Discovery and Diffusion of Knowledge in an Endogenous Social Network[1]

Myong-Hun Chang
Cleveland State University

Joseph E. Harrington, Jr.
Johns Hopkins University

The authors explore the evolution of the structure and performance of a social network in a population of individuals who search for local optima in diverse and dynamic environments. Individuals choose whether to innovate or imitate, and in the latter case, from whom to learn. The probabilities of these possible actions respond to an individual's past experiences using reinforcement learning. Among some of the authors' more interesting findings is that a population's performance is not monotonically increasing in either the reliability of the communication network or the productivity of innovation.

INTRODUCTION

The acceleration of scientific progress in 17th-century Western Europe is often attributed to the founding of learned societies and to improved communication among contemporary researchers: "Here was a widely dispersed population of intellectuals, working in different lands, using different vernaculars—and yet a community. What happened in one place was quickly known everywhere else, partly thanks to a common language of learning, Latin; partly to a precocious development of courier and mail services; most of all because people were moving in all directions. In the

[1] The comments of Rich Burton, Jon Harford, Nicolaj Siggelkow, the *AJS* referees, seminar participants at DePaul, and conference participants at the 2002 CASOS Conference (Carnegie Mellon) and the 2004 SCE Workshop on Complex Behavior in Economics (Aix-en-Provence, France) are gratefully acknowledged as are general discussions on the topic of networks with Rob Axtell and Steve Durlauf. This research is supported in part by the National Science Foundation (grant SES-0078752). Direct correspondence to Myong-Hun Chang, Department of Economics, Cleveland State University, Cleveland, Ohio, 44115. E-mail: m.chang@csuohio.edu

0002-9602/2005/11004-0003$10.00

seventeenth century, these links were institutionalized . . . in the form of learned societies with their corresponding secretaries, frequent meetings, and periodical journals" (Landes 1998, pp. 204–5). The emergence of such formal and informal networks of scientists, triggered by improvements in communication technology and transportation, marked the beginning of the institutionalization of scientific investigation which eventually paved the way to the Industrial Revolution.

The latter half of the 20th century has seen another major explosion in the formation of such networks and informal communities. Thanks to the advent of the Internet and the World Wide Web, the time delay in acquiring information has shrunk from weeks (or at best, days and at worst, months) to minutes. In the communities of research scientists, learning what others are working on and what methods they are deploying previously required mailing letters requesting working papers, frequent trips to professional meetings and workshops, and laborious search in the dust-covered library stacks.[2] All this has changed dramatically in recent years. A letter sent using the postal service is replaced with an e-mail which arrives at the desktop of the other party within seconds, or by directly downloading the paper from the individual's Web site. What is more, there exist ongoing projects that facilitate dissemination of research over the Internet via a decentralized database of working papers, journal articles, and software components.[3] Downloading a working paper is now merely a matter of pointing a cursor and clicking. The library, a physical institution that was once the hallmark of academic community, is gradually disappearing into the network of connected computers which allow us to search, check out, and read material at our desktop PCs.

The 300 years of separation notwithstanding, these two episodes share a common element: the prominent role played by the social learning network in achieving and sustaining scientific progress at the macrolevel from the chance occurrences of local innovations scattered across wide geographic areas. Furthermore, these issues are not unique to scientific processes but also apply to markets and organizations, for what at work

[2] See Liberman and Wolf (1997) who investigate knowledge flows from conferences.

[3] RePEc (Research Papers in Economics) at http://repec.org, an international collaborative project in economics, links various sources of research publications in over 30 countries and enable their distribution via electronic media. As of November 2002, they report that the "RePEc database holds over 176,000 items of interest, over 85,000 of which are available online." Also, as of November 2002, SSRN (Social Science Research Network) at http://www.ssrn.com claims to have in its eLibrary an abstract database of over 44,900 working papers and forthcoming papers as well as an electronic paper collection of over 25,400 downloadable full-text documents.

is quite fundamental—finding better solutions to a stream of problems.[4] Given this relationship between local innovations and social learning via community networks, what are the implications for an individual's decision to engage in independent innovation or to access the social network for observing the ideas of others? If the networks themselves are an outcome of interactive choices among individuals as to whom to observe and whom to ignore, what are the determinants of their emergent structure? Will an improvement in communication technology be sufficient to generate superior performance at the individual and the community level, or are there other complementary factors essential for a social network to realize these potential gains? Our objective is to explore these issues by building an agent-based computational model of an evolving social network with potentially innovative individuals and analyzing its emergent structure as well as its long-run performance.

Given the richness of these processes, the model we develop is stylized and is designed to encompass three generic features of these knowledge processes. First, agents are faced with problems to solve which we model as an effort to achieve a local optimum in the space of possible things that one can do.[5] This optimum is uniquely defined for each individual, where similarity among optima is to be interpreted as similarity in the problems being solved. Second, agents must choose how to allocate their effort between individual learning (innovation) and social learning (imitation). When they engage in learning from others, an agent decides from whom to learn, which takes the form of establishing links in a social network. These choices are made probabilistically, and the probabilities are adjusted over time based on personal experience. This modeling structure allows us to examine the emergent structure of the social network in terms of how observation probabilities are distributed across individuals as well as to track the evolving choice between innovation and imitation for each individual. Our model also captures the fact that the

[4] See Podolny, Stuart, and Hannan (1996) for an analysis of knowledge production via a network in the semiconductor industry.

[5] In defense of this view, we appeal to Kuhn's (1962) interpretation of "normal science as puzzle-solving" in which a scientist is someone driven by the desire to search for a solution (a goal yet unknown) to the problem at hand: "Bringing a normal research problem to a conclusion is achieving the anticipated in a new way, and it requires the solution of all sorts of complex instrumental, conceptual, and mathematical puzzles. The man who succeeds proves himself an expert puzzle-solver, and the challenge of the puzzle is an important part of what usually drives him on" (Kuhn 1962, p. 36).

actual outcomes of these efforts tend to have a substantial random aspect.[6] Third, knowledge creation and diffusion occurs in the context of a changing environment as represented by stochastic movement in optima. This feature captures the fact that individuals do not typically solve one problem once and for all, but instead must solve a sequence of problems which, to varying degrees, can at least partially be solved using solutions to previous problems.

This model is used to explore how the innovativeness of individuals and the reliability of the communication technology impact network structure and performance, and how those relationships depend on the characteristics of the environment. Several interesting properties emerge. When communication technology is sufficiently poor, the results are quite intuitive. For example, performance is improved by enhancing the reliability of communication and making agents more productive in innovation. However, when communication is sufficiently efficacious, further enhancement of it can actually be detrimental to performance. The intuition is that the increased social learning from a more reliable network leads to local homogenization of agents in terms of their solutions to problems. This lack of diversity within the social network results in a population of agents who are ill equipped to adapt to a changing environment. Second, when reliability is of moderate quality, a rise in the productivity of innovation can be deleterious to average performance in the population. The key insight here is that imitating others is both a social good and a social detriment. Engaging in imitation fails to add to the stock of knowledge, but it does serve to spread worthwhile ideas. Thus, an agent who imitates is improving the value of the network, which can enhance collective performance. When network reliability is moderate, agents may be engaging in too much individual learning from a societal perspective; it might be better to tap into the network to pass along ideas and develop more useful links (which can only be achieved through experience). Making agents more innovative induces them to use the network even less and thereby exacerbates the problem of a poorly developed network and inadequate sharing of ideas. A third property is that when network reliability is low, the social network is less structured when agents are more innovative. The simple reason is that they choose

[6] See Mokyr (1990, p. 152): "Over most of human history, technological change did not take place, as it does today, in specialized research laboratories paid for by research and development budgets and following strategies mapped out by corporate planners well-informed by marketing analysts. Technological change occurred mostly through new ideas and suggestions occurring if not randomly, then certainly in a highly unpredictable fashion. Demand conditions may have affected the rate at which these ideas occurred, and may have focused them in a particular direction, but they did not determine whether a society would be technologically creative or not."

to engage in more individual learning. Since they then access the network less, they are less effective in developing worthwhile links. When network reliability is high, however, those two measures—the structure of the social network and the capacity for innovation—move together. This is driven by the fact that the quality of a network, in terms of the population's possessing a diverse array of useful ideas, rises with the productivity of innovation. But that quality cannot be adequately tapped if reliability is low. A more structured network emerges when both the reliability and the quality (which depends on the capacity for innovation) of the network are high. While innovation and imitation are substitutes at the individual level, they are complements at the population level through the mechanism of social learning.

Related Work

While there is a wealth of research exploring the implications of a network, there is much less formal theoretical work that endogenously derives social learning networks. An early exception is the work of Carley (1990, 1991). She develops a simulation model in which interagent interaction leads to shared knowledge, which then determines the likelihood of further interactions. Her model differs from ours in two crucial ways. First, she concentrates purely on the social network as the learning mechanism; individuals do not innovate. The dynamic behavior of the network is then solely driven by the endowed differences in knowledge across individuals. Second, she specifies the likelihood of interaction between two agents as being determined directly by the amount of information they share; the larger the set of common information, the higher the probability that they will interact again. In contrast, our model takes a more bottom-up approach. We assume that the agents are motivated by private goals. For each agent, the probability of observing another agent is adjusted on the basis of the past value of other agents' information in attaining her goal. This approach highlights our perspective that the likelihood of an interaction between two individuals is an *emergent* outcome and should be modeled as such.

More recent work is based on the specification of goals for agents, and it is the striving for attainment of these goals ("utility maximization") that determines a network's structure. Jackson and Wolinsky (1996) characterize the properties of a stable network—one in which no agent wants to create or destroy links—while the dynamics of network formation is explored in Bala and Goyal (2000), Hummon (2000), and Jackson and Watts (2002). They assume agents adaptively adjust their networks so as to improve performance. Though utilizing a distinct approach, a predecessor to these discussions is Huberman and Hogg (1995), but it differs from our model in several substantive ways. First, the payoff to an agent from a given network is exogenously fixed, while in our model, it evolves with what an agent can learn from other agents. Second, it does not consider an agent's choice of innovation versus imitation, but focuses exclusively on network structure. Third, Huberman and Hogg's study does not allow agents to adapt their search process in response to experience, while in our model, this is done through reinforcement learning.

THE MODEL

Agents, Tasks, Goals, and Performance

Consider a social system consisting of M individuals. Each individual $i \in \{1, 2, \ldots, M\}$ engages in an operation that can be broken down into N separate tasks. There are several different methods that can be used for each task. The method chosen by an agent for a given task is represented by a sequence of d bits (0 or 1) such that there are 2^d possible methods available for each task. In any period t, an individual i is then fully characterized by a binary vector of $N \cdot d$ dimensions. Denote it by $\underline{z}_i(t) \in \{0, 1\}^{Nd}$ so that $\underline{z}_i(t) \equiv (\underline{z}_i^1(t), \ldots, \underline{z}_i^N(t))$ and $\underline{z}_i^h(t) \equiv (z_i^{h,1}(t), \ldots, z_i^{h,d}(t)) \in \{0, 1\}^d$ is individual i's chosen method in task $h \in \{1, \ldots, N\}$. An example with $N = 24$ and $d = 4$ is given below:

task (h)	#1	#2	#3	. . .	#24
task methods ($\underline{z}_i^h(t)$) :	1101	0010	1001	. . .	1110

Note that there are 16 ($= 2^4$) different methods or bit configurations for each task. What is shown above for a given task represents a particular method chosen out of the 16 available methods. Given that the operation is completely described by a vector of 96 ($= 24 \times 4$) bits, there are then 2^{96} possible bit configurations for the overall operation.

In measuring the degree of heterogeneity between two methods vectors, $\underline{z}_i$ and $\underline{z}_j$, we shall use "Hamming distance," which is defined as the number of positions for which the corresponding bits differ:

$$D(\underline{z}_i, \underline{z}_j) \equiv \sum_{h=1}^{N} \sum_{k=1}^{d} |z_i^{h,k} - z_j^{h,k}|. \quad (1)$$

Each individual possesses a goal vector which may be different from one period to the next. Let $\hat{\underline{z}}_i(t) \in \{0, 1\}^{Nd}$ be the goal vector of agent i in period t. That $\hat{\underline{z}}_i(t)$ can differ across agents implies diversity in agents' problems, or alternatively, in their environments, so that the optimum differs. The degree of turbulence in task environments is captured by intertemporal variability in $\hat{\underline{z}}_i(t)$.

