems are sufficient to make a crucial point: when the step size is proportional to the state, as it is in the models of this paper, the appropriate notion of “negative drift” to use for determining whether the system converges to zero is not the usual supermartingale condition that $E(x_{t+1}|x_t) < x_t$, but rather that the log of the system has negative drift; that is, that $E(\log(x_{t+1})|x_t) < \log(x_t)$. In particular, a system with positive drift in the usual sense can still be certain to converge to zero. This explains the role of the logarithms in the various expressions in our theorems.

To see why this is so, consider the case where x_t is a Markov process on $[0, \infty)$ and the functions H_t are exactly linear, so that $H_t(x) = \gamma_t x$. Let $z_t = \log(x_t)$, and note that if $x_t \rightarrow 0$, then $z_t \rightarrow -\infty$. Moreover, the evolution of z_t is given by

$$z_{t+1} = \begin{cases} z_t + \log(\gamma_1) & \text{with prob } p \\ z_t + \log(\gamma_2) & \text{with prob } 1-p \end{cases}$$

so that z_t is the sum of the i.i.d. increments $\log(\gamma_1)$ and $\log(\gamma_2)$, with probabilities p and $1-p$, respectively. If $p \log(\gamma_1) + (1-p) \log(\gamma_2) = \log(\gamma_1^p \gamma_2^{1-p}) > 0$, then by the strong law of large numbers z_t converges to ∞ with probability 1, and hence x_t converges to ∞ , which is the basis of the proof of (a) and (c). If $p \log(\gamma_1) + (1-p) \log(\gamma_2) < 0$, then by the strong law z_t converges to $-\infty$, and so x_t converges to 0, which is the basis of the proof of (b).

Once the analysis of the linear system is understood, it remains to explain why the linear approximation methods we use are valid. Basically, if the log has negative drift, then not only does the exactly linear system converge to zero, there is a positive probability that it will converge without ever moving to the right of its initial position. Along such paths the error in ignoring terms of higher order in x is bounded, and we can conclude that the approximately linear system converges to zero as well.

To conclude this section, we return to our model to give the linearized form of its dynamics, which is

$$(2) \quad x_{t+1} = \begin{cases} (1 - \alpha + \alpha(N+1)q^+(N))x_t + o(x_t) & \text{with prob. } p \\ (1 - \alpha + \alpha(N+1)q^-(N))x_t + o(x_t) & \text{with prob. } 1-p, \end{cases}$$

where

$$q^+(N) = \Phi \left\{ \frac{\theta}{\sigma\sqrt{(1 + 1/N)}} \right\}$$

is the probability that the payoff to a random individual using g exceeds the average payoff of N individuals using f , conditional on $\theta_t = \theta$, and

$$q^-(N) = \Phi \left\{ \frac{-\theta}{\sigma\sqrt{(1 + 1/N)}} \right\}$$

denotes the same probability conditional on $\theta_t = -\theta$.

To understand these equations intuitively, note that in period $t + 1$ the number of g -users will consist of the $(1 - \alpha)x_t$ who previously used g and did not consider switching, plus a fraction of those who did consider switching. When x_t is near zero, the number of players who considered switching and observed both choices f and g is $\alpha(N + 1)x_t$ to first order in x . (This is the sum of almost all of the g -users who consider switching, who have mass αx_t plus a share Nx_t of the $\alpha(1 - x_t)$ f -users who consider switching.) The sample of each of these players most likely consists of one g -user and N f -users. With such a sample, the probability that the average payoff to g is higher than the average payoff to f is $q^+(N)$ or $q^-(N)$ depending on the realization of the common shock.

Lemma 1 shows that this approximation can be used to determine the behavior in the neighborhood of zero. Roughly speaking, the true system has positive probability of converging to zero if and only if the system in (2) has positive probability of converging to zero.

IV. CONFORMITY VERSUS DIVERSITY WITH TWO EQUALLY EFFICIENT CHOICES

This section begins our analysis of word-of-mouth learning by examining the special case in which $p = 1/2$, so that the payoffs to the two choices have the same distribution, and hence are equally efficient.⁵ Hence there is nothing for the players to learn, although the players do not realize this. By abstracting from the forces that favor the adoption of a superior product, this simple model allows

⁵ Note that the mean payoffs to the two choices still fluctuate—the two choices have the same ex ante mean but one or the other can do better ex post in a given period

us to highlight the effect of learning rules in determining when a market is likely to be dominated by a single product, practice, or technology. Also, the simplicity of this case makes it a useful introduction to the more general analysis of the next section.

Theorem 1 below shows that the system exhibits conformity when players gather very little information, and exhibits diversity when the word-of-mouth communication is more extensive. To understand the mechanism behind this result, it helps to note first that when each player talks to exactly one other player, there will be equally many players using f who hear about g and using g who hear about f . Given the symmetry of the payoffs, the popularity of g will thus move up or down in equal increments depending on the realization of the common shock. For example, when g is unpopular, the dynamics might dictate that g will either decrease in popularity by 50 percent or increase in popularity by 50 percent. Because $(\frac{1}{2})(\frac{3}{2}) = \frac{3}{4} < 1$, the system has negative "log drift" in the sense explained in the previous section, and so such dynamics imply that unpopular actions tend to die out.

In contrast, when players obtain more word-of-mouth information, unpopular technologies tend to make comebacks. Suppose that g is very unpopular and the sample size N is large. While nearly all g -users who consider switching in a given period will see an f -user, the size of this group will still be small. Because the number of f -users is much larger, there will be many more f -users who hear about g than vice versa. Hence, favorable draws of the common shock obtained by g -users will cause a large number of players to switch to g , while unfavorable draws will cause only a few players to switch away from g . As a result, unpopular technologies will not die out, and the system will exhibit diversity.

THEOREM 1. When $p = \frac{1}{2}$,

(a) The state cannot converge to an endpoint if

$$1 > \frac{-\log(1 + \alpha((N+1)q^-(N) - 1))}{\log(1 + \alpha((N+1)q^+(N) - 1))} \equiv Q(\alpha, N).$$

(b) The state converges to an endpoint with probability 1 if $1 < Q(\alpha, N)$. In this case the system has positive probability of converging to either endpoint.

(c) There is a unique $N^*(\alpha)$ such that $Q(\alpha, N^*(\alpha)) = 1$. Further, N^* is increasing in α , and $N^*(\alpha) > 1 \forall \alpha$ provided that $\theta > 0$.

(d) For $N < N^*(\alpha)$ the system exhibits herding; i.e., the state converges to an endpoint with probability one. For $N > N^*(\alpha)$ the state does not converge, and so the system exhibits diversity.

Proof. For part (a), rewrite equation (2) as

$$(2') \quad x_{t+1} = \begin{cases} x_t(1 + \alpha[(N+1)q^+(N) - 1]) & \text{with probability } p \\ x_t(1 + \alpha[(N+1)q^-(N) - 1]) & \text{with probability } 1 - p. \end{cases}$$

If $(N+1)q^-(N) > 1$, then in the neighborhood of 0 the state increases at each date regardless of the realization of θ . Clearly, the state cannot converge to zero in this case. If $(N+1)q^-(N) < 1$ and $Q(\alpha, N) < 1$, applying part (a) of Lemma 1 yields the same conclusion.⁶ Symmetry then implies that the state cannot converge to one.

