

The dynamics of social innovation

H. Peyton Young¹

Department of Economics, University of Oxford, Oxford OX1 3UQ, United Kingdom

Edited by Paul Robert Milgrom, Stanford University, Stanford, CA, and approved June 16, 2011 (received for review January 18, 2011)

Social norms and institutions are mechanisms that facilitate coordination between individuals. A social innovation is a novel mechanism that increases the welfare of the individuals who adopt it compared with the status quo. We model the dynamics of social innovation as a coordination game played on a network. Individuals experiment with a novel strategy that would increase their payoffs provided that it is also adopted by their neighbors. The rate at which a social innovation spreads depends on three factors: the topology of the network and in particular the extent to which agents interact in small local clusters, the payoff gain of the innovation relative to the status quo, and the amount of noise in the best response process. The analysis shows that local clustering greatly enhances the speed with which social innovations spread. It also suggests that the welfare gains from innovation are more likely to occur in large jumps than in a series of small incremental improvements.

convergence time | learning | game dynamics

Social institutions are forms of capital that, together with physical and human capital, determine the rate of economic growth. Notable examples include private property rights and legal systems to protect them, accounting methods, forms of corporate governance, and the terms of economic contracts. However, in contrast with the literature on technological progress, relatively little is known about the ways in which new institutions are created and how they become established within a given social framework. In this paper I discuss one approach to this problem using methods from evolutionary game theory.

In the abstract, an institution can be viewed as a set of rules that structure a given type of interaction between individuals (1). These rules can be simple coordination devices, such as which hand to extend in greeting or who goes through the door first. Or they can be very elaborate, such as rituals of courtship and marriage, cycles of retribution, performance criteria in employment contracts, and litigation procedures in the courts. We want to know how a particular set of rules becomes established as common practice and what process describes the displacement of one set of rules by another.

The viewpoint we adopt here is that new norms and institutions are introduced through a process of experimentation by individuals and that the value of adoption by a given individual is an increasing function of the number of his neighbors who have adopted. Under certain conditions that we discuss below, this trial-and-error process will eventually trigger a general change in expectations and behaviors that establishes the new institution within society at large. However, it can take a very long time for this to happen even when the new way of doing things is superior to the status quo.

One reason for this inertia is lack of information: it may not be immediately evident that an innovation actually is superior to the status quo, due to the small number of prior instances and variability in their outcomes. Thus, it may take a long time for enough information to accumulate before it becomes clear that the innovation is superior (2). A second reason is that an innovation as initially conceived may not work very well in practice; it must be refined over time through a process of learning by doing (3). A third reason is that social innovations often exhibit increasing returns. Indeed this property is characteristic of social

innovations, because they require *coordinated* change in expectations and behaviors by multiple individuals. For example, an individual who invents a new form of legal contract cannot simply institute it on his own: First, the other parties to the contract must enter into it, and second, the ability to enforce the contract will depend on its acceptance in society more generally. Institutions exhibit strongly increasing returns precisely because of their function as coordination technologies. This property is of course not limited to social innovations; it is also a feature of technological innovations with positive network externalities (4–6), and the results in this paper apply to these situations as well.

The increasing returns feature of social innovation has important implications for the dynamics governing institutional change. Of particular relevance is the social network that governs people's interactions. The reason is that, when a social innovation first appears, it will typically gain a foothold in a relatively small subgroup of individuals that are closely linked by geography or social connections. Once the new way of doing things has become firmly established within a local social group, it propagates to the rest of society through the social network. Thus, a key determinant of the speed with which institutional change occurs is the *network topology* and in particular the extent to which interactions are "localized".