The individuals are uninformed about $\hat{\underline{z}}_i(t)$ *ex ante*, but engage in search to get as close to it as possible. Given N tasks with d bits in each task and the goal vector $\hat{\underline{z}}_i(t)$, the period-t performance of individual i is then measured by $\pi_i(t)$, where

$$\pi_i(t) \equiv N \cdot d - D(\underline{z}_i(t), \hat{\underline{z}}_i(t)). \tag{2}$$

The performance of a social system is measured by how close the individuals are to their respective goals. We let $\bar{\pi}(t)$ denote the population-average performance in period t,

$$\bar{\pi}(t) \equiv \frac{1}{M}\sum_{i=1}^{M} \pi_i(t). \tag{3}$$

Modeling Innovation and Imitation

In any given period, an individual's search for an optimum is carried out through two distinct mechanisms, innovation and imitation. Innovation occurs when an individual independently discovers and considers for implementation a random method for a randomly chosen task. Imitation is when an individual selects someone and then observes and considers for implementation some randomly chosen task currently deployed by that agent.[7] Though a single act of innovation or imitation is assumed to be a single task, this is without loss of generality. If we choose to define a task to include d' dimensions, then the case of a single act of innovation or imitation involving two tasks can be handled by setting $d = 2d'$.[8] In essence, what we are calling a "task" is defined to be the unit of discovery or observation. The actual substantive condition is instead the relationship between d and N, as an agent's innovation or imitation involves a smaller part of the possible solution when d/N is smaller.

Whether obtained through innovation or imitation, an experimental method is actually adopted if and only if its adoption gets the agent closer to her goal by lowering the Hamming distance between her new methods

[7] It may be useful to think of innovation and imitation as being analogous to *mutation* and *crossover*, respectively, in evolutionary biology.

[8] There is a restriction in that an agent only has the option of adopting all d dimensions or none.

vector and her goal vector. For clarity, consider the following example with $N = 5$ and $d = 2$:

agent i's goal vector: 01|10|10|01|01

agent i's current methods vector: 01|01|11|00|11

The relevant operation has five tasks. In each task, there are four distinct methods that can be tried: (0,0), (0,1), (1,0), and (1,1). The Hamming distance between i's current methods vector and her goal vector is then five. Suppose i observes the method used for task 4 by another agent j ($\neq i$) whose methods vector is:

agent j's current methods vector: 10|10|11|01|01

Since j's method in task 4 is (0,1), when it is tried by agent i, her experimental methods vector becomes:

agent i's experimental methods vector: 01|01|11|01|11

which then reduces the Hamming distance to i's goal to four, hence, the experimental methods vector becomes i's new methods vector.

Endogenizing Choices for Innovation and Imitation

In each period, an individual may engage in either innovation or imitation by using the network. Exactly how does an individual choose between innovation and imitation, and if she chooses to imitate, how does she decide whom to imitate? We model this question as a two-stage stochastic decision process with reinforcement learning. Figure 1 describes the timing of decisions in our model. In stage 1 of period t, individual i is in possession of the current methods vector, $\underline{z}_i(t)$, and chooses to innovate with probability $q_i(t)$ and imitate with probability $1 - q_i(t)$. It she chooses to innovate then, with probability μ_i^{in}, she comes up with an idea which is a randomly chosen task, $h \in \{1, \ldots, N\}$, and a randomly chosen method, $\underline{z}_i^{h'}$, for that task such that the experimental method vector is $\underline{z}_i'(t) \equiv (\underline{z}_i^1(t), \ldots, \underline{z}_i^{h-1}, \underline{z}_i^{h'}, \underline{z}_i^{h+1}, \ldots, \underline{z}_i^N(t))$. $\mu_i^{in}$is a parameter that controls the productivity of an agent's innovation. This experimental vector is adopted by i if and only if its adoption lowers the Hamming distance to her current goal vector, $\hat{\underline{z}}_i(t)$. Otherwise, it is discarded:

$$\underline{z}_i(t+1) = \begin{cases} \underline{z}_i'(t) & \text{if } D(\underline{z}_i'(t), \hat{\underline{z}}_i(t)) < D(\underline{z}_i(t), \hat{\underline{z}}_i(t)) \\ \underline{z}_i(t) & \text{if } D(\underline{z}_i'(t), \hat{\underline{z}}_i(t)) \geq D(\underline{z}_i(t), \hat{\underline{z}}_i(t)). \end{cases} \quad (4)$$

Alternatively, with probability $1 - \mu_i^{in}$, the individual fails to generate an idea, in which case $\underline{z}_i(t+1) = \underline{z}_i(t)$.

Now suppose individual i chooses to imitate in stage 1. Given that she

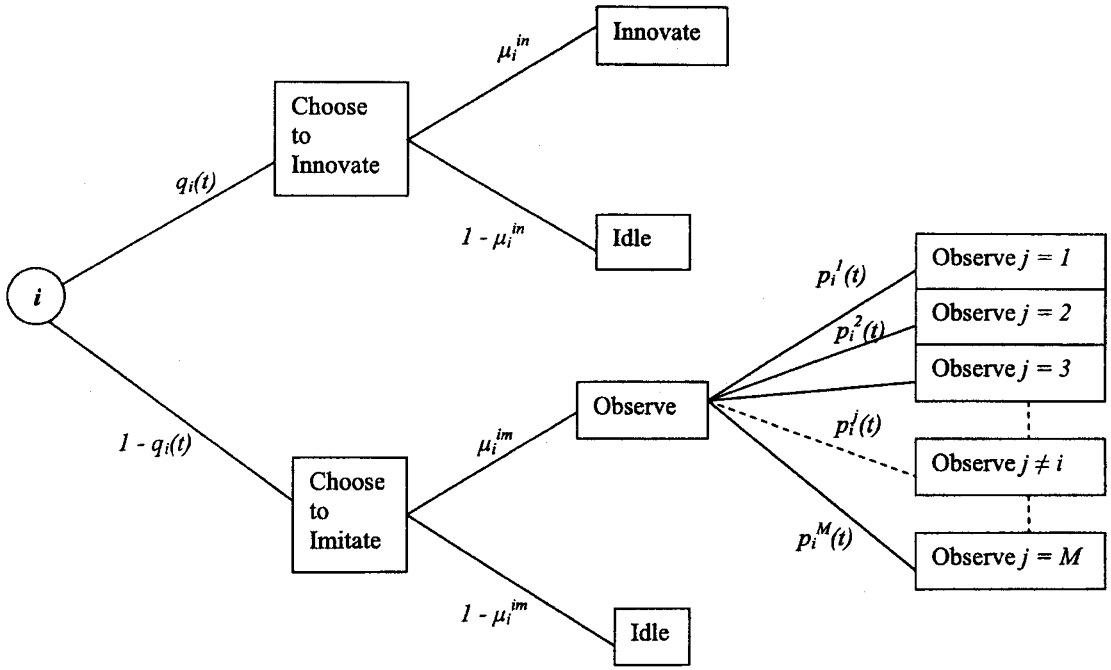

FIG. 1.—Decision sequence of individual i in period t

decides to imitate someone else, she taps into the network to make an observation. Tapping into the network is also a probabilistic event, in which with probability μ_i^{im}, the agent is connected to the network, while with probability $1 - \mu_i^{im}$, the agent fails to do so. μ_i^{im} measures the reliability of the communication technology and thereby the network. An agent that is connected then enters stage 2 of the decision process in which another agent must be selected to be studied and possibly imitated. Let $p_i^j(t)$ be the probability with which i observes j in period t so $\sum_{j \neq i} p_i^j(t) = 1$ for all i. If agent i observes another agent l, the observation involves a randomly chosen task h and the current method used by agent l in that task, $\underline{z}_l^h(t)$. Let $\underline{z}_i''(t) = (\underline{z}_i^1(t), \ldots, \underline{z}_i^{h-1}(t), \underline{z}_l^h(t), \underline{z}_i^{h+1}(t), \ldots, \underline{z}_i^N(t))$ be the experimental vector. Adoption or rejection of the observed method is based on the Hamming distance criterion:

$$\underline{z}_i(t+1) = \begin{cases} \underline{z}_i''(t) & \text{if } D(\underline{z}_i''(t), \hat{\underline{z}}_i(t)) < D(\underline{z}_i(t), \hat{\underline{z}}_i(t)) \\ \underline{z}_i(t) & \text{if } D(\underline{z}_i''(t), \hat{\underline{z}}_i(t)) \geq D(\underline{z}_i(t), \hat{\underline{z}}_i(t)). \end{cases} \tag{5}$$

If the agent fails to connect to the network, which occurs with probability $1 - \mu_i^{im}$, $\underline{z}_i(t+1) = \underline{z}_i(t)$.

The probabilities $q_i(t)$ and $\{p_i^1(t), \ldots, p_i^{i-1}(t), p_i^{i+1}(t), \ldots, p_i^M(t)\}$ are adjusted over time by individual agents according to a reinforcement learning rule. We adopt a version of the *experience-weighted attraction (EWA)* learning rule as described in Camerer and Ho (1999).[9] Using this rule, $q_i(t)$ is adjusted each period on the basis of evolving attraction measures, $B_i^{in}(t)$ and $B_i^{im}(t)$, for innovation and imitation, respectively. The evolution of $B_i^{in}(t)$ and $B_i^{im}(t)$ follows the process below:

$$B_i^{in}(t+1) = \begin{cases} \phi B_i^{in}(t) + 1 & \text{if } i \text{ adopted a method through innovation in } t \\ \phi B_i^{in}(t) & \text{otherwise,} \end{cases} \tag{6}$$

$$B_i^{im}(t+1) = \begin{cases} \phi B_i^{im}(t) + 1 & \text{if } i \text{ adopted a method through imitation in } t \\ \phi B_i^{im}(t) & \text{otherwise,} \end{cases} \tag{7}$$

where $\phi \in (0, 1]$. Hence, if the agent chose to pursue innovation and discovered and then adopted her new idea, then the attraction measure for innovation increases by one after allowing for the decay factor of ϕ on the previous attraction level. If innovation was pursued but was unsuccessful (which occurs if either she failed to generate an idea or suc-

[9] For a discussion of reinforcement learning mechanisms in general, see Sutton and Barto (2000).

ceeded in generating an idea which turned out to be useless), or the agent chose to pursue imitation instead, then her attraction measure for innovation is simply the attraction level from the previous period decayed by the factor ϕ. Likewise, a success or failure in imitation in t has the same influence on $B_i^{im}(t+1)$. Given $B_i^{in}(t)$ and $B_i^{im}(t)$, one then derives the choice probability of innovation in period t as follows:

$$q_i(t) = \frac{(B_i^{in}(t))^\lambda}{(B_i^{in}(t))^\lambda + (B_i^{im}(t))^\lambda} \tag{8}$$

where $\lambda > 0$. A higher value for λ means that a single success has more of an impact on the likelihood of repeating that activity (innovation or imitation).[10] The probability of imitation is, of course, $1 - q_i(t)$. The expression in equation (8) tells us that a favorable experience through innovation (imitation) raises the probability that an agent will choose to innovate (imitate) again in the future—a positive outcome realized from a course of action reinforces the likelihood of that same action's being chosen again.

The stage-2 attractions and probabilities are derived similarly. Let $A_i^j(t)$ be agent i's attraction to another agent j in period t. It evolves according to the rule below:

$$A_i^j(t+1) = \begin{cases} \phi A_i^j(t) + 1 & \text{if } i \text{ successfully imitated } j \text{ in } t \\ \phi A_i^j(t) & \text{otherwise,} \end{cases} \tag{9}$$

$\forall j \neq i$. The probability that agent i observes agent j in period t is adjusted each period on the basis of the attraction measures, $\{A_i^j(t)\}_{j \neq i}$:

$$p_i^j(t) = \frac{(A_i^j(t))^\lambda}{\sum_{h \neq i} (A_i^h(t))^\lambda} \tag{10}$$

$\forall j \neq i$, $\forall i$, where $\lambda > 0$.