For part (b), suppose that $1 < Q(\alpha, N)$ and let Λ be the (tail) event that x_t does not converge to either 0 or 1, and suppose that Λ has positive probability. By Paul Levy's zero-or-one law (see, e.g., Chung [1974], p. 341), $\lim_{t \rightarrow \infty} \text{prob}(\Lambda | x_t) = 1_\Lambda$ a.e., where 1_Λ is the indicator function on Λ . Lemma 1(b) shows that there is an $\epsilon > 0$ and an $x' > 0$ such that $\text{Prob}(\Lambda | x_t < x' \text{ or } x_t > 1 - x') < 1 - \epsilon$. Hence if $\{x_t\}$ does not converge to an endpoint, then with probability one there must be a T such that $x_t \in (x', 1 - x')$ for all $t > T$. But for all x_t in this interval, the state can move to some $x_{t'} < x'$ in a finite number of steps, and the probability that this occurs is bounded away from zero. Thus, by theorem 9.5.2 of Chung, there is probability zero that the state remains in $(x', 1 - x')$ infinitely often without making infinitely many transitions to the region $[0, x']$, and so the set Λ must be empty.

For part (c), define $Q^*(\alpha, N) = (1 - \alpha + \alpha(N+1)q^+(N))(1 - \alpha + \alpha(N+1)q^-(N)) = (1 - \alpha)^2 + \alpha(1 - \alpha)(N+1) + \alpha^2(N+1)^2q^+(N)q^-(N)$, where the last equality follows from $q^+(N) + q^-(N) = 1$. Since Q^* is obtained by multiplying through $Q = 1$ by its denominator and then exponentiating, $Q^*(\alpha, N) = 1$ if and only if $Q(\alpha, N) = 1$. We first check that for any α there is at least one N that solves $Q^*(\alpha, N) = 1$. This follows from the computations $Q^*(\alpha, 1) = 1 - (2\alpha q^+(1) - \alpha)^2 \leq 1$ and $\lim_{N \rightarrow \infty} Q^*(\alpha, N) = \infty$, and the observation that Q^* is a continuous function.

Appendix 2 shows that $Q^* = 1$ implies $\partial Q^*/\partial N > 0$, and hence there is a unique $N^*(\alpha)$ such that $Q^*(\alpha, N^*(\alpha)) = 1$. Appendix 2 also

⁶ Lemma 1 does not apply if $(N+1)q^-(N) > 1$, as the lemma supposes that $\gamma_2 < 1$

shows that $\partial Q^*(\alpha, N^*(\alpha))/\partial \alpha < 0$ so that N^* is an increasing function. Finally, since $Q^*(\alpha, 1) < 1$, we have $N^*(\alpha) > 1$ provided that $\theta > 0$.

QED

The shape of the curve $N^*(\alpha)$ depends on the relative size of the aggregate and idiosyncratic shocks, which we measure by $\eta = \theta/\sigma$. When N is below the curve $N^*(\alpha)$, the system converges to a state of conformity, where all consumers use the same product. Intuitively, the fewer other consumers each consumer samples, the less likely the consumer is to observe and then adopt a product with a small market share, so that when N is small a product with a very small share will tend to be driven from the market. If N lies above this curve, the state does not converge (provided that $\theta > 0$), so that the market shares are always changing, and do not approach a state where one product captures the entire market. Figure I displays the graph of $N^*(\alpha)$ for the case $\eta = 2$.

A particularly strong form of diversity obtains when N is very large. Let \bar{N} be the solution to $\bar{N} = 1/q^-(\bar{N}) - 1$. (The solution

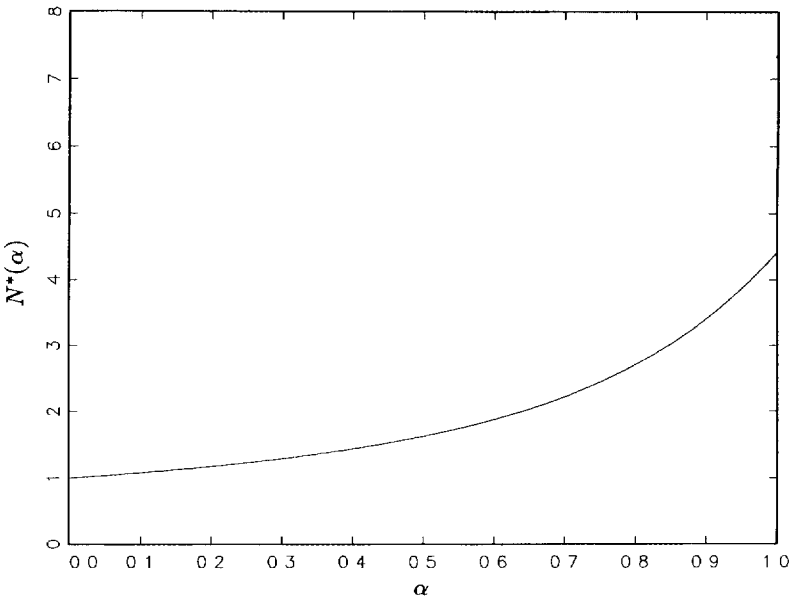


FIGURE I

exists and is unique.⁷) If $N > \bar{N}$ (which is about 40 when $\eta = 2$ and about 5 when $\eta = 1$), then not only does the system not converge to conformity, it never even enters a neighborhood of a state where everyone uses the same product. This is because, when N is this large, the state moves toward the center when it is near an endpoint, even if the common shock favors the more prevalent choice. This occurs because the support of the idiosyncratic noise is large enough that, even when the common shock favors product g , some users of product f will receive payoffs that exceed that period's mean payoff to g , and those g -users whose only observation of f is favorable will switch, so each f -user who receives a favorable draw will cause N g -users to switch.

As shown in Theorem 1, $N^*(\alpha)$ is increasing in α . The intuition for this is that the system is more likely to lock on when it takes large jumps; and so as α increases larger N 's become consistent with the system locking on to a state of conformity.

In this simple case of equally efficient choices, we can also determine the relationship between the relative sizes of the aggregate and idiosyncratic shocks, as measured by η , and the tendency for the system to exhibit conformity or diversity. When there are only common shocks, ($\eta = \infty$), we have $q^+(N) \equiv 1$ and $q^-(N) \equiv 0$, so that $N^*(\alpha) = 1/(1 - \alpha)$. When η is very small, so the idiosyncratic shocks dominate, q^+ and q^- are both approximately $1/2$, and so $N^*(\alpha)$ is slightly greater than 1 for all α . These two curves are graphed in Figure II, along with the previously displayed curve for $\eta = 2$. Note that the long-run behavior of the system is independent of η when α is small: the system then exhibits conformity for $N = 1$, and diversity for $N \geq 2$.