The relationship between speed of diffusion and network structure has been investigated in a variety of settings. One branch of the literature is concerned with dynamics in which each agent responds *deterministically* to the number of his neighbors who have adopted the innovation. From an initial seeding among a subset of agents, the question is how long it takes for the innovation to spread as a function of network structure (7–11). A second branch of the literature explores dynamics in which agents respond to their neighbors' choices *probabilistically*; that is, actions with higher expected payoffs are more likely to be chosen than actions with low expected payoffs. Here the question is how long it takes in expectation for the innovation to spread in the population at large (12–15). The rate of spread hinges on topological properties of the network that are different from those that determine the speed of convergence in the nonstochastic case. A third branch of the literature examines the issue of how the network structure itself evolves as agents are matched and rematched to play a given game (16–19).

In this paper we assume that agents adjust their behavior probabilistically and that the network structure is fixed. In contrast to earlier work in this vein, however, we emphasize that it is not only the network topology that determines how fast an innovation spreads. Rather, the speed depends on the interaction between three complementary factors: *i*) the *payoff gain* represented by the innovation in relation to the status quo; *ii*) the *responsiveness* or *rationality* of the agents, that is, the probability

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, "Dynamics of Social, Political, and Economic Institutions," held December 3–4, 2010, at the Arnold and Mabel Beckman Center of the National Academies of Sciences and Engineering in Irvine, California. The complete program and audio files of most presentations are available on the NAS Web site at www.nasonline.org/Sackler_Dynamics.

Author contributions: H.P.Y. performed research and wrote the paper.

The author declares no conflict of interest.

This article is a PNAS Direct Submission.

¹E-mail: peyton.young@economics.ox.ac.uk.

that they choose a (myopic) best response given their information; and *iii*) the presence of *small autonomous enclaves* where the innovation can gain an initial foothold, combined with pathways that allow the innovation to spread by contagion once it has gained a sufficient number of distinct footholds. We also use a somewhat subtle notion of what it means for an innovation to “spread quickly”. Namely, we ask how long it takes in expectation for a high proportion of agents to adopt the innovation and stick with it with high probability. This latter condition is needed because, in a stochastic learning process, it is quite possible for an innovation to spread initially, but then go into reverse and die out. We wish to know how long it takes for an innovation to reach the stage where it is in widespread use and continues to be widely used with high probability thereafter.

The main results of this paper can be summarized as follows. First we distinguish between fast and slow rates of diffusion. Roughly speaking, an innovation *spreads quickly* in a given class of networks if the expected waiting time to reach a high level of penetration and stay at that level with high probability is bounded above independently of the number of agents; otherwise the innovation *spreads slowly*. Whether it spreads quickly or slowly depends on the particular learning rule used, the degree of rationality of the agents, the gain in payoff from the innovation, and the structure of the network. In the interest of keeping the model simple and transparent, we restrict our attention to situations that can be represented by a two-person coordination game with two actions—the status quo and the innovation. Although this assumption is somewhat restrictive, many coordination mechanisms are essentially dyadic, such as principal-agent contracts, the division of property rights between married couples, and rules governing bilateral trade. We also focus mainly on log-linear response processes by individuals, because this class of rules is particularly easy to work with analytically. Among the key findings are the following:

1. If the agents’ level of rationality is too low, the waiting time to spread successfully is very long (in fact it may be infinite) because there is too much noise in the system for a substantial proportion of the agents to *stay coordinated* on the innovation once it has spread initially. However, if the level of rationality is too high, it takes an exponentially long time in expectation for the innovation to gain a foothold anywhere. (Here “rationality” refers to the probability of choosing a myopic best response to the actions of one’s neighbors. Under some circumstances a forward-looking rational agent might try to *induce* a change in his neighbors’ behavior by adopting first, a situation that is investigated in ref. 20.) Hence only for intermediate levels of rationality can one expect the waiting time to be fast in an absolute sense and to be bounded independently of the number of agents. In this respect the analysis differs substantially from that in Montanari and Saberi (15).
2. Certain topological characteristics of the network promote fast learning. In particular, if the agents fall into small clusters or enclaves that are mainly connected with each other as opposed to outsiders, then learning will be fast, assuming that the level of rationality is neither too low nor too high. One reason why clustering may occur in practice is homophily—the tendency of like to associate with like (9, 21). Moreover, recent empirical work suggests that some behaviors actually do spread more quickly in clustered networks than in random ones (22). However, not everyone has to be contained in a small cluster for learning to be fast: It suffices that the innovation be able to gain a foothold in a *subset* of clusters from which it can spread by contagion.
3. For convergence to be fast, it is not necessary for the agents to be contained in enclaves that are small in an absolute sense; it suffices that everyone be contained in a subgroup of