There are two distinct sets of probabilities in our model. One set of probabilities, $q_i(t)$ and $\{p_i^j(t)\}_{j \neq i}$, are endogenously derived and evolve over time in response to the personal experiences of agent i. Another set of probabilities, μ_i^{in} and μ_i^{im}, are exogenously specified and imposed on the model as parameters. They represent the state of existing technologies and control the potential capabilities of individual agents to innovate independently or to imitate someone else in the population via social

[10] For analytical simplicity, we assume φ and λ to be common to all individuals in the population.

learning.[11] Of particular interest is understanding how these parameters influence the structure and performance of the network and the rate of innovation which we will measure by the population-average level of innovation:

$$\bar{q}(t) \equiv \frac{1}{M}\sum_{i=1}^{M} q_i(t). \tag{11}$$

Modeling Turbulence in Task Environment

Central to the performance of a population is how it responds to an evolving environment, or, if we cast this in the context of problem solving, an evolving set of problems to be solved. It is such change that makes innovation and the spread of those innovations through a social network so essential. Change or turbulence is specified in our model by first assigning an initial goal vector, $\underline{\hat{z}}_i(0)$, to each agent and then specifying a dynamic process by which it shifts over time. In order to provide the possibility of a substantive network's forming, it is critical to allow for some persistent similarity in goals across subsets of agents. This is achieved by initially partitioning the population into a fixed number of groups; those agents belonging to the same group tend to have more similar goals—they are working on similar problems—than those belonging to different groups. As such, for any given agent, there are two broad sources for social learning—another agent in the same group and another agent in a different group. We expect the efficacy of social learning to depend on the tightness of the goals within a given group relative to the tightness of the goals between different groups. To explore this issue, we distribute the initial goal vectors of agents in a sequence of steps described below.

Letting $s \in \{0, 1\}^{Nd}$, define $\delta(s, \kappa) \subset \{0, 1\}^{Nd}$ as the set of points that are exactly Hamming distance κ away from s. The set of points *within* Hamming distance κ of s is defined as

$$\Delta(s, \kappa) \equiv \bigcup_{i=0}^{\kappa} \delta(s, i). \tag{12}$$

[11] One may view these as being completely determined by available technologies outside of our model. For instance, the size of the manuscript libraries may have limited the magnitude of μ^{im} in ancient Alexandria, while the invention of the printing press presumably raised it in the 15th century: "Limits set by the very largest manuscript libraries were also broken. Even the exceptional resources which were available to ancient Alexandrians stopped short of those that were opened up after the shift from script to print. The new open-ended information-flow that commenced in the fifteenth century made it possible for fresh finding to accumulate at an ever-accelerated pace" (Eisenstein 1979, pp. 517–18).

$\Delta(s, \kappa)$ is a set whose "center" is s.

Suppose there are J groups in the population and let us randomly allocate the M agents into these groups. Let a_k be the set of agents belonging to group $k \in \{1, 2, \ldots, J\}$. We define g_k as the seed vector used to generate the initial goal vectors for all agents in a_k,

$$\underline{\hat{z}}_i(0) \in \Delta(g_k, \kappa) \quad \forall i \in a_k, \forall k \in \{1, 2, \ldots, J\}. \tag{13}$$

All agents in a_k then have goal vectors that lie within Hamming distance κ of the group seed vector g_k. The diversity among groups is modeled by allowing their group seed vectors to differ. Specifically, we define a seed vector U for the entire population and randomly select the group seed vectors from $\Delta(U, X)$; the intergroup tightness of the goals is controlled through X, the maximum Hamming distance between a group seed vector and the population seed vector. Of course, the intragroup tightness of the goals is controlled with κ as described in (13). Figure 2 depicts how these sets are related to one another. With $J = 4$, the four group seed vectors are chosen from the set $\Delta(U, X)$. Taking the seed vector for group 2, g_2, $\Delta(g_2, \kappa)$ is the set of vectors that are within Hamming distance κ of g_2. The initial goal vector for individual i, $\underline{\hat{z}}_i(0)$, is then an element of this set.

In period t, agent i has the current goal vector of $\underline{\hat{z}}_i(t)$. In period $t + 1$, her goal stays the same with probability σ and changes with probability $(1 - \sigma)$. The shift dynamic of the goal vector is guided by the following stochastic process. The goal in $t + 1$, if different from $\underline{\hat{z}}_i(t)$, is then chosen *iid* from the set of points that lie both within the Hamming distance ρ of $\underline{\hat{z}}_i(t)$ and within Hamming distance κ of the original group seed vector g_k. Hence, defining $\Lambda(\underline{z}_i(t), \rho, g_k, \kappa)$ as the set of points from which the goal in $t + 1$ is chosen, we have

$$\Lambda(\underline{\hat{z}}_i(t), \rho, g_k, \kappa) \equiv (\Delta(\underline{\hat{z}}_i(t), \rho) \backslash \underline{\hat{z}}_i(t)) \cap \Delta(g_k, \kappa). \tag{14}$$

Figure 2 shows $\Lambda(\underline{\hat{z}}_i(t), \rho, g_k, \kappa)$ as the doubly shaded area which is the intersection of $\Delta(\underline{\hat{z}}_i(t), \rho)$ and $\Delta(g_k, \kappa)$, minus $\underline{\hat{z}}_i(t)$. Consequently,

$$\begin{cases} \underline{\hat{z}}_i(t + 1) = \underline{\hat{z}}_i(t) & \text{with probability } \sigma \\ \underline{\hat{z}}_i(t + 1) \in \Lambda(\underline{\hat{z}}_i(t), \rho, g_k, \kappa) & \text{with probability } 1 - \sigma. \end{cases} \tag{15}$$

The goal vector for agent i who belongs to group k then stochastically fluctuates while remaining within Hamming distance ρ of his current goal *and* Hamming distance κ of the group's initial seed vector. The former condition allows us to control the possible size of the change while the

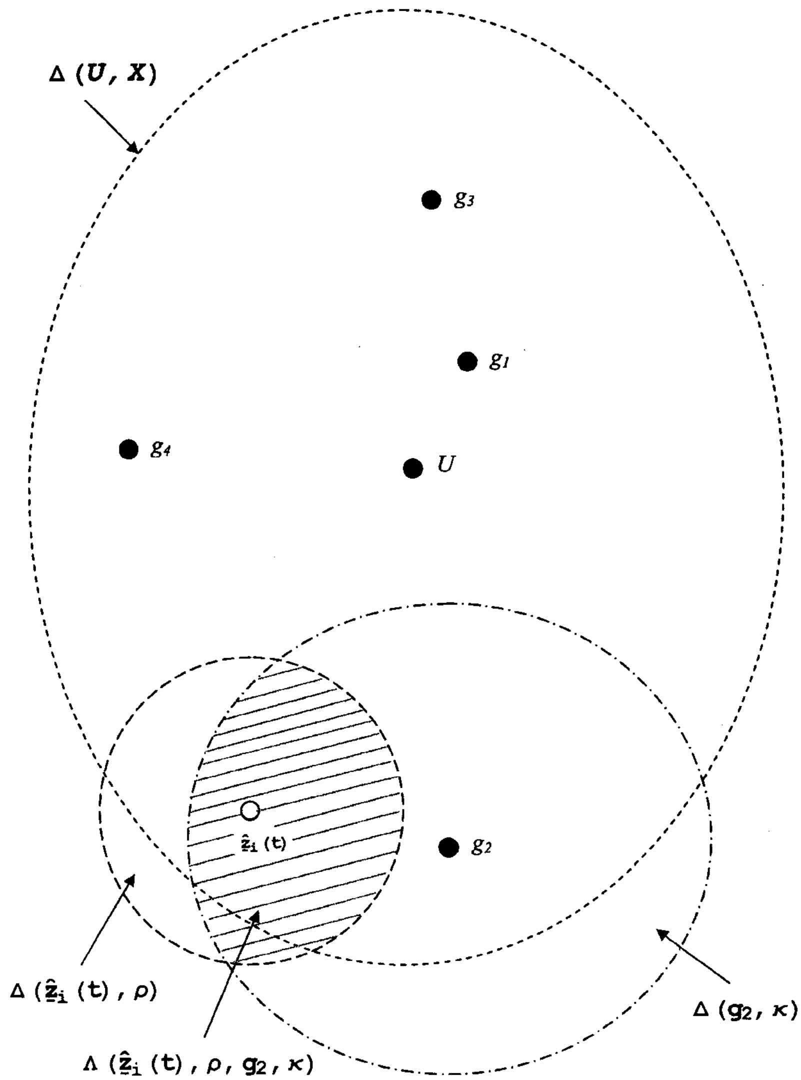

Fig. 2

latter condition allows us to maintain the intragroup tightness of goals.[12]

[12] If we had instead allowed an agent's goal to be a random walk without constraining it to $\Delta(g_k, \kappa)$, then eventually, the distance between one agent's goal and another agent's goal would be independent of whether they are in the same group so that intragroup tightness would have been lost.

The lower σ and the greater ρ, the more frequent and variable is the change, respectively, in an agent's goal vector. The higher κ is, the lower the intragroup goal congruence. The higher X is, the greater the intergroup diversity in terms of their goals.

Measuring the Network Structure

A social network emerges when individuals rely on observation and imitation of others. A main structural characteristic of a social network is its concentration: Does an individual learn from many or from a relatively narrow set of other individuals? In our context, this question can be addressed by observing the distribution of $p_i^j(t)$'s. If an individual is equally likely to imitate any other member of the population—so that $p_i^j(t) = 1/(M-1)$—there is no order in the social network as imitation is completely (that is, uniformly) random. Alternatively, if the probability of observing another agent is concentrated on a single individual—so that $p_i^j(t) = 1$ for some j—then there is a maximal degree of order in the network.

An appropriate measure for this purpose is Shannon's (1948) "entropy" which was originally defined in the context of information theory as an inverse measure of the information content of a message. In adapting this measure for the context at hand, the entropy measure for the social network of agent i is defined to be:

$$E_i(t) \equiv -\sum_{\forall j \neq i} p_i^j(t) \cdot \log_2 p_i^j(t). \tag{16}$$

The value for $E_i(t)$ can range from a minimum of zero to a maximum of $\log_2(M-1)$.[13] By taking an average of $E_i(t)$ over all individuals in the population, we obtain the *mean entropy* of the network system:

$$\overline{E}(t) \equiv \frac{1}{M}\sum_i E_i(t). \tag{17}$$

We say that the population network structure becomes more ordered (random) as $\overline{E}(t)$ decreases (increases).

[13] $E_i(t) = 0$ when the probability mass is entirely concentrated on one particular individual $k \neq i$ so that $p_i^k(t) = 1$ and $p_i^j(t) = 0$ $\forall j \neq k, i$. Conversely, maximal entropy is attained at $E_i(t) = \log_2(M-1)$ if observation is equally likely for all individuals, $p_i^j(t) = 1/(M-1)$ $\forall j \neq i$.

SIMULATION DESIGN

We consider a population of 20 individuals with identical capacities for learning:[14] $\mu_i^{in} = \mu^{in}$ and $\mu_i^{im} = \mu^{im}$ $\forall i$. There are four ($J = 4$) distinct groups into which these agents are allocated. Assuming $N = 24$ and $d = 4$, there are 96 total bits in a methods vector.[15] The initial methods vectors, $\underline{z}_i(0)$'s, are independent draws from $\{0, 1\}^{Nd}$. In any time period t, the state variables for individual i are then $\underline{z}_i(t)$, $\hat{\underline{z}}_i(t)$, $\{A_i^j(t)\}_{\forall j \neq i}$, $B_i^{in}(t)$, and $B_i^{im}(t)$. The parameters of the model are μ^{im}, μ^{in}, ϕ, and λ, which govern an individual agent's decision making, and X, κ, σ, and ρ, which control the environment by specifying intragroup and intergroup goal congruence as well as the dynamics of the task environments. So that results are not driven by the peculiarities of the initial methods vectors, we will focus on the steady-state behavior of the social system with the particular intent of understanding how it depends on these agent and environment parameters. The only exceptions are that $\phi = 1$ and $\lambda = 1$ and remain fixed. Finally, the initial attraction stocks are set at $A_i^j(0) = 1$ $\forall i$, $\forall j \neq i$, and $B_i^{in}(0) = B_i^{im}(0) = 1$ $\forall i$. Hence, an individual is initially equally attracted to innovation and imitation and has no inclination to observe one individual over another *ex ante*. Table 1 provides a comprehensive list of parameters along with the set of values considered over the course of the simulations.