The conclusion that, with substantial inertia, conformity obtains only for the extreme case of $N = 1$ might seem to indicate that conformity is unlikely, at least absent an efficiency explanation. Yet even this degree of conformity may be surprising, given that players have no memory and always gravitate toward the choice which did better in the previous period. Our interpretation of this result is that the combination of limited word-of-mouth communication and our implicit "must-see-to-use" assumption

⁷ It is sufficient to show that $(N + 1)q^-(N)$ is increasing whenever $(N + 1)q^-(N) = 1$. The calculations in Appendix 2 show that $(d/dN)(N + 1)q^-(N) = q^-(N) - u\phi(u)/2N$. Substituting $-1 + 1/q^-(N)$ for N , and using $u\phi(u) < \sqrt{2\pi}e$, it is straightforward to check that this expression is positive

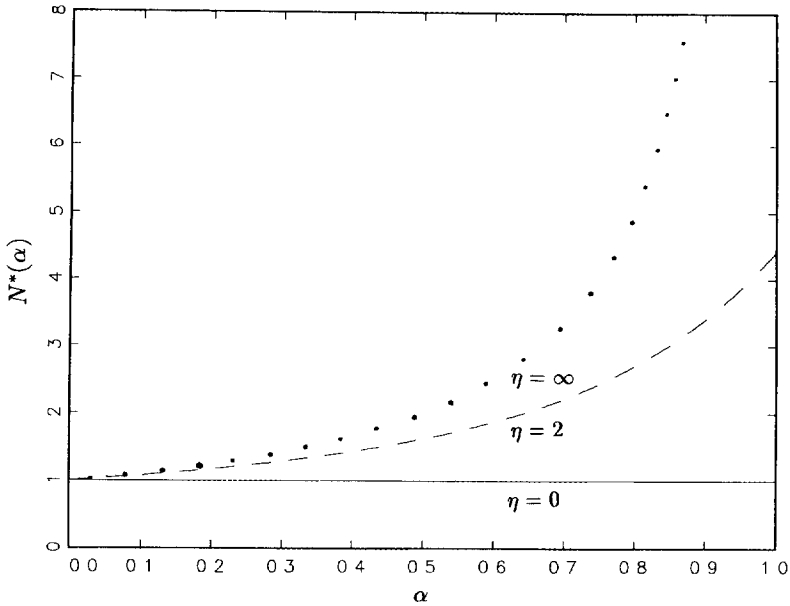


FIGURE II

generate a tendency toward conformity that is not present when players sample more broadly.

Appendix 2 shows that $\partial Q^*(\alpha, N^*(\alpha)) / \partial \eta < 0$, so as the importance of the common shocks increases, the system converges to uniformity for a larger set of parameter values. This should be fairly intuitive: because both choices are equally good, in the absence of a common shock there is no force that favors one choice over the other. Indeed, in the extreme case $\eta = 0$ where common shocks are absent, the system is deterministic, and converges to $x = 1/2$ for any $N \geq 2$. (For $N = 1$ the system becomes $x_{t+1} = x_t$.)

V. PAYOFF-RELEVANT CHOICES AND EFFICIENT SOCIAL LEARNING

We now turn our attention to the diffusion of unequal technologies. In contrast to the situations studied in the previous section, the long-run outcome here may have substantial efficiency consequences. Technological change has generally been regarded as accounting for a significant fraction of growth, and we believe

that models of learning may help explain why the diffusion of innovations has so often been found to be slow and incomplete.⁸

We now refine conformity into two subcategories: we will say that a population exhibits inefficient herding if in the long run everyone uses the inferior technology, and that it exhibits “efficient social learning” if in the long run everyone uses the superior choice.

Formally, we examine here the general model of Section II for the case of $p > \frac{1}{2}$, so that technology g is superior in the sense that it is better more than half of the time. For example, this superior technology could take the form of a new agricultural practice that outperforms the old one for a wide range of weather conditions, or a new management practice that is superior for a large set of economic conditions.

The behavior of our model is most easily understood by thinking of the diffusion process as being affected by two forces. First, the fact that g is better more often tends to make g grow in popularity, although given the players’ inability to aggregate information, this force is of limited strength. Second, we have the tendency of word-of-mouth communication to produce conformity or diversity depending on the number of “friends” that the agents contact. As we will see, these forces can combine to produce efficient social learning despite the naive behavior we have assumed at the individual level. This occurs when the tendency of limited word-of-mouth communication to produce conformity is sufficiently strong so that in combination with the force of superior payoffs it allows g to dominate, while at the same time the tendency for conformity is too weak to allow inefficient herding on f .

Theorem 2 provides a complete characterization of the behavior of the model. In general, all three behaviors are possible. As we will see later, however, the model typically exhibits diversity when the word-of-mouth communication is extensive, and efficient social learning when players receive less information.

THEOREM 2. Suppose that $p > \frac{1}{2}$.

- (i) If $[p/(1-p)] < Q(\alpha, N)$, the system exhibits inefficient herding: the state converges to an endpoint with probability one, and converges to everyone using the inferior product f with positive probability from all initial conditions.

8 See, e.g., Griliches [1957, 1994] and Mansfield [1968] for empirical studies

- (ii) If $[(1 - p)/p] < Q(\alpha, N) < [p/(1 - p)]$, the system exhibits efficient social learning; i.e., it converges with probability one to the state where everyone uses g .
- (iii) If $[(1 - p)/p] > Q(\alpha, N)$, the system exhibits diversity, as the state does not converge to either endpoint. Moreover, if $\eta \neq 0$, the system does not converge to any deterministic limit; more precisely, $\text{Prob}(\exists x \text{ s.t. } x_t \rightarrow x) = 0$.

Proof of Theorem 2.

- (i) This follows Lemma 1 in the same way as the proof of Theorem 1(a). (To analyze convergence to $x = 1$, make the change of variables $y = 1 - x$.)
- (ii) From Lemma 1(c), there is an $\underline{x} > 0$ such that $\text{Prob}[x_t < \underline{x} \forall t | x_0] = 0$ for all x_0 . From Lemma 1(b) there is an $\epsilon > 0$ and a $x^\epsilon < 1$ such that $\text{Prob}[x_t \rightarrow 1 | x_0 \geq x^\epsilon] \geq \epsilon$. Moreover, for any $x \in [\underline{x}, x^\epsilon]$, there is probability bounded away from zero that x_t increases in every period until it reaches a state greater than or equal to x^ϵ , and the number of steps in this transition is finite. Hence whenever the system is below x^ϵ , it has positive probability of passing above it; and each time it does so, it has probability at least ϵ of converging to 1. Hence from Theorem 9.5.4 of Chung [1974], the system converges to 1 almost surely.
- (iii) From Lemma 1(c) the system cannot converge to any point below some $\underline{x} > 0$ or to any point above some $\bar{x} < 1$, and the system has steps that are bounded away from zero in the interval $[\underline{x}, \bar{x}]$.

QED

Theorem 2 shows that, as in the previous section, the long-run behavior of the system for a fixed level of η is summarized by size of $Q(\alpha, N)$. Figure III displays numerical plots of the curves \bar{N} and \underline{N} defined by

$$Q(\alpha, \bar{N}(\alpha)) = (1 - p)/p$$

$$Q(\alpha, \underline{N}(\alpha)) = p/(1 - p)$$

for the parameter values $p = 0.6$ and $\eta = 2$. We have not been able to prove that \underline{N} and \bar{N} must be curves in general, but this seems to be the case in our numerical computations.⁹

⁹ The results of the last section and a continuity argument show that \underline{N} must be a curve for p 's sufficiently near to $1/2$.