bounded (possibly large) size that has a sufficiently high proportion of its interactions with other members of the group as opposed to outsiders. Various natural networks have this property, including those in which agents are embedded more or less uniformly in a finite Euclidean space and are neighbors if and only if they are within some specified distance of one another. (This result follows from Proposition 2 below; a similar result is proved in ref. 15.)

4. The payoff gain from the innovation relative to the status quo has an important bearing on the absolute speed with which it spreads. We call this the *advance* in welfare represented by the innovation. If the advance is sufficiently large, *no* special topological properties of the network are required for fast learning: It suffices that the maximum degree is bounded.
5. An innovation that leads to a small advance will tend to take much longer to spread than an innovation with a large advance. A consequence is that successful innovations will tend to occur in big bursts, because a major advance will tend to overtake prior innovations that represent only minor advances.

Model

Let Γ be a graph with vertex set V and edge set E , where the edges are assumed to be undirected. Thus, E is a collection of unordered pairs of vertices $\{i, j\}$, where $i \neq j$. Assume that there are n vertices or *nodes*. Each edge $\{i, j\}$ has a weight $w_{ij} = w_{ji} > 0$ that we interpret as a measure of the *mutual influence* that i and j have on one another. For example, w_{ij} may increase the closer that i and j are to each other geographically. Because we can always assume that $w_{ij} = w_{ji} = 0$ whenever $\{i, j\}$ is *not* an edge, the graph Γ is completely specified by a set V consisting of n vertices and a symmetric nonnegative $n \times n$ matrix of weights $W = (w_{ij})$, where $w_{ii} = 0$ for all i .

Assume that each agent has two available choices, A and B . We think of B as the *status quo behavior* and A as the *innovative behavior*. The *state* of the evolutionary process at time t is a vector $x^t \in \{A, B\}^n$, where x_i^t is i ’s choice at time t . Let G be a symmetric two-person game with payoff function $u(x, y)$, which is the payoff to the player who chooses x against an opponent who chooses y . We assume that $u(x, x) > u(y, x)$ and $u(y, y) > u(x, y)$; that is, G is a coordination game. The game G induces an n -person *network game* on Γ as follows: The payoff to individual i in any given period results from playing the game against each of his neighbors once, where i ’s current strategy x_i^t is *unconditional* on which neighbor he plays. The payoff from any given match is weighted according to the influence of that neighbor. For each agent i , let N_i denote the set of i ’s neighbors, that is, the set of all vertices j such that $\{i, j\}$ is an edge. Thus, the payoff to i in state x is

$$U_i(x) = \sum_{j \in N_i} w_{ij} u(x_i, x_j). \quad [1]$$

To give this a concrete interpretation, suppose that B is a form of contractual negotiation that relies solely on a verbal understanding and a handshake, whereas A requires a written agreement and a witness to the parties’ signatures. If one side insists on a written agreement (A) whereas the other side views a verbal understanding as appropriate (B), they will fail to coordinate because they disagree on the basic rules of the game. Note that this metagame (agreeing on the rules of the game) can be viewed as a pure coordination game with zero off-diagonal payoffs: There is no inherent value in adopting A or B unless someone else also adopts it. This situation is shown below, where α is the added benefit from the superior action:

$$\begin{array}{cc|cc} & & A & B \\ \hline A & 1 + \alpha, 1 + \alpha & & \\ B & 0, 0 & & 1, 1 \end{array} \quad [2]$$

[In the case of a symmetric 2×2 coordination game with non-zero off-diagonal payoffs, the adoption dynamics are qualitatively similar, but they tend to select the risk-dominant equilibrium, which may differ from the Pareto-dominant equilibrium (13, 23, 24).]