In each case, the model is run for 20,000 periods. For each period, $q_i(t)$ and $\{p_i^j(t)\}_{j \neq i}$ are collected as well as the performance of each individual, $\pi_i(t)$. Population averages were then computed for these time series, thereby giving us $\overline{E}(t)$, $\bar{q}(t)$, and $\bar{\pi}(t)$ for each replication. For each parameter configuration, we performed 20 replications—each of 20,000 periods in length—where we had fresh realization of the random variables including the initial method vectors, the agents' choices (recall that they are probabilistic), the outcome of innovation and imitation, the seed vectors, and agents' optima. The time series for the relevant variables, such as the entropy and performance measures, were then averaged over the 20 replications so as to generate the final time series reported here. Hereinafter, $\overline{E}(t)$, $\bar{\pi}(t)$, and $\bar{q}(t)$ will denote the final time series averaged over these 20 replications.

STRUCTURE OF NETWORKS

For the baseline simulation, we considered $\mu^{in} = .5$ and $\mu^{im} = .5$ so that when in the innovation mode, an agent generates an idea 50% of the

[14] The source code for the simulation design, written in C^{++}, is available upon request from Myong-Hun Chang.

[15] The search space then contains over 7.9×10^{28} possibilities.

TABLE 1
List of Parameters

Notation	Definition	Baseline Value	Parameter Value	Parameter Values Considered
μ^{in}	Exogenous rate of innovation	.5		{0, .25, .5, .75, 1}
μ^{im}	Exogenous rate of imitation	.5		{0, .05, .1, .15, .2, . . ., .9, .95, 1}
X	Intergroup goal diversity	16		{0, 4, 8, 16, 32, 64, 96}
κ	Intragroup goal diversity	16		{4, 8, 16, 32}
σ	Intertemporal goal stability	.75		{.5, .75, .95, .99}
ρ	Intertemporal goal variability	4		{1, 4, 8}
M	Number of agents		20	
J	Number of groups		4	
ϕ	Attraction decay factor		1	
λ	Agent's sensitivity to attraction		1	
$A_i^j(0)$, $\forall i$, $\forall j \neq i$	i's attraction to j in $t = 0$		1	
$B_i^{in}(0)$, $\forall i$	i's attraction to innovation in $t = 0$		1	
$B_i^{im}(0)$, $\forall i$	i's attraction to imitation in $t = 0$		1	

time, and when in the imitation mode, an agent observes another agent 50% of the time. The stability of the task environment is set at $\sigma = .75$, so that the local optimum for an individual shifts 25% of the time, and $\rho = 4$ so that in case of such a shift, four or less randomly chosen bits in the goal vector will flip. We also assume moderate intragroup and intergroup diversity by setting $\kappa = 16$ and $X = 16$, respectively.

Figure 3 shows the results from the baseline simulation of the endogenous network. These results are typical based on many other parameter configurations. The mean performance of the population, shown in figure 3(A), increases rapidly early on and then converges to a steady state. Toward understanding how the social network evolves, figure 3(B) plots the entropy measure. Recall that at $t = 0$, the attractions are all identically equal to one, so consequently, $p_i^j(0) = 1/(M - 1)\ \forall j \neq i,\ \forall i$. The social network is then starting from a point of maximum entropy of 4.25. As is shown in figure 3(B), it then monotonically declines to about 3.64 by the end of the 20,000 periods; 3.64 is equivalent to the entropy of a fully random network with only 13.5 agents.[16] The network is then becoming increasingly ordered as each agent's imitation is being concentrated on an ever smaller set of other agents.

Given that agents focus their attention on an increasingly narrower set of other agents, to whom are they attracted? Is learning mutual so that agent i, who chooses to learn from agent j, is also a primary source of knowledge for agent j? How strong is the tendency for mutual learning within groups? These questions can be answered only by going beyond the measure of entropy and delving into the relationships between the stage-2 choice probabilities of agents in the population. For this purpose, we impose extra structure on the simulation by randomly allocating 20 individuals to four groups and then holding their group memberships fixed across the 20 replications.[17] The output of this exercise is reported in figure 4 and table 2. Letting $\bar{p}_i^j(t)$ denote the time series of $p_i^j(t)$ averaged

[16] As there are five agents in each agent's group, it is not then true that each of these groups form their own networks. This does not occur for a variety of reasons. First, it is possible that agents in other groups might be more similar though that will not hold on average. Second, even if agents outside of one's group are less similar than those in one's group, randomness in what agents adopt and what is observed could cause an agent to learn more from nongroup agents and thus induce the establishment of stronger links with them.

[17] This structure is imposed only on the simulations reported in this section. For all other simulations, the individuals are reallocated (randomly) to the groups at the start of each replication.

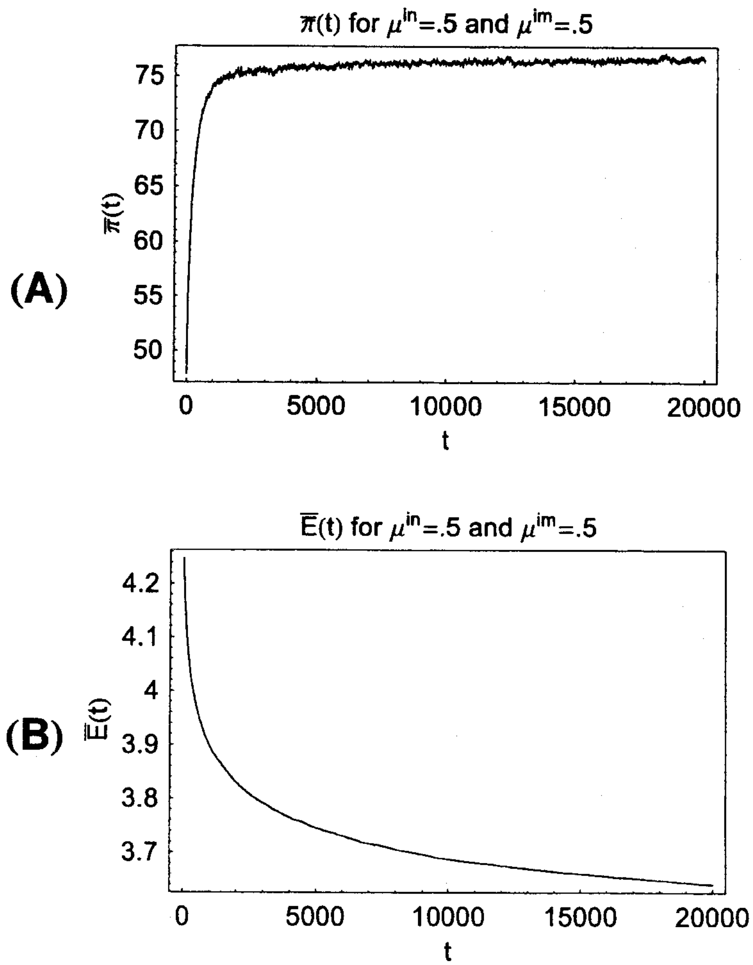

FIG. 3.—Baseline time series

over the 20 replications, agent i's steady-state choice probability of observing agent j is defined by:

$$\tilde{p}_i^j \equiv \frac{1}{5000} \sum_{t=15,001}^{20,000} \bar{p}_i^j(t). \tag{18}$$

Figure 4 plots pairs of choice probabilities, $(\tilde{p}_i^j, \tilde{p}_j^i)$, as points in a probability space for all i and j, $i \neq j$. Since there are 20 agents in our experiment and each agent can then observe 19 other agents, there are 190 distinct points in total. For each point, the horizontal coordinate is the

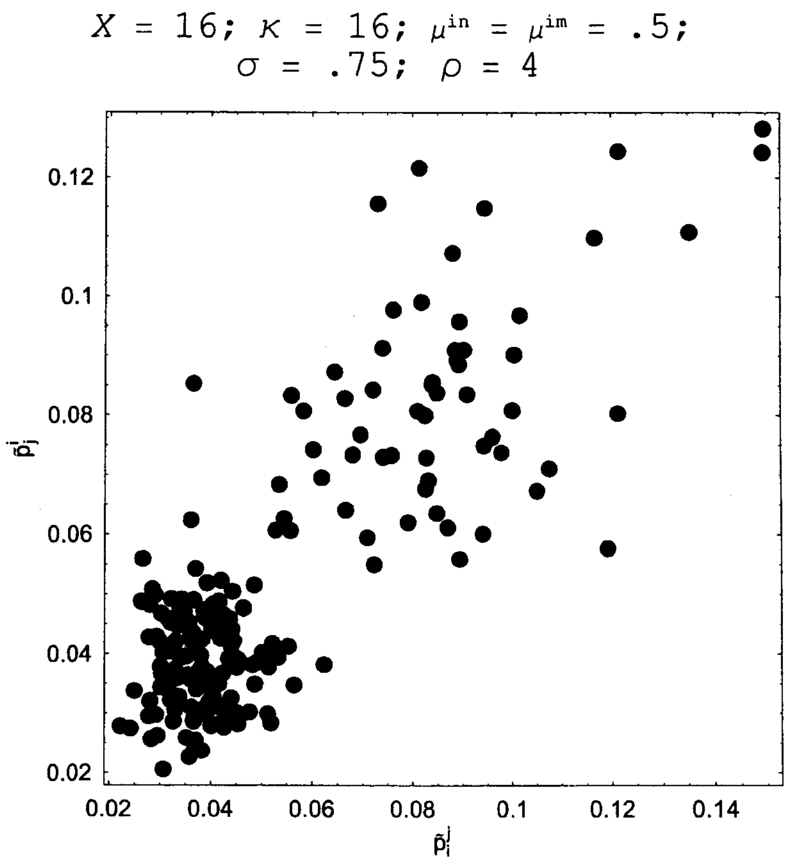

FIG. 4.—($\tilde{p}_i^j$, $\tilde{p}_j^i$) plot

probability agent i observes agent j and the vertical coordinate is the probability agent j observes agent i. First, note that these probabilities range widely from 0.02 to 0.15, though there is a concentration of these points at lower probabilities, around 0.04. More important, the plot clearly indicates that there is a positive relationship between $\tilde{p}_i^j$ and $\tilde{p}_j^i$ such that a relatively high (low) $\tilde{p}_i^j$ is paired with a relatively high (low) $\tilde{p}_j^i$. In other words, social learning tends to be mutual. For a more thorough verification of this property, we computed correlations between $\tilde{p}_i^j$ and $\tilde{p}_j^i$; see table 2. The correlation is positive in all cases and is generally rather high. Furthermore, the correlation appears to increase in X—see table 2(B) for when X is raised to 32—and decrease in κ—see table 2(C) for when κ is lowered to 8. The stronger the mutuality in learning, the greater the intergroup diversity in goals, and the lower the intragroup diversity.

TABLE 2
CORRELATION BETWEEN p_i^j AND p_j^i

μ^{in}	μ^{im}: .25	.5	.75	1
	A. $X = 16$; $\kappa = 16$			
0	.352657	.34733	.29233	.480061
.25	.807451	.773048	.664245	.506862
.5	.586809	.823839	.806357	.689279
.75	.406306	.791755	.799773	.784047
1	.268945	.681329	.860758	.821619
	B. $X = 32$; $\kappa = 16$			
0	.540246	.60035	.736072	.630133
.25	.901541	.862012	.835666	.795593
.5	.786882	.941283	.906933	.878898
.75	.546645	.892148	.93532	.911084
1	.385492	.837522	.930428	.949636
	C. $X = 16$; $\kappa = 8$			
0	.649251	.718233	.741049	.706987
.25	.875383	.851206	.83393	.845132
.5	.870919	.920375	.857695	.837231
.75	.735823	.922655	.907393	.90641
1	.57091	.912438	.905253	.910782

To delve further into network structure, let us impose a bit more structure without losing any generality. Suppose agents are initially allocated to different groups in a nonrandom way so that group 1 contains agents 1 through 5, group 2 contains agents 6 through 10, group 3 contains agents 11 through 15, and group 4 contains agents 16 through 20. The steady-state observation probabilities, $\tilde{p}_i^j$, are described in figure 5, where agent i's (the observer's) identity is represented along the vertical frame and agent j's (the target's) identity is represented along the horizontal frame. A scaled value of $\tilde{p}_i^j$ is then captured by the grey shading of the cell located on ith row and jth column. The lighter (darker) the shading, the higher (lower) the probability with which agent i observes agent j. Since an agent observes herself with zero probability, the diagonal cells are shaded completely black. The most striking feature of the figure is the emergence of four adjoining square (5 × 5) blocks of cells along the diagonal having lighter shades. Each block measures the probabilities of various agents' observing members of the same group. Clearly, not only is learning mutual,