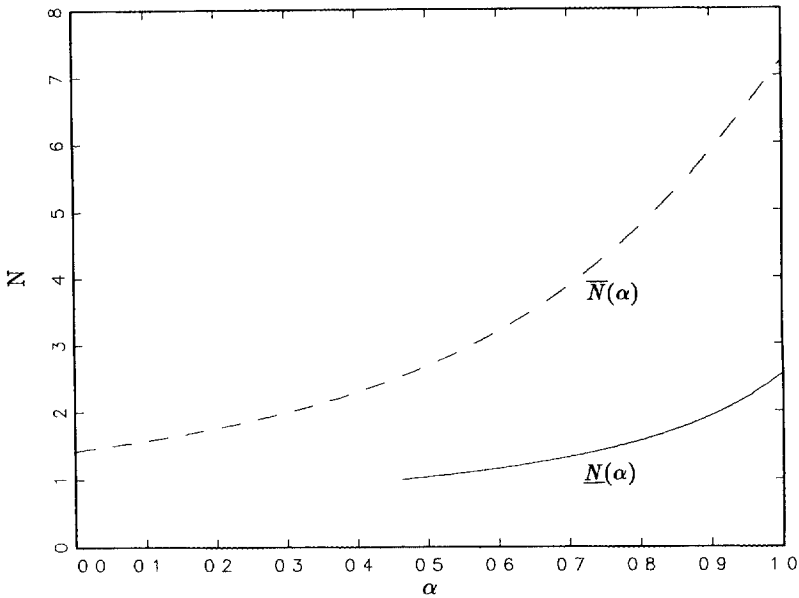


FIGURE III

Looking at Figure III, in the top region where $N > \bar{N}(\alpha)$ the system exhibits diversity, as before. In the bottom region, where $N < \tilde{N}(\alpha)$, the system exhibits conformity or “two-sided herding,” with the state eventually converging to one of the endpoints, and both endpoints having positive probability of being selected. This behavior is much like the conformity case of the previous section; the one new feature is that the state can converge to the inferior choice. However, unlike in the previous section, this two-sided herding cannot occur for very small values of α as now $N \leq \tilde{N}(\alpha)$ may require that $N < 1$: when f , say, is an inferior choice, there is a smaller set of parameter values where the state can converge to f .

The most novel aspect of Figure III is the intermediate region where $\tilde{N}(\alpha) < N < \bar{N}(\alpha)$. Here the system exhibits a form of efficient social learning, in the sense that all of the agents eventually adopt and stick with the superior product. Intuitively, because players using a word-of-mouth decision rule are unlikely to switch to an unpopular product (at least if N is not too large), the state has a tendency to be absorbed by an endpoint whenever a neighborhood of that endpoint is reached. A payoff difference

between the products provides an additional force favoring the superior choice, which increases the probability of converging to the superior choice, and can reduce or eliminate the probability of converging to the opposite, inferior, extreme

Figure IV illustrates the way the three regions change as p increases from 0.55 to 0.70, again with $\eta = 2$. As seen in the figure, increasing p increases $p/(1-p)$, which shifts the curve N down, and decreases $(1-p)/p$, which shifts the curve \bar{N} up. Thus, the "efficient learning region" $\underline{N}(\alpha) < N < \bar{N}(\alpha)$ where the state converges to the better choice increases, and the two other regions both shrink.

Three properties of the efficient learning region are worth noting. First, the curve $N^*(\alpha)$ of the previous section (which corresponds to $p = 1/2$) lies within the efficient learning region for any $p > 1/2$, and indeed $N = N^*(\alpha)$ is the unique value of N which leads to the efficient choice for all values of p . We will therefore refer to sampling $N^*(\alpha)$ people as the uniformly optimal rule.

Second, $\underline{N}(\alpha)$ is less than 1 when α is not too large, where "too large" is 0.2 for $p = 0.55$ and increases in p . Since $\bar{N}(\alpha) > N^*(\alpha) > 1$, $N = 1$ leads to efficient learning for a wide range of α 's provided that the two choices are not too similar.¹⁰

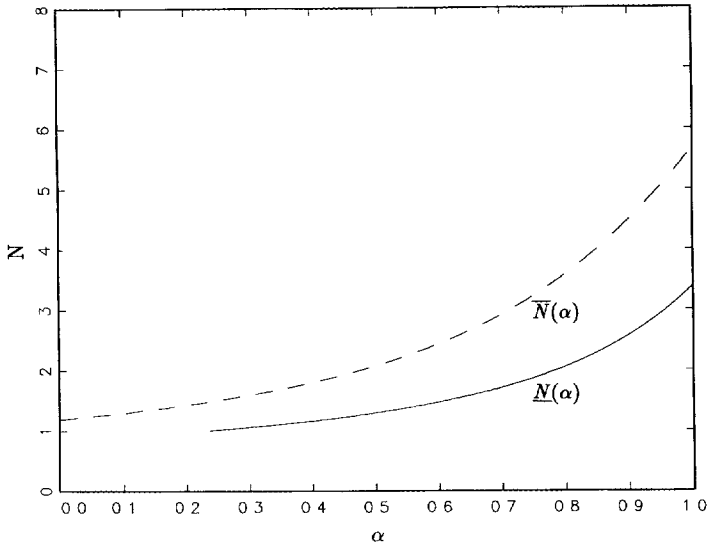
Third, $Q(\alpha, N) < (1-p)/p$ for all $N > \bar{\bar{N}}$, where $\bar{\bar{N}}$ is the number defined in the last section such that $N > \bar{\bar{N}}$ implies the state moves toward the center with probability one when it is sufficiently near the endpoints. Hence, the curve provides an upper bound on the efficient learning region which is independent of p . When N exceeds this level, the forces that prevent convergence are strong enough to prevent efficient learning even if one technology is always better than the other; i.e., if its mean payoff is better with probability one.¹¹

The long-run dynamics of this model are fairly similar to those in the model with explicit popularity weighting we considered in our earlier paper [Ellison and Fudenberg 1993], where players were assumed to be less willing to switch to an unpopular choice. In that model, as in this one, there are three regions of parameter space. With very little popularity weighting, the system oscillates,

10 More formally, $Q(\alpha, 1)$ is continuous in α , and $\lim_{\alpha \rightarrow 0} Q(\alpha, 1) = 1$, so for sufficiently small α we have $Q(\alpha, 1) \in ((1-p)/p, p/(1-p))$, and $N = 1$ yields efficient social learning. Here "sufficiently small" is about 0.5(!) for $p = 0.65$, and about 0.2 for $p = 0.55$, as p approaches 0.5 the required α 's shrink to 0.

11 But note that $\bar{\bar{N}}$ depends on the magnitude θ of the payoff difference as θ increases, the probability q^- decreases.