Learning

How do agents adapt their expectations and behaviors in such an environment? Several models of the learning process have been proposed in the literature; here we discuss two particular examples. Assume that each agent updates at random times according to the realization of a Poisson process with unit expectation. The processes are assumed to be independent among agents. Thus, the probability is zero that two agents update simultaneously, and each agent updates once per unit time period in expectation. We think of each update by some agent as defining a “tick” of the clock (a *period*), and the periods are denoted by $t = 1, 2, 3, \dots$.

Denote the state of the process at the start of period t by x^{t-1} . Let $\epsilon > 0$ be a small error probability. Suppose that agent i is the unique agent who updates at time t . Assume that, with probability $1 - \epsilon$, i chooses an action that maximizes his expected payoff, and with probability ϵ he chooses an action uniformly at random. This process is the *uniform error model* (12, 23, 24).

An alternative approach is the *log-linear model* (13). Given a real number $\beta \geq 0$ and two actions A and B , suppose that agent i chooses A with probability

$$\frac{e^{\beta U_i(A, x_{-i}^{t-1})}}{e^{\beta U_i(A, x_{-i}^{t-1})} + e^{\beta U_i(B, x_{-i}^{t-1})}}. \quad [3]$$

In other words, the log probability of choosing A minus the log probability of choosing B is β times the difference in payoff; hence the term “log-linear learning”. [This model is also standard in the discrete choice literature (25).] The parameter β measures the *responsiveness* or *rationality* of the agent: the larger β is, the more likely it is that he chooses a best reply given the actions of his neighbors, i.e., the less likely it is that he experiments with a suboptimal action.

More generally, we may define a *perturbed best response learning process* as follows. Let G be a 2×2 game with payoff matrix as in eq. 2, and let Γ be a weighted undirected network with n nodes, each occupied by an agent. For simplicity we continue to assume that agents update asynchronously according to i.i.d. Poisson arrival processes. Let $\epsilon > 0$ be a *noise parameter*, and let $q_i^\epsilon(x^{t-1})$ be the probability that agent i chooses action A when i is chosen to update and the state is x^{t-1} . We assume that i) $q_i^\epsilon(x^{t-1})$ depends only on the choices of i 's neighbors; ii) $q_i^0(x^{t-1}) = 1$ if and only if A is a strict best response to the choices of i 's neighbors; iii) $q_i^\epsilon(x^{t-1})$ does not decrease when one or more of i 's neighbors switches to A ; and iv) $q_i^\epsilon(x^{t-1}) \rightarrow q_i^0(x^{t-1})$ in an exponentially smooth manner; that is, there is a real number $r \geq 0$ such that $0 < \lim_{\epsilon \rightarrow 0} q_i^\epsilon(x^{t-1})/\epsilon^r < \infty$. (In the case of log-linear learning these conditions are satisfied if we take $\epsilon = e^{-\beta}$.) This process defines a Markov chain on the finite set of states $\{A, B\}^n$, whose probability transition matrix is denoted by P^ϵ .

Speed of Diffusion

Let \vec{B} be the initial state in which everyone is following the status quo behavior B . How long does it take for the innovative behavior A to become widely adopted? One criterion would be the expected waiting time until the first time that everyone plays A . Unfortunately this definition is not satisfactory due to the sto-

chastic nature of the learning process. To understand the nature of the problem, consider a situation in which β is close to zero, so that the probability of playing A is only slightly larger than the probability of playing B . No matter how long we wait, the probability is high that a sizable proportion of the population will be playing B at any given future time. Thus, the expected waiting time until everyone first plays A is not the relevant concept. A similar difficulty arises for any stochastic learning process: If the noise is too large, it is unlikely that everyone is playing A in any given period.