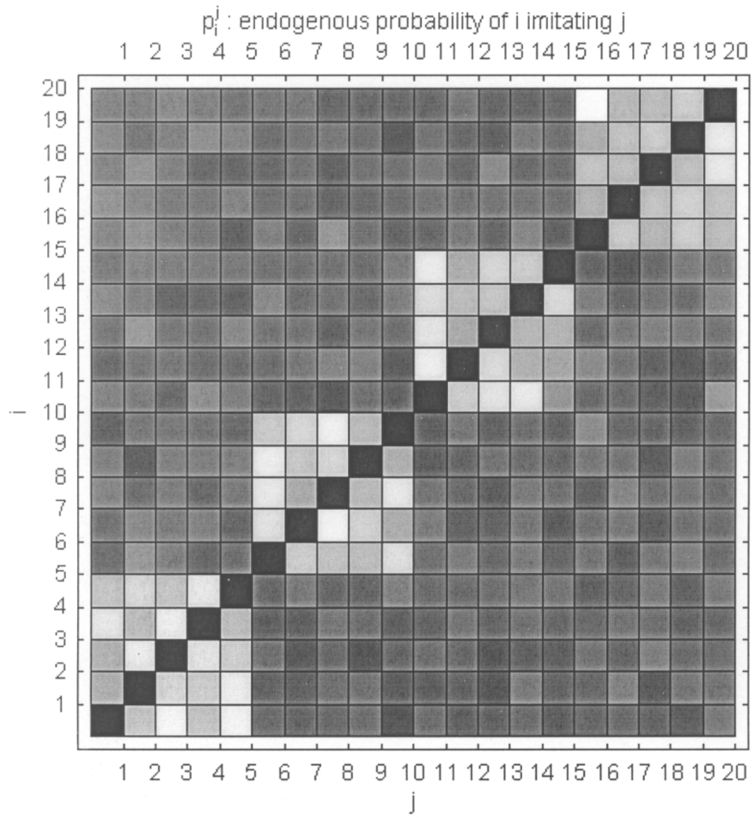

FIG. 5.—p_i^j with nonrandom agent allocation ($X = 16$; $\kappa = 16$; $\mu^{in} = \mu^{im} = .5$; $\sigma = .75$; $\rho = 4$).

it is also more active among agents sharing similar goals.[18] Considered together with the results reported in table 2, it indicates that such intragroup mutual learning is more intensive when the groups are more segregated and isolated from one another.

[18] As noted earlier, Carley (1991) constructs a model of endogenous social groups in which the probability of interaction between agent i and agent j is solely determined by how much information they share. However, in her formulation, the agents are not goal driven and the resulting mutual learning between members of the same group arises from the specification of these interaction probabilities at the social level, rather than from the deliberate choices made at the individual agent level.

COMPARATIVE DYNAMICS

Having explored the emergent structure of networks, the next step is to investigate how performance is influenced by the reliability of the communication technology supporting the network and the productivity of agents in engaging in innovation. Furthermore, we must investigate how this relationship is impacted by features of the environment such as the turbulence in the task environments and interagent goal diversity. So as to extract away the impact of initial conditions as well as to focus on the long-run properties of the system, our analysis will concentrate on the key endogenous variables at the steady state. Figure 3(A) shows that the social system reaches the steady state by $t = 10{,}000$. The speed of convergence to the steady state varies, depending on the parameter values. By taking our measurements as per-period averages over the time periods between 15,000 and 20,000, we are confident that our observations represent the system's behavior on the steady state.[19] To this end, define $\tilde{E}$, $\tilde{\pi}$, and $\tilde{q}$ as the steady-state values of the endogenous variable:

$$\tilde{E} \equiv \left(\frac{1}{5000}\right) \sum_{t=15{,}001}^{20{,}000} \overline{E}(t),\ \tilde{\pi} \equiv \left(\frac{1}{5000}\right) \sum_{t=15{,}001}^{20{,}000} \bar{\pi}(t),\ \tilde{q} \equiv \left(\frac{1}{5000}\right) \sum_{t=15{,}001}^{20{,}000} \bar{q}(t). \tag{19}$$

[19] Note that the performance measure is an average over the 5,000 periods from $t = 15{,}000$ to 20,000 with that being averaged over 20 replications. Two conditions need to be satisfied in order for this to be an appropriate measure for comparative dynamics: (1) the long-run output must not be sensitive to minor changes in initial conditions, and (2) the output must converge to a steady state by $t = 15{,}000$ such that the per-period average from 15,000 to 20,000 approximates the mean of the limiting distribution of the underlying stochastic process. For this purpose, we examined the output from a set of randomly selected replications using different initial conditions. For any given set of parameter values, the output generated from these runs is quite similar, and we found no evidence that the long-run population averages are sensitive to initial conditions. This suggests that the stochastic process generating $\bar{\pi}(t)$ is ergodic; that is, the *transient distribution* of the output in period t converges to a stationary distribution in the limit as $t \to \infty$ for any initial conditions. While the limit distribution is independent of initial conditions, the rate at which the transient distribution converges to it is generally not. A computationally efficient procedure would then call for identifying the minimum time index, $\hat{t}$, such that the transient distributions are approximately the same from that point on. We have examined the time series of $\bar{\pi}(t)$ averaged over 20 replications for the baseline parameter values of $X = 16$, $\kappa = 16$, $\rho = 4$, and $\sigma = .75$ for all $\mu^{in} \in \{.25, .5, .75, 1\}$ and all $\mu^{im} \in \{0, .05, .1, \ldots, .9, .95, 1\}$. The simulations looked very similar to the one captured in figure 3(A). In all cases, the convergence occurred by $t = 15{,}000$. See Law and Kelton (2000) for further discussions on how to identify the steady state.

In the remaining part of this article, the focus is on exploring the impact of μ^{im}, μ^{in}, σ, ρ, X, and κ on $\tilde{E}$, $\tilde{\pi}$, and $\tilde{q}$.[20]

Impact of the Reliability of the Network

Figure 6 plots $\tilde{E}$, $\tilde{q}$, and $\tilde{\pi}$ for various values of $\mu^{im} \in \{0, .05, .1, .15, .2, .25, \ldots, .9, .95, 1\}$, given the baseline parameter values, $\mu^{in} = .5$, $\sigma = .75$, $X = 16$, $\kappa = 16$, and $\rho = 4$. Figure 6(A) shows that when there is minimum reliability, $\mu^{im} = 0$, the mean entropy at the steady state remains at its initial value of 4.25 which corresponds to a fully random network. The mean entropy gradually drops as μ^{im} increases and then is essentially constant once μ^{im} achieves a sufficiently high level (.4 in this case). Further improvements in reliability mildly raise entropy for reasons that are unclear.

Figure 6(B) shows that the endogenous probability of choosing innovation, $\tilde{q}$, is near its maximum value of one when the reliability of the communication technology is at its minimum, $\mu^{im} = 0$. As μ^{im} rises, the rate of innovation correspondingly declines as a more reliable network encourages agents to engage in more imitation as well as forming a more useful network (as reflected in falling entropy). However, the impact of a marginal gain in μ^{im} on $\tilde{q}$ varies, depending on the level of μ^{im}. When μ^{im} is low, $\tilde{q}$ declines at an increasing rate up to $\mu^{im} \simeq 0.35$. Beyond that point, $\tilde{q}$ declines at a decreasing rate.

Most interesting is figure 6(C), where the steady-state performance of the social system, $\tilde{\pi}$, is plotted as a function of the network reliability, μ^{im}. The plot has the S-shape of a logistic curve, where the performance initially rises at an increasing rate until reaching an inflection point, at which the marginal gain is maximal, and then ending up in a region of diminishing returns. Performance achieves its maximum level at a point of maximum reliability of the network.

The logistic curve shape of $\tilde{\pi}$ with respect to μ^{im} is universal to all parameter configurations considered. However, a notable property of figure 6(C) is not universal. Figure 7 plots steady-state performance as a function of μ^{im} for the same parameters as in figure 6 except that μ^{in} is lowered to .25 and κ to 4; agents are less productive at innovation and the degree of intragroup goal congruence is higher. Performance is now declining in network reliability when reliability is sufficiently great (specifically, $\mu^{im} \geq .3$). Quite surprisingly, an improvement in the communi-

[20] It is important to remember that the system never settles down as agents' goals are always stochastically changing, and as a result, agents always changing their behavior and their methods vector. These reported variables—$\tilde{E}, \tilde{\pi}, \tilde{q}$—are then to be interpreted as the means of the stationary distribution for this stochastic process.

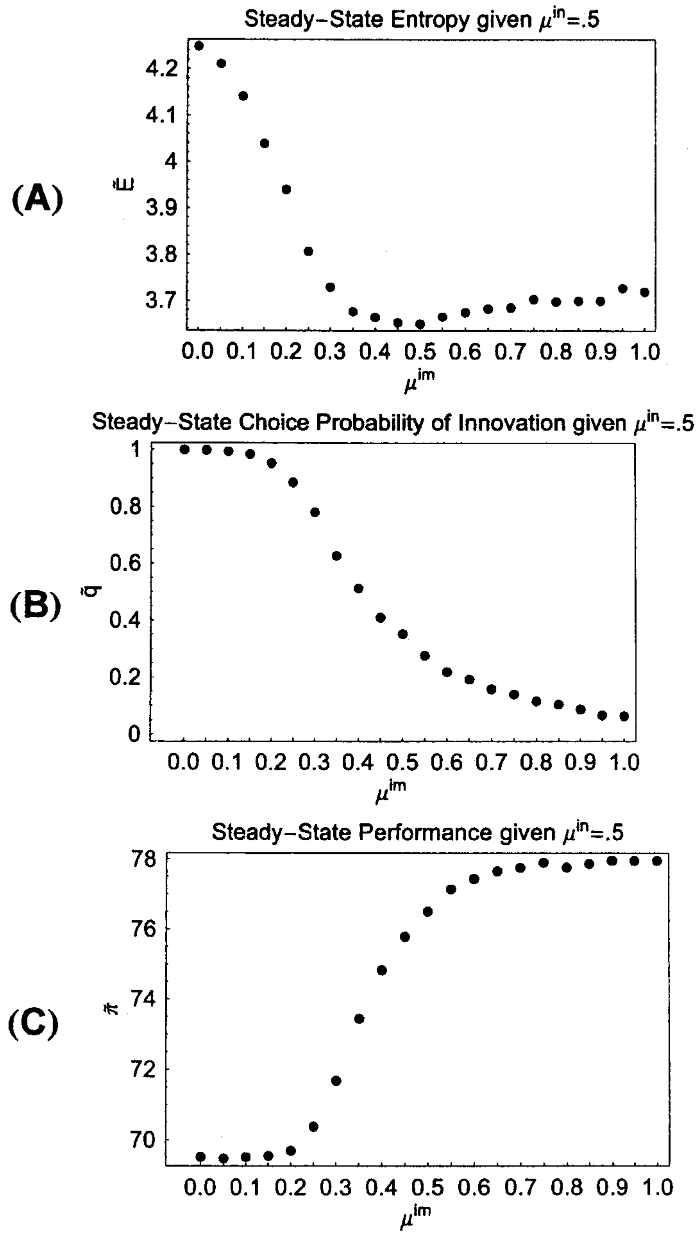

FIG. 6.—Baseline steady states

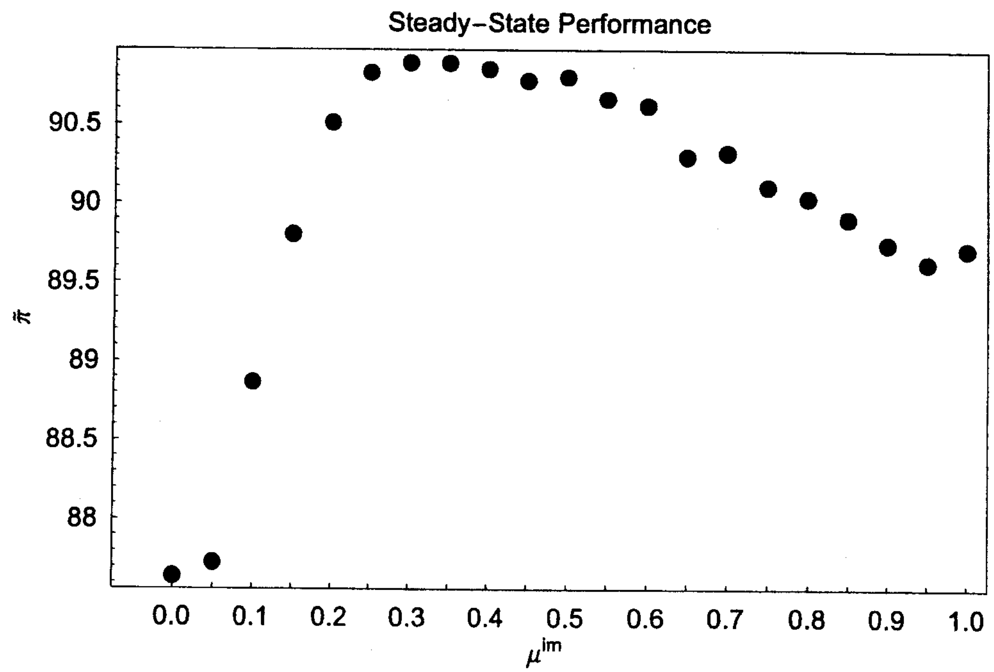

FIG. 7.—Nonmonotonicity of steady-state performance ($X = 16$; $\kappa = 4$; $\mu^{in} = .25$; $\sigma = .75$; $\rho = 4$)

cation technology leads to a deterioration in performance. Further simulations reveal that the nonmonotonicity of $\tilde{\pi}$ in μ^{im} is a general property for some part of the parameter space. In particular, the simulation results reported in figures 8 and 9 show that $\tilde{\pi}$ declines in μ^{im} when σ is low, ρ is high, and/or X is high.