$p = 0.55$



$p = 0.70$

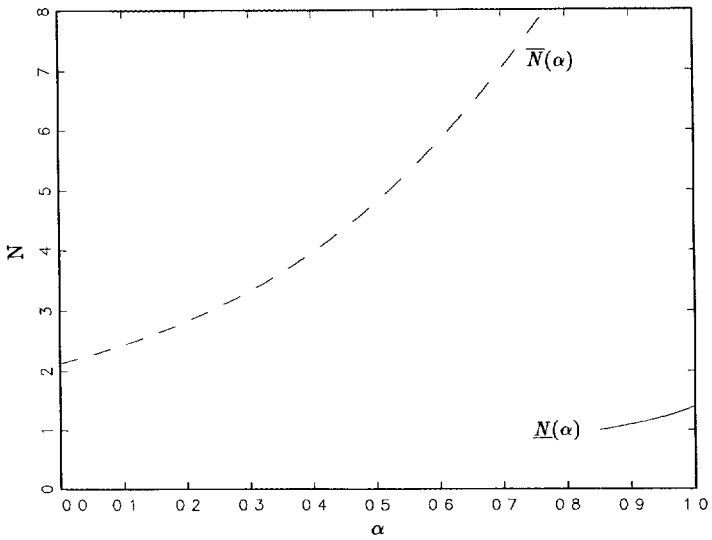


FIGURE IV

just as it does here when N is large. (Indeed, in the limit as N and η go to infinity, the word-of-mouth model converges to the zero-popularity-weighting case of the previous paper.) With very substantial popularity weighting, the system always converges, but can converge to the wrong choice, which corresponds to small N here.¹² Finally, for an intermediate amount of popularity weighting, the state always converges to the better choice. Unlike with word-of-mouth decisions, though, there is a single value of the popularity weighting that ensures the long-run outcome is efficient.

VI. MORE COMPLEX DECISION RULES

To test the robustness of our qualitative conclusions, we now consider two more complex variants of the word-of-mouth decision rule.

1. Popularity Weighting

The first variant we consider supposes that players use popularity weighting of the kind we considered in Ellison and Fudenberg [1993]: players who observe the average payoffs \bar{u}_i^g and \bar{u}_i^f chose g if $\bar{u}_i^g - \bar{u}_i^f > m(1 - 2x_i)$. Note that this rule supposes that agents observe the aggregate popularity x_i in the entire population, but learn the payoffs of the agents they talk to directly. This seems plausible to us, as popularity is easier to communicate than a level of satisfaction, and also may be well proxied by the consumer's observations of, e.g., the relative numbers of different brands of cars on the road to work.

Intuitively, popularity weighting increases the forces favoring convergence to an endpoint, and so should increase the set of parameters for which convergence occurs. Thus, beginning from a situation of no popularity weighting, i.e., $m = 0$, a small increase in m , holding N constant, should increase the efficiency of the system if (α, N) lies just above the curve \bar{N} that determines the upper boundary of the optimal learning region, and an increase in m with N constant should decrease the long-run efficiency of the system starting at a point (α, N) that lies just above the curve \underline{N} that determines the lower boundary of the optimal learning region. Moreover, a nonnegligible amount of popularity weighting could shift the curve \underline{N} up enough so that it lies above the point $(0, 1)$, reversing our conclusion that $N = 1$ is optimal when the inertia is

¹² But, unlike the previous paper, this case may not arise if α is small

high. Consequently, when players use popularity weights, larger sample sizes may be optimal.

All of these intuitions turn out to be correct. To state the conclusions formally, let

$$q^+(N, m) = \Phi\left(\frac{\theta - m}{\sigma\sqrt{1 + 1/N}}\right) \quad \text{and} \quad q^-(N, M) = \Phi\left(\frac{-\theta - m}{\sigma\sqrt{1 + 1/N}}\right)$$

be the analogs of q^+ and q^- when players use popularity weighting, let

$$Q(\alpha, N, m) = \frac{-\log(1 + \alpha((N + 1)q^-(N, m) - 1))}{\log(1 + \alpha((N + 1)q^+(N, m) - 1))},$$

and let $N^*(\alpha, m)$ be the solution to $Q(\alpha, N, m) = 1$. (Calculations similar to those of Section IV show that a solution exists and is unique.) As in Section V, $N = N^*(\alpha, m)$ is the one sampling rule that guarantees efficient social learning uniformly over all values of p .

THEOREM 3. Suppose that $p > 1/2$.

- (i) If $[p/(1 - p)] < Q(\alpha, N, m)$, the state converges to an endpoint with probability one, and exhibits inefficient herding, as the state converges to the inferior choice with positive probability from initial conditions sufficiently close to 0. If m is large enough, this inefficient herding can occur with probability one for some initial conditions.
- (ii) If $[(1 - p)/p] < Q(\alpha, N, m) < [p/(1 - p)]$, there is efficient social learning: the system converges with probability one to the state where everyone uses g .
- (iii) If $[(1 - p)/p] > Q(\alpha, N, m)$, the system does not converge to either endpoint. Moreover, if $\eta \neq 0$, the system does not converge to any deterministic limit; more precisely, $\text{Prob}(\exists x \text{ s.t. } x_t \rightarrow x) = 0$.
- (iv) $Q(\alpha, N, m)$ is increasing in m wherever it is defined and nonnegative, so that a small increase in m , holding other parameters constant, either causes a change from case (iii) to case (ii) or a change from case (ii) to case (i), or has no effect on which of the three cases prevails. Moreover, the uniformly optimal sampling rule $N^*(\alpha, m)$ is increasing in m .

Since the analysis closely follows that of the previous sections, we omit the proof. Figure V graphs \bar{N} and \underline{N} from $m = 0$ and $m = \theta$, with $\sigma = 0.5\theta$ and $p = 0.6$. (Thus, the case $m = 0$ repeats Figure

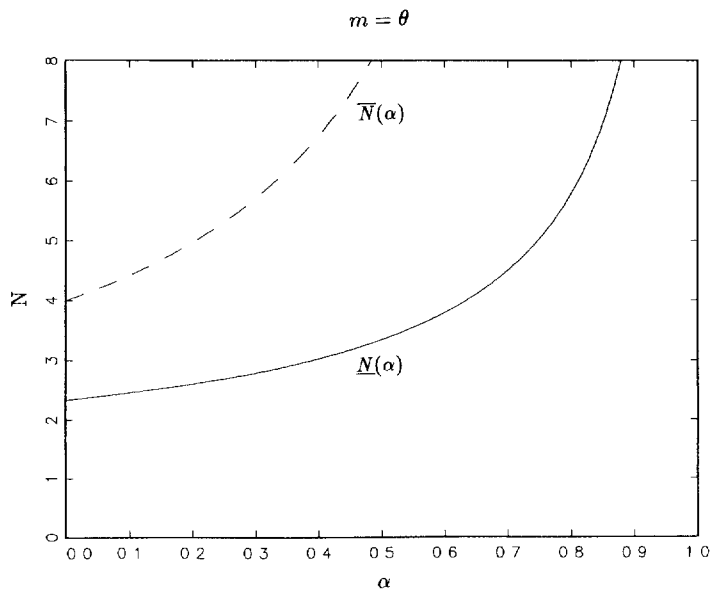
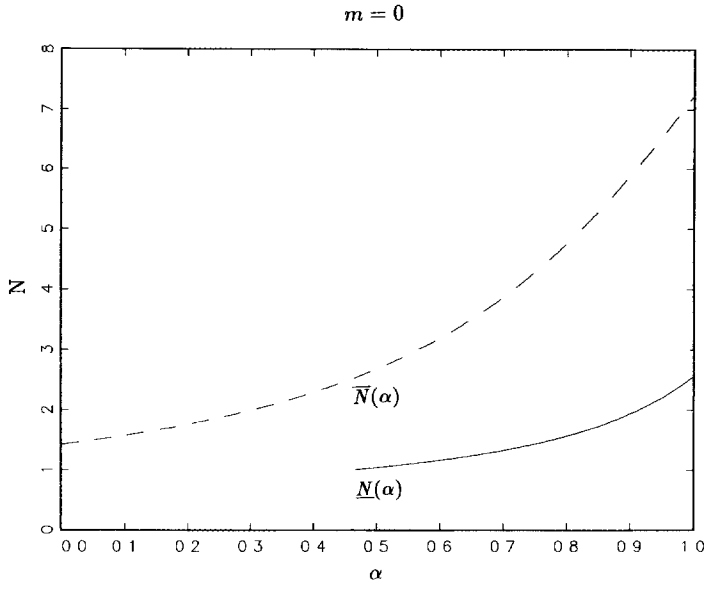