We are therefore led to the following definition. Given a stochastic learning process P^ϵ on a graph Γ (not necessarily log-linear learning), for each state x let $a(x)$ denote the *proportion* of agents playing A in state x . Given a target level of *penetration* $0 < p < 1$, define

$$T(P^\epsilon, \Gamma, \alpha, p) = E[\min \{t : a(x^t) \geq p \ \& \ \forall t' \geq t, P(a(x^{t'}) \geq p) \geq p\}] \quad [4]$$

In words, $T(P^\epsilon, \Gamma, \alpha, p)$ is the expected waiting time until at least p of the agents are playing A , and the probability is at least p that at least this proportion plays A at all subsequent times. The waiting time depends on the learning process P^ϵ (including the specific level of noise ϵ), as well as on the graph Γ and the size of the advance (the payoff gain) α . Note that the higher the value of p is, the smaller the noise must be for the waiting time to be finite.

To distinguish between fast and slow learning as a function of the number of agents, we consider families of networks of different sizes, where the *size* of a network is the number of nodes (equivalently, the number of agents).

Fast Versus Slow Learning. Given a family of networks \mathcal{G} and a size of advance $\alpha > 0$ (the payoff gain from the innovation), learning is *fast* for \mathcal{G} and α if, for every $p < 1$ there exists $\epsilon_p > 0$ such that for each $\epsilon \in (0, \epsilon_p)$,

$$T(P^\epsilon, \Gamma, \alpha, p) \text{ is bounded above for all } \Gamma \in \mathcal{G}. \quad [5]$$

Otherwise learning is *slow*; that is, for every $\epsilon > 0$ there is an infinite sequence of graphs $\Gamma_1, \Gamma_2, \dots, \Gamma_n, \dots \in \mathcal{G}$ such that $\lim_{n \rightarrow \infty} T(P^\epsilon, \Gamma_n, \alpha, p) = \infty$.

Autonomy

In this section we describe a general condition on families of networks that guarantees fast learning. Fix a network $\Gamma = (V, W)$, a learning process P^ϵ , and an advance $\alpha > 0$. Given a subset of vertices $S \subseteq V$, define the *restricted learning process* \vec{P}_S^ϵ as follows: All nodes $i \notin S$ are *held fixed* at B whereas the nodes in S update according to the process P^ϵ . Let $(\vec{A}_S, \vec{B}_{V-S})$ denote the state in which every member of S plays A and every member of $V - S$ plays B .

The set S is *autonomous* for $(P^\epsilon, \Gamma, \alpha)$ if and only if $(\vec{A}_S, \vec{B}_{V-S})$ is strictly stochastically stable for the restricted process \vec{P}_S^ϵ . [A state is *strictly stochastically stable* if it has probability one in the limit as the noise goes to zero (24, 26, 27).]

Proposition 1. *Given a learning process P^ϵ , a family of networks \mathcal{G} , and an innovation with advance $\alpha > 0$, suppose that there exists a positive integer s such that for every $\Gamma \in \mathcal{G}$, every node of Γ is contained in a subset of size at most s that is autonomous for $(P^\epsilon, \Gamma, \alpha)$. Then learning is fast.*

Concretely this means that given any target level of penetration $p < 1$, there is an upper bound on the noise, $\epsilon_{p,\alpha}$, such that for any given ϵ in the range $0 < \epsilon < \epsilon_{p,\alpha}$, the expected waiting time until at least p of the agents play A (and continue to do so with probability at least p in each subsequent period) is bounded above independently of the number of agents n in the network.

Note that in these examples fast learning does not arise because neighbors of neighbors tend to be neighbors of one another. In fact, a d -dimensional lattice has the property that *none* of the neighbors of a given agent are themselves neighbors. Rather, fast learning arises from a basic fact of Euclidean geometry: The ratio of “surface” to “volume” of a d -dimensional cube goes to zero as the cube becomes arbitrarily large.

A d -dimensional lattice illustrates the concept of autonomy in a very transparent way, but it applies in many other situations as well. Indeed one could argue that many real-world networks are composed of relatively small autonomous groups, because people tend to cluster geographically, or because they tend to interact with people of their own kind (homophily), or for both reasons.