Property 1.—When the reliability of the network (μ^{im}) is sufficiently low, steady-state performance ($\tilde{\pi}$) is increasing in reliability. When the reliability of the network is sufficiently high and the task environment is sufficiently volatile (σ low and/or ρ high) and/or the goal diversity among groups is sufficiently great (X high), performance is decreasing in reliability.

Why does the reliability of the network have a deleterious effect on performance? Recall that there are two search strategies in our model, innovation (individual learning) and imitation (social learning). For a given value of μ^{in}, an increase in μ^{im} raises the rate at which existing useful practices diffuse in the population, which certainly improves the short-run performance of an average agent. However, in the longer run, its impact is more subtle because agent behavior regarding innovation, imitation, and network links is adapting. In particular, a more reliable network means that new ideas are diffused faster, imitation tends to be substituted for innovation, and the greater use of imitation allows a more structured network to form. All of these effects have the implication of making agents in a network more alike as they observe and adopt the practices of others. When the task environment is relatively stable, this lack of diversity is not a problem, and it is more critical that the faster social learning allowed by a higher value of μ^{im} (and the commensurate fall in entropy) speeds up convergence to local optima. But when the task environment is sufficiently volatile, agents have to modify their tasks continually. For that reason, homogenization of the network is seriously harmful as it is less likely that anyone in the network will have useful tasks that would serve the new environment well. While individually, it may make sense to utilize the network more intensively when it is more reliable, the ensuing lack of diversity that occurs—as imitation crowds out innovation—is detrimental from a collective perspective. This finding explains why, under certain conditions, steady-state performance can decline in response to a more reliable network. Better communication does not necessarily mean higher performance when agents can adjust their mix of innovation and imitation and their network links.

As for the role of intergroup diversity (as measured by X), first note that when X is low, the optima of agents from different groups tend to overlap to a greater extent. This overlap leads to social learning that is more global in nature in that an agent tends to observe other agents in different groups more frequently than when optima are tightly clustered

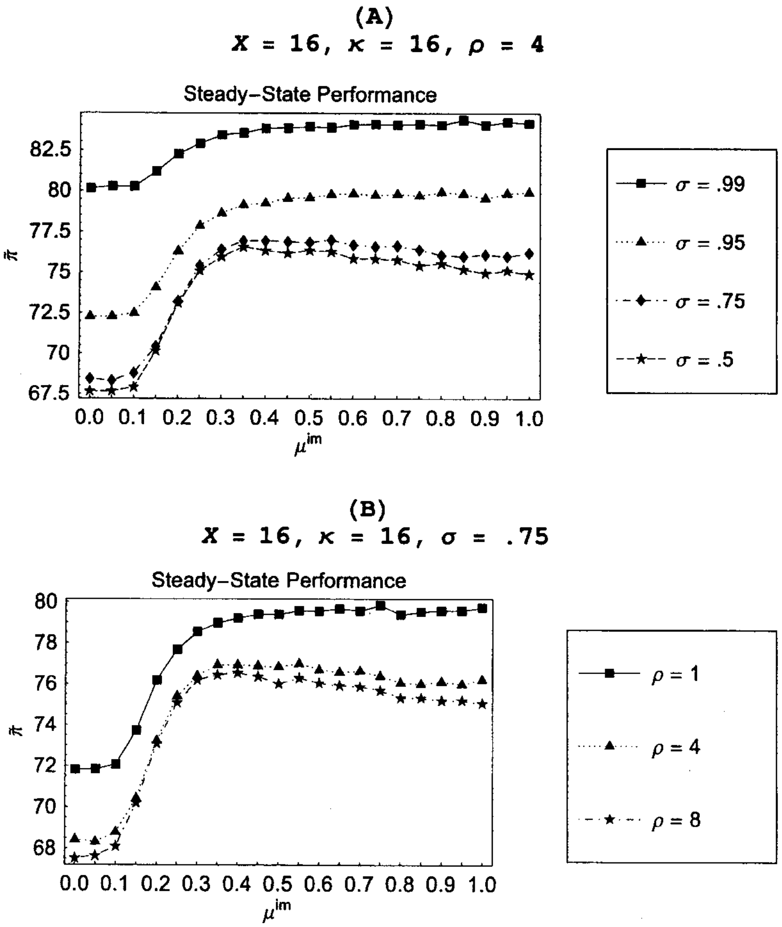

FIG. 8.—Impact of σ and ρ ($X = 16$; $\kappa = 16$)

around distinct groups. The implication of such global learning in our context is that it allows interagent diversity to survive over time instead of leading the social system to a collection of isolated homogeneous clusters of individuals who are unable to adapt flexibly to changing environments because of the lack of diversity.[21] This is portrayed in figures 9(A)–(B).

What we have seen in this section is the critical role that persistent

[21] Note that for a higher value of X, there is more intergroup diversity in the population. However, this diversity is not particularly useful for the agents, as their goals are very different and the likelihood of learning anything useful from the agents in other groups is much lower.

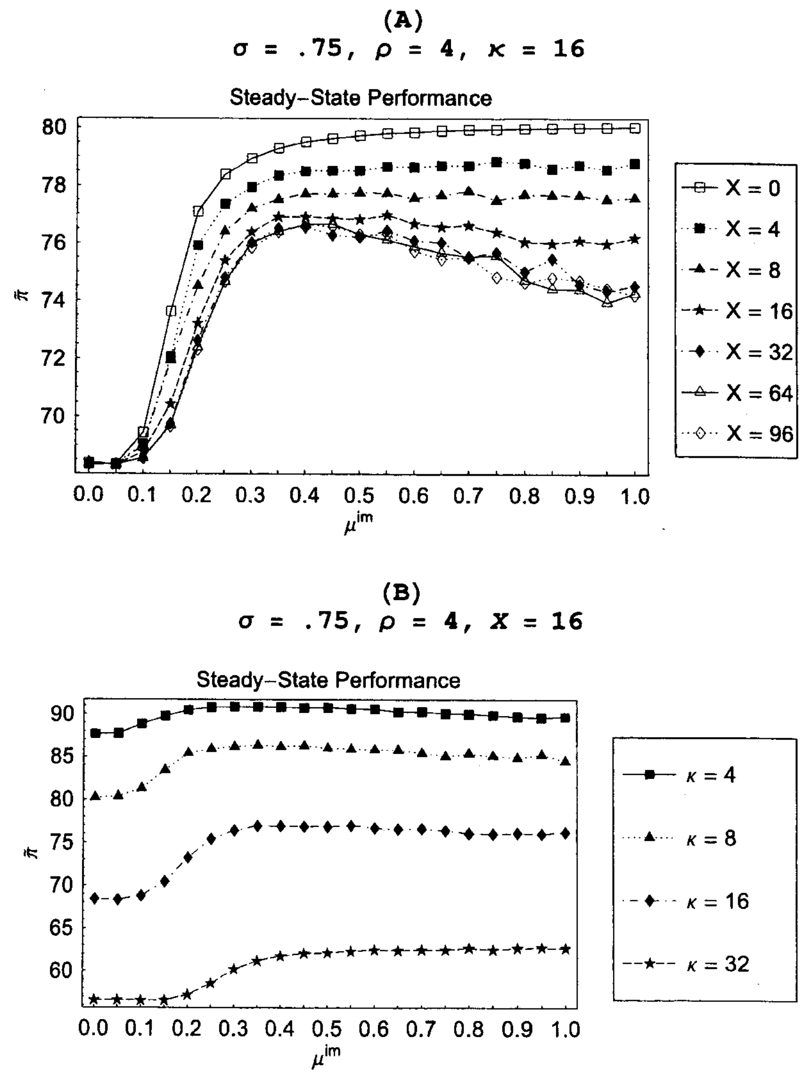

FIG. 9.—Impact of X and κ ($\sigma = .75$; $\rho = 4$)

diversity plays in promoting steady-state performance in the context of a changing environment. A simple improvement in the reliability of the network may harm long-run performance because of how agents adapt their rates of innovation and imitation. Enhanced communication technology can induce too much imitation with a resulting loss of diversity in responding to future environments. In that case, an individual agent's capacity to carry on independent innovation is crucial in supplying the necessary fuel for the effective operation of the social networks. To that issue, we turn next.

Impact of the Productivity of Innovation

Given the relationship between network reliability and performance, the next task is to explore how this relationship depends on the capacity of agents to innovate. Obviously, agents who are more productive at innovation have an immediate virtue through enhancing individual learning. But productivity at innovation is also beneficial from a social perspective as it implies a wider variation in methods in the population which is the raw material for, in the short run, responding to a changing environment. This suggests that more innovative agents are also beneficial from a social learning perspective. However, as the results below will show, it is also necessary that the network be well developed for diffusing ideas, and that depends not only on the exogenously determined reliability, but also on the endogenously determined structure of the network. The extent to which a higher capacity for innovation improves performance depends on how well the network is developed, which depends on how extensively agents use it, which depends on their innovation-imitation mix, which depends on the capacity for innovation; thus, we come full circle. The ensuing results are far from transparent at this point.

Figure 10(A) reports the frequency with which agents choose to innovate on average, $\tilde{q}$, as a function of network reliability, μ^{im}, for various innovation capabilities, $\mu^{in} \in \{0, .25, .5, .75, 1\}$. The impact of μ^{in} on the frequency with which agents try to innovate is monotonic: $\tilde{q}$ rises as μ^{in} is increased. Populations with more innovative members will indeed engage in more innovation, irrespective of network reliability. The steady-state entropy of the social network, $\tilde{E}$, is plotted as a function of μ^{im} in figure 10(B) for varying values of μ^{in}. Note that as μ^{in} goes up, entropy falls more slowly in μ^{im} but tends to converge to a lower steady-state value. Consequently, an increase in μ^{in} causes the network to be less ordered when μ^{im} is low, while it results in a more ordered network when μ^{im} is high.

Property 2.—When network reliability is low, the order of the social network and the capacity for innovation are inversely related. When network reliability is high, the order of the social network and the capacity for innovation are positively related.