FIGURE V

III.) Note that with popularity weighting in this example, the optimal sample size N is about three. The one qualitative change caused by popularity weights is that for small N and large enough m the system is certain to lock on to the nearest endpoint whenever it starts out sufficiently close to it.

2. Infrequent Sampling and Social Memory

In the models we have considered so far, word-of-mouth sampling generated information on the performance of the various choices in the previous period. These models thus implicitly supposed that the N players each agent sampled have all observed the payoffs in the previous period, presumably because they themselves had used the products then. As a further check on the robustness of our conclusions, we now consider a model in which individuals observe the payoffs less frequently, so that the individuals in an agent's sample may report observations from a number of different past periods, and the sampling process incorporates a form of "social memory" that can reduce the effect of the common shocks. This social memory might be expected to give a clearer advantage to the product that is better on average, as it reduces the probability that unrepresentative draws push the system in the wrong direction. Social memory might also be expected to make larger sample sizes more efficient than in the model without social memory, where large sample sizes preclude fully efficient social learning. Without social memory a large sample can still be misleading due to an unrepresentative value of the common shock, but a large sample with a large social memory should lead almost all of the agents to choose the better action. As we will see, both of these intuitions are essentially correct.

Allowing for social memory requires players to compare payoffs to f in one period with those to g in another. For this reason, we find it easiest to specialize the distribution of payoffs so that the common shock affects the payoffs to f and g symmetrically. That is, we suppose that $\bar{g}_t = \theta_t/2$, $\bar{f}_t = -\theta_t/2$.

We begin with the case of a two-period social memory. For simplicity, we assume that a fraction α of the players consider switching every other period. This leads to a Markov process with three realizations for the common shocks, 0, 1, and 2, corresponding to the number of times the common shock took on the value θ in the previous two periods. In a neighborhood of $x = 0$, the system

then evolves according to

$$x_{t+1} = \begin{cases} x_t(1 - \alpha + \alpha(N + 1)q^+(N)) + o(x_t) & \text{with probability } p^2 \\ x_t(1 - \alpha + \alpha(N + 1)/2) + o(x_t) & \text{with probability } 2p(1 - p) \\ x_t(1 - \alpha + \alpha(N + 1)q^-(N)) + o(x_t) & \text{with probability } (1 - p)^2. \end{cases}$$

In this system the transitions for realization 0 and 2 are the same as those for the realizations $-\theta$ and θ , respectively, in the original model. In realization 1, players who see both actions played are equally likely to think that either is better.

This system takes on a particularly simple form if $N = 1$, for then in realization 1 the system does not move (To see why, note that there are just as many f -players who observe g -players as g -players who observe f -players, and that in realization 1, the probability that one draw of g beats one draw of f is $1/2$.) The same arguments as in the proof of Proposition 2 show that there is efficient social learning with $p > 1/2$ if

$$p^2/(1 - p)^2 > Q(\alpha, 1) > (1 - p)^2/p^2,$$

which is the previous condition with p^2 replacing p , and $(1 - p)^2$ replacing $(1 - p)$. Since this is a weaker condition, $N = 1$ will still yield efficient social learning provided that α is not too large.

With larger sample sizes, the behavior of the model is more complicated, because it is no longer true that the system remains constant ($x_{t+1} = x_t$) following the realization of one draw of each value of the shock. For instance, with very large samples, almost all players will have both f -users and g -users in their samples, and since in realization 1 the payoffs of f and g are drawn from the same distribution, half of these players will adopt each technology.¹³ This provides an additional force favoring diversity, instead of the null transition that occurs in state 1 when $N = 1$. For this reason, it is no longer necessarily true that efficient learning occurs for any parameter values that yield efficient learning with one-period social memory, and indeed we have found examples where increasing social memory from 1 to 2, holding other parameters constant, causes long-run behavior to switch from efficient learning to diversity. However, increasing the social memory cannot allow inefficient herding to occur when it would not occur previously.¹⁴

13 Remember that the distribution of the idiosyncratic shocks is normal and so is symmetric

14 For the special case of two-period social memory, we can show that the long-run outcome is as if the regions shown in Figure III are shifted down, so that

Now we turn to the limit case of an infinite social memory, which we model by supposing that the distribution of experience in the population is deterministic, with a fraction p of the g -users having experience based θ^+ and fraction $(1 - p)$ having experience based on θ^- . Hence the aggregate system is deterministic and its evolution is given by

$$(3) \quad x_{t+1} = (1 - \alpha)x_t + \alpha \text{Prob}(\hat{g} > \hat{f} | N, x_t),$$

where the probability term in (3) is the probability that the average realized payoff to g exceeds the average realized payoff to f in a sample of $N + 1$ players drawn from a population where x_t of the players use g . (If no players in the sample use g , g 's average realized payoff in the sample is minus infinity.)

In a neighborhood of $x = 0$, equation (3) can be approximated by

$$x_{t+1} = x_t[1 - \alpha + \alpha(N + 1)(p\bar{q}^+(N) + (1 - p)\bar{q}^-(N))] + o(x_t),$$

where $\bar{q}^+(N)$ is the probability that a single draw from the favorable distribution of g exceeds the average of N draws from the distribution of f .

Thus, the state converges to 0 if $p\bar{q}^+(N) + (1 - p)\bar{q}^-(N) < 1/(1 + N)$, and the state cannot converge to zero if the reverse inequality holds. If $p < 1/2$, this inequality is satisfied for $N = 1$. To see this, note that

$$\begin{aligned} p\bar{q}^+(1) + (1 - p)\bar{q}^-(1) &= p^2\Phi\left(\frac{\theta}{\sigma\sqrt{2}}\right) + 2p(1 - p)\frac{1}{2} \\ &\quad + (1 - p)^2\left(1 - \Phi\left(\frac{\theta}{\sigma\sqrt{2}}\right)\right). \end{aligned}$$

This expression will be less than $1/2$ if the (weighted) average of the first and third term is less than $1/2$. Because $p^2 < (1 - p)^2$, $\Phi(\theta/\sigma\sqrt{2}) > 1/2$, and the fact that the unweighted average of these two terms is $1/2$, this is true. Hence under infinite social memory $N = 1$ yields efficiency for any value of the inertia parameter α , while $N = 1$ and large α can lead to inefficient herding without social memory. Intuitively, the social memory averages out the unrepresentative draws that can otherwise lead to inefficient herding when α is large.

some parameters that previously corresponded to inefficient herding lead to efficient learning, and some parameters that corresponded to efficient learning lead to diversity. The calculations are available on request.