To understand the difference between a network with small autonomous groups and one without, consider the pair of networks in Fig. 2. Fig. 2, *Upper* shows a tree in which every node other than the end nodes has degree 4, and there is a “hub” (not shown) that is connected to all of the end nodes. Fig. 2, *Lower* shows a graph with a similar overall structure in which every node other than the hub has degree 4; however, in this case everyone (except the hub) is contained in a clique of size 4. In both networks all edges are assumed to have weight 1.

Suppose that we begin in the all- B state in both networks, that agents use log-linear learning with $\beta = 1$, and that the size of the advance is $\alpha > 2/3$. Let each network have n vertices. It can be shown that the waiting time to reach at least 99% A (and for this to hold in each subsequent period with probability at least 99%) is unbounded in n for the network on the top, whereas it is bounded independently of n for the network on the bottom. In the latter case, simulations show that it takes <25 periods (on average) for A to penetrate to the 99% level independently of n . The key difference between the two situations is that, in the network with cliques,




Fig. 2. Two networks of degree 4, except for a hub (not shown) that is connected to every end node (dashed lines). All edge weights equal 1.

the innovation can establish a toehold in the cliques relatively quickly, which then causes the hub to switch to the innovation also.

Note, however, that fast learning in the network with cliques does *not* follow from Proposition 2, because not every node is contained in a clique. In particular, the hub is connected to all of the leaves, the number of which grows with the size of the tree, so it is not in an r -close-knit set of bounded size for any given $r < 1/2$. Nevertheless learning is fast: Any given clique adopts A with high probability in bounded time; hence a sizable proportion of the cliques linked to the hub switch to A in bounded time, and then the hub switches to A also. In other words, fast learning occurs through a combination of autonomy and contagion, a topic that we explore in the next section.

Autonomy and Contagion

Contagion expresses the idea that once an innovation has become established for some core group, it spreads throughout the network via the best reply dynamic. Morris (8) was the first to study the properties of contagion in the setting of local interaction games and to formulate graph-theoretic conditions under which contagion causes the innovation to spread throughout the network (for more recent work along these lines see refs. 9–11). Whereas contagion by itself may not guarantee fast learning in a stochastic environment, a combination of autonomy and contagion does suffice. The idea is that autonomy allows the innovation to gain a foothold quickly on certain key subsets of the network, after which contagion completes the job.

Consider a subset S of nodes, all of which are playing A , and choose some $i \notin S$. Let α be the size of the advance of A relative to B . Then A is a strict best response by i provided that

$$(1 + \alpha) \sum_{j \in S} w_{ij} > \sum_{j \notin S} w_{ij}. \quad [13]$$

Letting $r = 1/(\alpha + 2)$, we can write this as follows:

$$\sum_{j \in S} w_{ij} > r \sum_{j \in V} w_{ij} = rd_i. \quad [14]$$

Recall that for any vertex i and subset of vertices S , $d(i, S) = \sum_{j \in S} w_{ij}$ is the total weight on the edges linking i to a member of S . Given a graph $\Gamma = (V, W)$, a real number $r \in (0, 1/2)$, and a subset of vertices S , define the *first r -orbit* of S as follows:

$$O_r^1(S) = S \cup \{i \notin S : d(i, S) > rd_i\}. \quad [15]$$

Similarly, for each integer $k > 1$ recursively define the k th r -orbit by

$$O_r^k(S) = O_r^{k-1} \cup \{i \notin O_r^{k-1} : d(i, O_r^{k-1}) > rd_i\}. \quad [16]$$

Now suppose that $r = 1/(\alpha + 2)$. Suppose also that everyone in the set S is playing A . If the learning process is deterministic best response dynamics, and if everyone updates once per time period, then after k periods everyone in the k th r -orbit of S will be playing A . Of course, this argument does not show that log-linear learning with asynchronous updating will produce the same result. The key difficulty is that the core set S of A -players might unravel before contagion converts the other players to A . However, this problem can be avoided if *i*) the core set S reaches the all- A state within a bounded period and *ii*) S is autonomous and hence its members *continue to play A with high probability*. These conditions are satisfied if the core set is the union of autonomous sets of bounded size. We therefore have the following result.