In explaining this property, we know from figure 10(A) that an increase in μ^{in} induces the population to lean toward innovation relative to imitation regardless of network reliability. As there is less social learning, reinforcement learning implies the network is less ordered. Examination

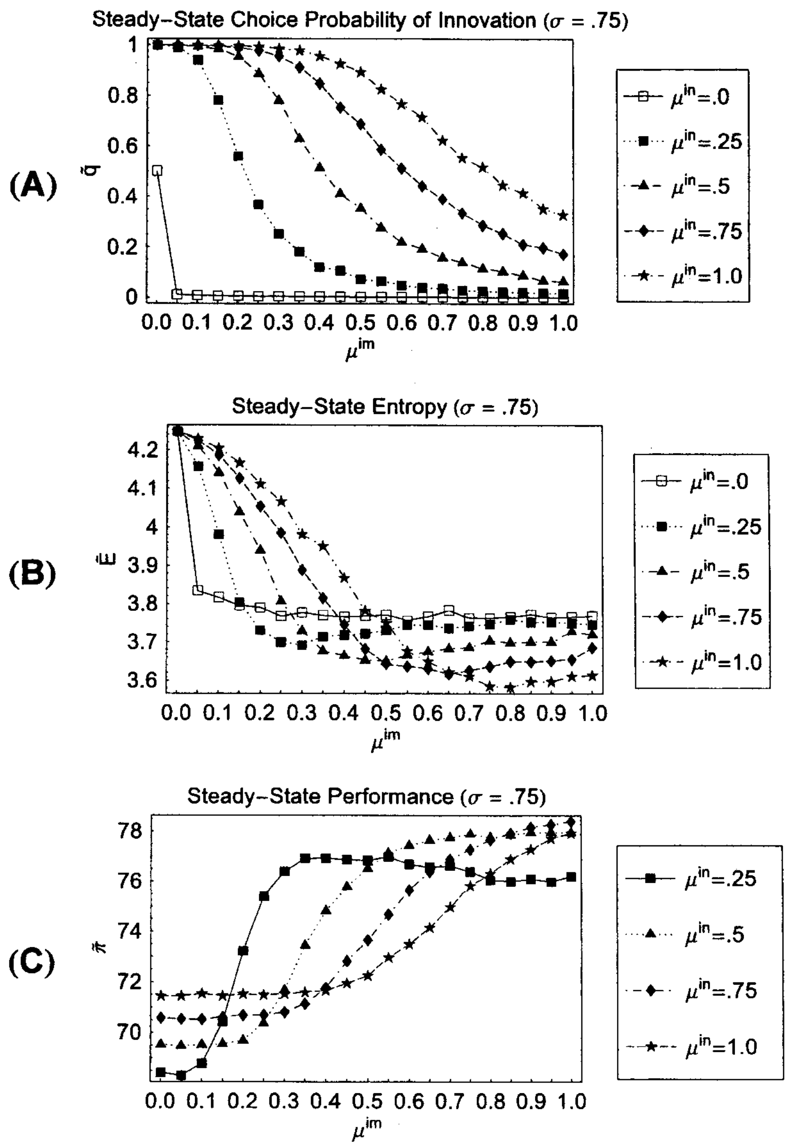

FIG. 10.—$\tilde{q}$, $\tilde{E}$, and $\tilde{\pi}$ ($\sigma = .75$; $X = 16$; $\kappa = 16$; $\rho = 4$)

of figure 10(B) reveals this is the case when network reliability is low as entropy is rising in μ^{in} when $\mu^{im} \leq .35$. However, when network reliability is high, a second force comes into play because a higher value for μ^{in} also raises the value of learning from others since they are more likely to have discovered useful methods. Thus, the quality of the network—in terms of agents' having a diverse array of useful ideas—rises with the productivity of innovation. Innovation is the fuel for improvements in system-wide performance as it creates variations in the population which can be diffused among its members through the social network. But that quality cannot be adequately tapped if reliability is low. Thus, when μ^{im} is low, agents find that accessing the network is unproductive, as it is hard to observe other agents' ideas. Agents then respond to being more innovative by simply substituting away from imitation to innovation. However, when μ^{im} is high, agents find that they can better access the ideas of others, and since agents are highly innovative, there are many worthwhile ideas floating about. Increased quality and reliability of the network then work together to result in a more ordered network even though innovation is more productive. The key point is that the value of imitation and the value of developing links in the network are dependent on the innovativeness of agents. Innovation and imitation are substitutes in search, but through the mechanism of social learning, they are also complements.

Having described how the capacity for innovation impacts network structure, let us next turn to its effect on population performance. Figure 10(C) reports how μ^{in} influences the relationship between steady-state performance, $\tilde{\pi}$, and network reliability, μ^{im}. Note that the $\tilde{\pi}$ curve retains the *S*-shape for all values of μ^{in}. However, the position of the curve for varying values of μ^{in} exhibits an interesting pattern: as μ^{in} rises, the $\tilde{\pi}$ curve shifts up *and* to the right. This pattern leads to the rather surprising result that steady-state performance can decrease with the capacity for innovation and thus, performance can be maximized by having moderately innovative agents.

When μ^{im} is sufficiently low, a population mainly improves on the basis of individual learning through innovation because of the poor reliability of the social network. As innovation is the primary source of improvement, greater innovativeness enhances performance as shown in figure 10(C). On the other hand, for sufficiently high μ^{im}, the population extensively deploys both innovation and imitation. As network reliability achieves its maximum value, there are diminishing returns to network reliability so that the mix of innovation and imitation is fairly insensitive. Once again, higher levels of innovativeness tend to produce higher performance (though the relationship is not quite as strong as when μ^{im} is low). Things get interesting for intermediate values of μ^{im} because the effect of μ^{in} becomes more complex. To begin with, imitation is both a social good and a social detriment. By imitating others rather than developing new ideas, an agent is not adding to the stock of knowledge. However, by helping to spread worthwhile ideas, an agent is improving the value of the network, and thereby, social performance. When network reliability is moderate, agents may be engaging in too much individual learning from a societal perspective. It might be better to tap into the network to pass along ideas and develop more useful links (which can only be achieved through experience). Making agents more innovative induces them to use the network even less, thereby exacerbating this problem. We then find that when network reliability is moderate, greater innovativeness can be deleterious from a societal perspective. By inducing the population to concentrate more on diffusing local innovations across the social system, a reduced capacity for innovation can lead to superior performance.

Figures 11–14 explore how this nonmonotonicity of performance with respect to innovation productivity depends on the characteristics of the environment. The ensuing properties are summarized below.

Property 3.—When the reliability of the network is sufficiently low, performance is increasing in the capacity for innovation. When the reliability of the network is moderate and the task environment is sufficiently volatile (σ low and/or ρ high) and/or the individual goals of the agents are not too dissimilar (X low and/or κ low), performance is maximized at a moderate capacity for innovation.

The volatility in the task environment is captured by $1 - \sigma$, the frequency of the goal change, and ρ, the variability of the goal change. Figure 11 shows performance as a function of μ^{im} for $\mu^{in} \in \{.25, .5, .75, 1\}$ and $\sigma \in \{.75, .95, .99\}$. Imitation and the development of a useful network are more valuable when the environment is changing more frequently, because individual learning is woefully inadequate for keeping up with these changes. One needs the diversity of ideas of the population at large, and as a complement, a well-developed network. It is in such a situation that making agents more innovative will encourage them to engage in more individual learning when the population of agents taken as a whole would do better by engaging in more imitation and the development of better links. We then find that, under a moderately reliable network, a greater capacity for innovation decreases performance when the environment is volatile ($\sigma = .75$) but enhances performance when the environment is stable ($\sigma = .99$). For similar reasons, more innovativeness can be deleterious when ρ is high, because the changes in the environment tend to be bigger on average (holding fixed the frequency with which those changes occur); see figure 12.

The impact of intergroup and intragroup goal diversity on the relationship between innovativeness and performance is reported in figures 13 and 14. Of course, performance decreases both in X and κ since goal

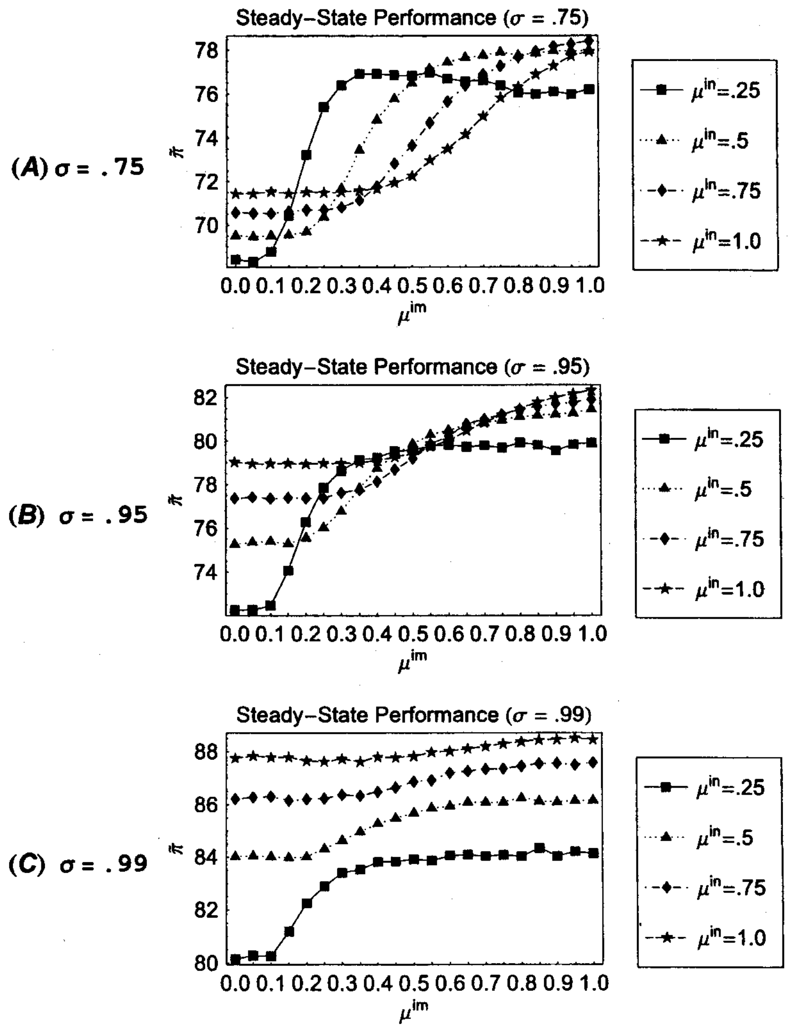

FIG. 11.—Impact of σ ($X = 16$; $\kappa = 16$; $\rho = 4$)

heterogeneity is deleterious (though heterogeneity in tasks is advantageous). Speaking to property 3, the opportunities for social learning and development of a useful network diminish as goals become more diverse both globally and locally. Thus, the deleterious effect of a greater capacity for innovation is less pronounced as agents rely more upon individual learning. When individual learning is the dominant method of search, performance rises monotonically in the ease with which new ideas are generated.

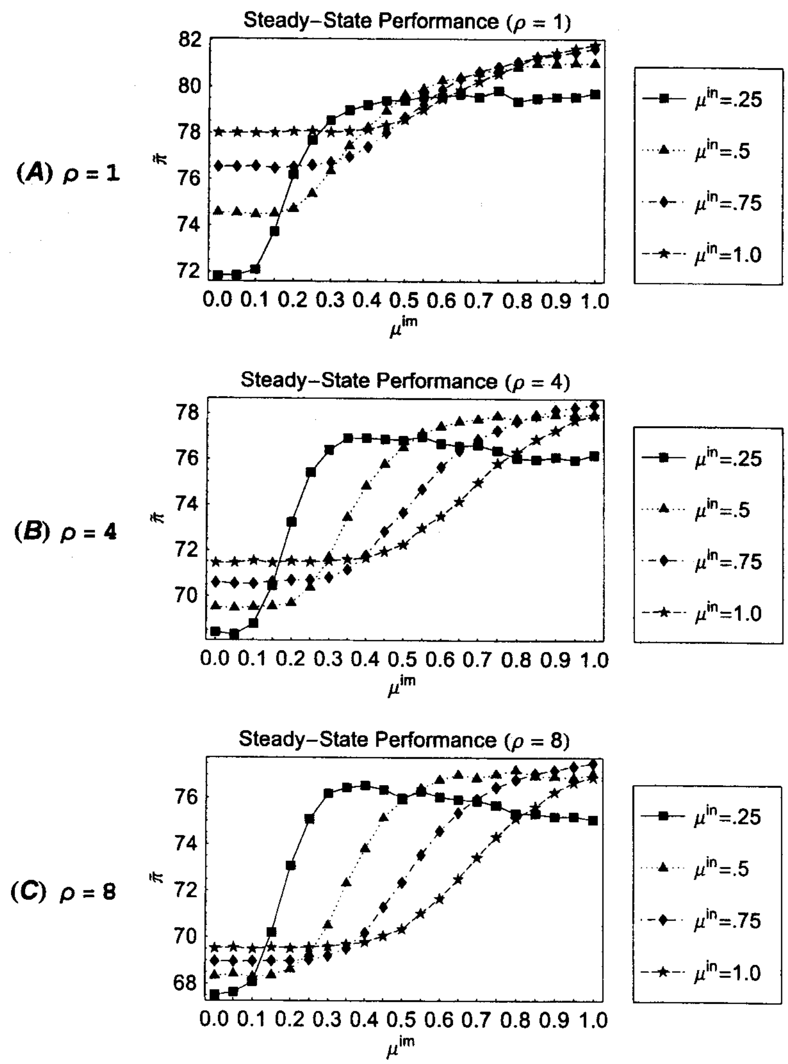

FIG. 12.—Impact of ρ ($\sigma = .75$; $X = 16$; $\kappa = 16$)