There is still diversity as N grows, since \bar{q}^+ is at least $1/2$, and so the inequality above cannot be satisfied for any $N > -1 + 2/p$. We should note, though, that this observation may say more about the inclusiveness of our notion of diversity than about the long-run behavior of the system. When N is large and $p < 1/2$, the system does not converge to zero, but its steady state is very near to this point. The intuition for this is simple. If players receive large samples from the population distribution, and that distribution corresponds to the long-run distribution of θ , then most players should receive representative samples and hence make the payoff-maximizing choice. Thus, while large samples continue to preclude exact efficiency, the efficiency cost of large samples is low when the social memory is large

(More formally, suppose that $p < 1/2$, and that there is a sequence of steady states x_N^* that remains bounded away from zero as N goes to infinity. Then for any $\epsilon > 0$, there is an N' such that if the sample size N is greater than N' and the system is at the steady state x_N^* , at least $(1 - \epsilon)$ of the players will have samples that contain many draws of both f and g , and, moreover, have average payoffs in the sample for each technology that are within $(2p - 1)\theta/3$ of their theoretical average payoffs. Hence, at least $(1 - \epsilon)$ of the players will choose the better technology, and so x_N^* can be at most ϵ , which contradicts the hypothesis that the steady states were bounded away from 0.)

VII. CONCLUSION

This paper has studied a simple model of word-of-mouth communication. We find that, despite the naive play of individuals, this type of information flow may lead to efficient learning on the social level, and that social learning is often most efficient when communication between agents is fairly limited. When the choices are equally good on average, so that efficiency is not an issue, either herding or diversity can occur, with herding more likely with small sample sizes. These qualitative conclusions are largely independent of the composition of the noise, and are fairly robust to some obvious changes in the specification of the model although popularity weighting can make larger sample sizes optimal.

In addition to our specific conclusions, we hope that our analysis demonstrates that boundedly rational learning models are an interesting and tractable way to understand some aspects of social learning.

APPENDIX 1

LEMMA 1. Let x_t be a Markov process on $(0,1)$ with

$$x_{t+1} = \begin{cases} H_1(x_t) & \text{with probability } p \\ H_2(x_t) & \text{with probability } (1 - p) \end{cases}.$$

Suppose that $H_i(x_t) = \gamma_i x_t + o(x_t)$, with $\gamma_2 < 1 < \gamma_1$.

(a) If

$$\frac{p}{1-p} > -\frac{\log(\gamma_2)}{\log(\gamma_1)},$$

then x_t cannot converge to 0 with positive probability.

(b) If

$$\frac{p}{1-p} < -\frac{\log(\gamma_2)}{\log(\gamma_1)},$$

then there are strictly positive δ and ϵ such that $\text{Prob}[x_t \rightarrow 0 | x_0 \leq \delta] \geq \epsilon$.

(c) If

$$\frac{p}{1-p} > \frac{-\log(\gamma_2)}{\log(\gamma_1)},$$

there is a $x^* > 0$ such that for all $x_0 > 0$, $\text{Prob}[x_t < x^* \forall t | x_0] = 0$.

Proof of Lemma 1. (a) for any ϵ such that $0 < \epsilon < \gamma_2$, define $H_i^\epsilon(x) = \gamma_i x - \epsilon x$, and choose $\delta > 0$ such that, for $i = 1, 2$, $H_i^\epsilon(x) < H_i(x)$ whenever $x < \delta$. If for some initial condition x_0 there is positive probability that $x_t \rightarrow 0$, then from the same initial condition there is a T such that there is positive probability that of the event $\{x_t \rightarrow 0 \text{ and } x_t < \delta \text{ for all } t \geq T\}$. For each realization of the process $\{x_t\}$, let $i(t) = 1$ if $x_t = H_1(x_{t-1})$, and let $i(t) = 2$ if $x_t = H_2(x_{t-1})$.

Define a new stochastic process $\{y_t\}$ on $(0, \infty)$ by

$$\begin{aligned} y_t &= x_t & \text{for all } t \leq T \\ y_t &= H_{i(t)}^\epsilon(y_{t-1}) & \text{for all } t > T. \end{aligned}$$

We claim that since the event $\{x_t \rightarrow 0 \text{ and } x_t < \delta \text{ for all } t \geq T\}$ has positive probability, then so does the event $y_t \rightarrow 0$. To see this, note that on this event, $y_T < \delta$ and $y_T \leq x_T$. Suppose that for all t

from T to $T + s, y_{T+s} < \delta$ and $y_{T+s} \leq x_{T+s}$. Then

$$y_{T+s+1} = H_{i(T+s+1)}^\epsilon(y_{T+s}) \leq H_{i(T+s+1)}^\epsilon(x_{T+s}) \leq H_{i(T+s+1)}(x_{T+s}) = x_{T+s+1} < \delta,$$

and so by induction the first event implies the second one.

Next, if $y_t \rightarrow 0$, then the process z_t defined by $z_t = \log y_t$ must converge to $-\infty$. Moreover, we have $z_{T+s} = Y_T + \sum_{\tau=1}^s u_\tau$ where the u_τ are i.i.d. binomial random variables, taking on the values $\log(\gamma_1 - \epsilon)$ and $\log(\gamma_2 - \epsilon)$ with probabilities p and $(1 - p)$, respectively

Thus, the strong law of large numbers implies that the conditional probability that z_t converges to $-\infty$ must be 0 if $E(u_\tau) > 0$; i.e., if $p \log(\gamma_1 - \epsilon) + (1 - p) \log(\gamma_2 - \epsilon) > 0$. If, as hypothesized in part (a) of the lemma, $[p/(1 - p)] > [\log(\gamma_2)/\log(\gamma_1)]$, then this inequality holds true for all sufficiently small ϵ , which proves part (a).

The argument for part (c) of the Lemma is nearly identical. Again, let ϵ be such that $p \log(\gamma_1 - \epsilon) + (1 - p) \log(\gamma_2 - \epsilon) > 0$ and choose δ so that $H_i^\epsilon(x) < H_i(x)$ whenever $x < \delta$. Let $x^* = \delta$, and define $\{y_t\}$ as above with $T = 0$. If $x_t < x^* \forall t$, then $y_t < x^* \forall t$ as well. Because $\{y_t\}$ converges to infinity with probability one, this event has probability zero.