CONCLUDING REMARKS

Progress, whether scientific, economic, or social, is driven by innovation—which serves to produce a diversity of ideas—and imitation through a social network—which serves to diffuse these ideas. This study is, to our knowledge, the first to model all three primary forces—the discovery of new ideas, the observation and adoption of the ideas of others, and the

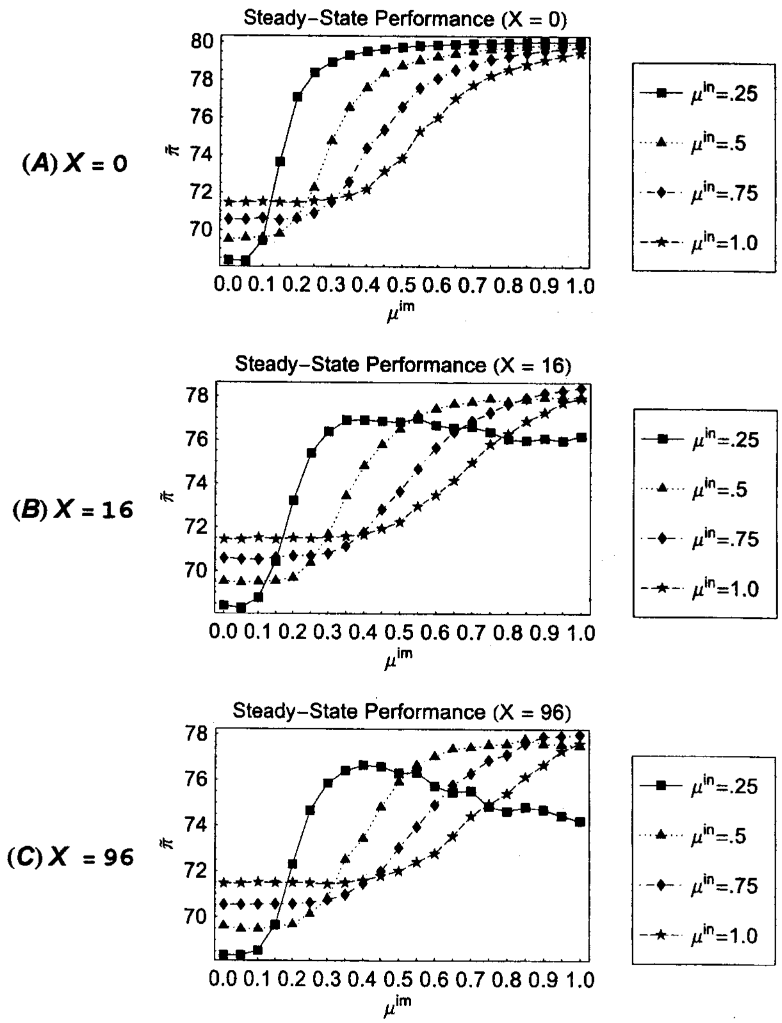

FIG. 13.—Impact of X (σ = .75; κ = 16; ρ = 4)

endogenous development of a social network. At the individual level, innovation and imitation are substitutes, as an agent can choose to allocate effort to discovering new ideas or to observing the ideas of others. However, at the social level, innovation and imitation are complements. More innovation provides a better pool of ideas that imitation can take advantage of through a social network. In our model, the quality of the social network is driven by three factors. The first factor is the exogenous

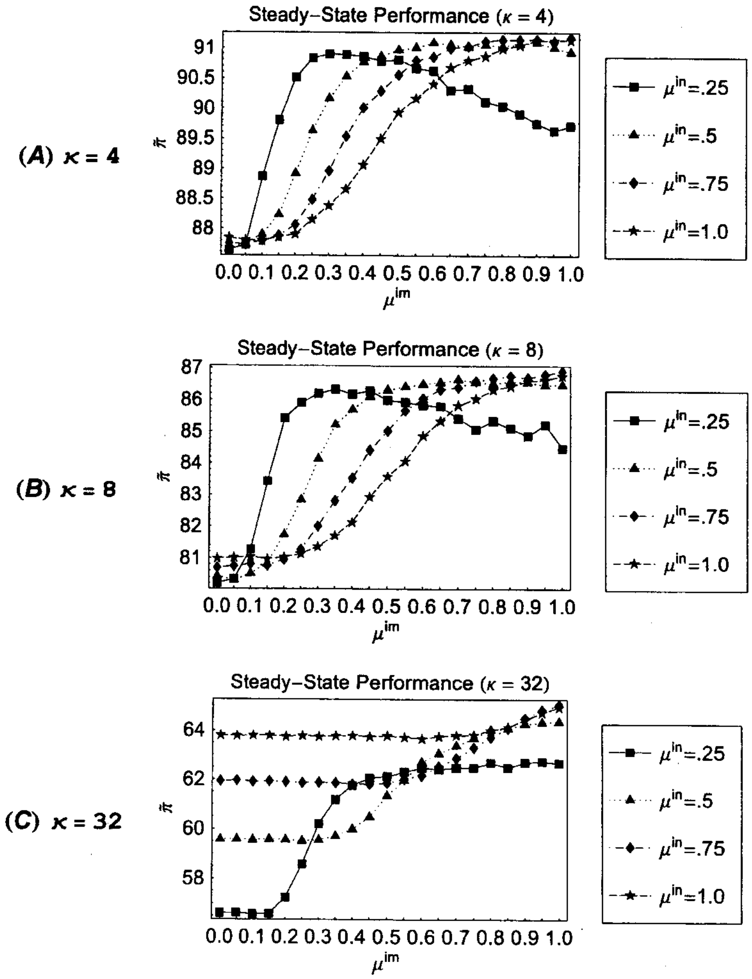

FIG. 14.—Impact of $\kappa\pi$ (σ = .75; X = 16; ρ = 4)

reliability of the communication technology which controls the likelihood that an agent observes another agent's idea. The second factor is the value of the ideas that are observed. This is the result not only of innovation, but also of how rapidly useful ideas diffuse, which depends on the regularity with which the network is accessed. The third factor is the quality of the links between agents, which is itself the product of past experience with the network, as agents reinforce those links that have proven to be a productive source of ideas in the past. These last two factors are not only endogenous, but also provide an increasing returns mechanism in that a network is of higher quality if it is used more frequently, and agents will use a network more frequently when it is of higher quality.

By endogenizing both the innovation-imitation mix and the structure of the social network, our analysis has been able to uncover some subtle determinants of the performance of a population, in particular, the role of the agents' capacity for innovation and the reliability of the communication technology for conveying ideas. While a more reliable communication technology is beneficial in that it leads to a greater diffusion of ideas, it has the undesirable by-product of reducing diversity in ideas as agents replace their own ideas with the most useful ones discovered and engage in more imitation and less innovation. This by-product is not seriously detrimental when the problems faced by agents are relatively stable, but when they are changing sufficiently quickly, then the loss of diversity means a loss of the raw material for adaptation. This results in performance being less when reliability is greater. From a social perspective, agents are engaging in too much imitation, though at an individual level, their innovation-imitation mix is appropriate.

A second subtle result concerns how performance depends on the capacity of agents for innovation. When that capacity is enhanced, agents invest more effort into innovation, which actually may cause performance of the population to decline. When network reliability is only moderately good, it may be socially (though not individually) beneficial to engage in more imitation. Even though this reduces the development of new ideas, it leads to a more developed social network and thus, greater diffusion of ideas. Making agents more innovative induces them to use the network even less, thereby exacerbating the inadequate development of the network. We then find that when agents endogenously determine how much effort to put into innovation and imitation, greater innovativeness can be deleterious from a societal perspective. By inducing the population to concentrate more on diffusing local innovations across the social system, a reduced capacity for innovation can lead to superior performance.

Let us conclude by discussing a few directions for future work. Though the discovery of ideas—through either innovation or imitation—was subject to random forces, it was assumed that the evaluation and implementation of those ideas was flawless. Alternatively, an agent's evaluation of an idea could be based on the true distance between an agent's goals and the new methods vector (after adopting the idea) plus some noise. Along similar lines, one could assume that what actually gets implemented is the true idea plus noise. This assumption is particularly pertinent when it comes to imitation, as one is trying to infer ideas from what people are doing or writing, and such a process is far from flawless. These modifications are not only a move toward greater realism, but also, both of these extensions would serve to introduce additional diversity in the set of methods present in the population. As our analysis showed that a superior communication technology can be detrimental because it reduces diversity, these factors are quite pertinent to the analysis.

On a grander level, a natural extension of our model is to allow agents to be heterogeneous in their capacity for innovating and imitating. Some agents are more creative and thereby more productive in generating new ideas when they choose to engage in the act of discovery. Other agents may be more sociable or more capable of understanding the ideas of others and thus find it easier to learn what other agents are doing. Such heterogeneity raises interesting questions about the properties of the social network. To what extent does this heterogeneity lead to more order? Are links strongest with certain types of agents? Do agents tend to form links with those who are most innovative or those who are most imitative (and thus may be best connected to others in the population)? Finally, in our model, it was assumed that an agent's goal—the characteristics of the problem being solved—was exogenous. However, there may be societal forces that determine what is considered to be a problem worthy of solving. In the context of a hierarchy, problems could be mandated from above, and in less structured contexts, imitation may occur not only with respect to solutions, but also the problems themselves. While it is unclear how the endogeneity of agents' goals should be modeled, it strikes us as important. Enriching our model in these directions should generate further insight into the forces underlying progress and growth.

REFERENCES

Bala, Venktaesh, and Sanjeev Goyal. 2000. "A Noncooperative Model of Network Formation." *Econometrica* 68:1181–1229.

Camerer, Colin, and Teck-Hua Ho. 1999. "Experience-Weighted Attraction Learning in Normal Form Games." *Econometrica* 67:827–74.

Carley, Kathleen. 1990. "Group Stability: A Socio-Cognitive Approach." Pp. 1–44 in *Advances in Group Processes: Theory and Research*, vol. 7, edited by Edward J. Lawler, Barry Markovsky, Cecilia Ridgeway, and Henry A. Walker. Greenwich, Conn.: JAI Press.

———. 1991. "A Theory of Group Stability." *American Sociological Review* 56:331–54.

Eisenstein, Elizabeth L. 1979. *The Printing Press as an Agent of Change: Communications and Cultural Transformations in Early-Modern Europe*. Cambridge: Cambridge University Press.

Huberman, Bernardo A., and Tad Hogg. 1995. "Communities of Practice: Performance and Evolution." *Computational and Mathematical Organization Theory* 1:73–92.

Hummon, Norman P. 2000. "Utility and Dynamic Networks." *Social Networks* 22: 221–49.

Jackson, Matthew O., and Alison Watts. 2002. "The Evolution of Social and Economic Networks." *Journal of Economic Theory* 106:265–95.

Jackson, Matthew O., and Asher Wolinsky. 1996. "A Strategic Model of Social and Economic Networks." *Journal of Economic Theory* 71:44–74.

Kuhn, Thomas S. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Landes, David S. 1998. *The Wealth and Poverty of Nations*. New York: W. W. Norton.

Law, Averill M., and David Kelton. 2000. *Simulation Modeling and Analysis*, 3d ed. New York: McGraw-Hill.

Liberman, Sofia, and Kurt Bernardo Wolf. 1997. "The Flow of Knowledge: Scientific Contacts in Formal Meetings." *Social Networks* 19:271–83.

Mokyr, Joel. 1990. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford: Oxford University Press.

Podolny, Joel M., Toby E. Stuart, and Michael T. Hannan. 1996. "Networks, Knowledge, and Niches: Competition in the Worldwise Semiconductor Industry, 1984–1991." *American Journal of Sociology* 102:659–89.

Shannon, Claude E. 1948. "A Mathematical Theory of Communication." *Bell System Technical Journal* 27:379–423, 623–56.

Sutton, Richard S., and Andrew G. Barto. 2000. *Reinforcement Learning: An Introduction*. Cambridge, Mass.: MIT Press.