(b) Suppose that

$$\frac{p}{1 - p} < - \frac{\log(\gamma_2)}{\log(\gamma_1)},$$

and choose ϵ such that

$$\frac{p}{1 - p} < - \frac{\log(\gamma_2 + \epsilon)}{\log(\gamma_1 + \epsilon)}.$$

Let $H_i^\epsilon(x) = \gamma_1 x + \epsilon x$, and choose $\delta > 0$ such that, for $i = 1, 2, H_i^\epsilon(x) > H_i(x)$ for $x < \delta$. Define a new stochastic process by $\{y_t\}$ on $(0, \infty)$ by

$$y_0 = x_0$$

$$y_t = H_{i(t)}^\epsilon(y_{t-1}) \text{ for all } t > 1$$

By the strong law, y_t converges to 0 with probability 1. Moreover, if $x_0 < \delta$, there is a positive probability of the event $\{y_t \rightarrow 0 \text{ and } y_t \leq \delta \forall t\}$. An induction like the one in part (a) shows

that this event implies that x_t converges to 0 as well, which completes the proof of part (b).

APPENDIX 2: SIGNING THE PARTIAL DERIVATIVES OF Q^*

(i) Proof that $(\partial Q^*/\partial N)(\alpha, N^*(\alpha)) > 0$. Recall (from the proof of Theorem 1) that $Q^*(\alpha, N) = (1 - \alpha)^2 + \alpha(1 - \alpha)(N + 1) + \alpha^2(N + 1)^2q^+(N)q^-(N)$. Taking the partial derivative yields

$$(A1) \quad \frac{\partial Q^*}{\partial N} = \alpha(1 - \alpha) + \alpha^2 2(N + 1)q^+(N)q^-(N) + \alpha^2(N + 1)^2q^+(N) \frac{\partial q^-(N)}{\partial N} + \alpha^2(N + 1)^2q^-(N) \frac{\partial q^+(N)}{\partial N}.$$

Using the change of variables

$$u = \frac{\theta}{\sigma\sqrt{1 + 1/N}},$$

we have

$$\frac{\partial}{\partial N} q^+(N) = \frac{\partial}{\partial N} \Phi(u) = \frac{u\phi(u)}{2N(N + 1)},$$

and

$$\frac{\partial}{\partial N} q^-(N) = -\frac{u\phi(u)}{2N(N + 1)},$$

where ϕ , the normal density, is the derivative of the normal c.d.f. Φ . Substituting into (A1) shows that it is sufficient to verify

$$(A2) \quad 1 - \alpha + 2\alpha(N + 1)\Phi(u)\Phi(-u) > \alpha \frac{N + 1}{2N} u\phi(u)(\Phi(u) - \Phi(-u))$$

when $N = N^*(\alpha)$.

Solving the quadratic equation that defines N^* , we find that

$$N^*(\alpha) + 1 = \frac{-\alpha(1 - \alpha) + \sqrt{\alpha^2(1 - \alpha)^2 + 4\alpha^3(2 - \alpha)\Phi(u)\Phi(-u)}}{2\alpha^2\Phi(u)\Phi(-u)}.$$

Substituting into (A2) then yields

$$\sqrt{(1 - \alpha)^2 + 4\alpha(2 - \alpha)\Phi(u)\Phi(-u)} > \alpha \frac{N^*(\alpha) + 1}{2N^*(\alpha)} \times u\phi(u)(\Phi(u) - \Phi(-u)).$$

Using $\alpha < 1$, $2 - \alpha > 1$, $(1 - \alpha)^2 > 0$, $(N + 1)/2N < 1$, $\Phi(u) > (\Phi(u) - \Phi(-u))^2$, and $\phi(u) = \phi(-u)$, it suffices that

$$(A3) \quad [u^2\phi(u)^2/\Phi(u)] < 4.$$

For $u > -2$, the left-hand side is at most $(u\phi(u))^2/\Phi(-2)$. Using $\phi' = -u\phi$, the numerator of this last expression is maximized at $u = 1$, so that the value of the left-hand side of (A3) is bounded by $(2\pi e\Phi(-2))^{-1}$, which is less than 4.

For $u \leq -2$ we use the fact that $\Phi(u) > \phi(u)|(u^2 - 1)/u^3|$ (see, e.g., Ross [1984], p. 161). Thus, $[u^2\phi(u)^2/\Phi(u)] < \frac{4}{3}u^3\phi(u)$. The maximum on the interval $u \leq -2$ occurs at -2 , where the left-hand side of (A3) is $32e^{-2}/3\sqrt{2\pi}$, which again is less than 4.

The other partial derivatives are easier to sign.

Writing

$$Q^*(\alpha, N) = 1 + \alpha[(N - 1) + \alpha((N + 1)q^+(N) - 1)((N + 1)q^-(N) - 1)],$$

we have

$$\begin{aligned} \frac{\partial Q^*}{\partial \alpha}(\alpha, N^*(\alpha)) &= \alpha[(N + 1)q^+(N) - 1)((N + 1)q^-(N) - 1)] \\ &\quad + (Q^*(\alpha, N^*(\alpha)) - 1)/\alpha \\ &= \alpha[(N + 1)q^+(N) - 1)((N + 1)q^-(N) - 1)] \\ &= 1 - N^*(\alpha) \leq 0, \end{aligned}$$

where the first equality comes directly from $Q^*(\alpha, N^*(\alpha)) = 1$, and the second uses that fact and the equation for Q^* . Note that the inequality is strict for all $\theta > 0$. Finally,

$$\frac{\partial Q^*}{\partial \eta}(\alpha, N^*(\alpha)) = \alpha^2(N + 1)^2 \frac{\partial}{\partial \eta}(q^+(N)q^-(N));$$

this is negative for $\theta > 0$ because $q^+(N) > \frac{1}{2}$, $q^+(N) = 1 - q^-(N)$, and $\partial q^+(N)/\partial \eta > 0$.

QED

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
HARVARD UNIVERSITY

REFERENCES

- Banerjee, Abhijit, "The Economics of Rumours," *Review of Economic Studies*, IX (1993), 309-27
Banerjee, Abhijit, and Drew Fudenberg, "A Bayesian Model of Word-of-Mouth Learning," 1994

- Chung, Kai Lai, *A Course in Probability Theory* (San Diego, CA Academic Press, 1974)
- Elhson, Glenn, and Drew Fudenberg, "Rules of Thumb for Social Learning," *Journal of Political Economy*, CI (1993), 612-43
- Griliches, Zvi, "Hybrid Corn: An Exploration of the Economics of Technological Change," *Econometrica*, XXV (1957), 501-22.
- , "Productivity, R&D, and the Data Constraint," *American Economic Review*, LXXXIV (1994), 1-23
- Kirman, Alan, "Ants, Rationality, and Recruitment," *Quarterly Journal of Economics*, CVIII (1993), 137-56
- Mansfield, Edwin, *The Economics of Technological Change* (New York: Norton, 1968)
- Ross, Sheldon, *A First Course in Probability* (New York, NY: Macmillan, 1984)
- Satterthwaite, Mark, "Consumer Information, Equilibrium Industry Price, and the Number of Sellers," *Bell Journal of Economics*, X (1979), 483-502
- Smallwood, Donald, and John Conlisk, "Product Quality in Markets Where Consumers Are Imperfectly Informed," *Quarterly Journal of Economics*, XCIII (1979), 1-23