Proposition 3. *Let \mathcal{G} be a family of graphs. Suppose that there exist positive integers s, k and a real number $r \in (0, 1/2)$ such that, for*

every $\Gamma = (V, W) \in \mathcal{G}$, there is a subset of vertices S such that i) S is the union of r -close-knit sets of size at most s and ii) $O_r^k(S) = V$. Then log-linear learning is fast on \mathcal{G} whenever $\alpha > (1/r) - 2$.

We illustrate with a simple example. Let the network consist of a circle of n agents (the rim) plus a central agent (the hub). Each agent on the rim is adjacent to the hub and to its two immediate neighbors on the rim. Note that the hub is not contained in an r -close-knit set of bounded size for any $r < 1/2$. However, for every $r < 1/2$, the hub is in the first r -orbit of the rim. Moreover, for every $r < 1/3$, there is an r -close-knit set of bounded size that consists of rim nodes; namely, choose any sequence of k adjacent rim nodes where $k > 1/(1 - 3r)$. It follows from Proposition 3 that learning is fast for this family of graphs whenever $\alpha > 1$.

Proposition 3 shows that for some families of networks, log-learning is fast provided that α is large enough, where the meaning of “large enough” depends on the network topology. When the degrees of the vertices are bounded above, however, there is a simple lower bound on α that guarantees fast learning independently of the particular topology.

Proposition 4. Let \mathcal{G} be a family of graphs with no isolated vertices and bounded degree $d \geq 3$. If $\alpha > d - 2$, then log-linear learning is fast.

Proof sketch. Given $\alpha > d - 2$, let k be the smallest integer that is strictly larger than $(\alpha + 2)/(\alpha + 2 - d)$. I claim that every connected set S of size k is autonomous. By Proposition 2 it suffices to show that any such set S is r -close-knit for some $r > 1/(\alpha + 2)$. In particular it suffices to show that

$$\forall S' \subseteq S, S' \neq \emptyset, \frac{d(S', S)}{d(S')} = \frac{d(S', S)}{\sum_{i \in S'} d_i} > 1/(\alpha + 2). \quad [17]$$

Since S is connected, it contains a spanning tree, hence $d(S, S) \geq |S| - 1$. Furthermore, $d \geq d_i$ for all i . It follows that

$$\frac{d(S, S)}{d(S)} \geq \frac{|S| - 1}{d|S|} > 1/(\alpha + 2), \quad [18]$$

where the strict inequality follows from the assumption that $|S| = k > (\alpha + 2)/(\alpha + 2 - d)$. This verifies eq. 17 for the case $S' = S$. It is straightforward to show that eq. 17 also holds for all proper subsets S' of S . We have therefore established that every connected set of size k is autonomous. Clearly every connected component of the network with size less than k is also auto-




Fig. 4. Expected number of periods to reach 99% playing A as a function of the size of advance α when $\beta = 1$, estimated from 10,000 simulations starting from the all- B state.

nomous. Therefore every vertex is contained in an autonomous set of size at most k . It follows that log-linear learning is fast on all networks in \mathcal{G} when $\alpha > d - 2$.

Fast learning says that the waiting time is bounded for a given family of networks, but it does not specify the size of the bound concretely. Actual examples show that the waiting time can be surprisingly short in an absolute sense. Consider an innovation with advance $\alpha = 1$, and suppose that all agents use log-linear learning with $\beta = 1$. Fig. 3 shows the expected waiting time to reach the 99% penetration level for two families of networks: circles where agents are adjacent to their nearest four neighbors and 2D lattices. (Thus, in both cases the networks are regular of degree 4.) The expected waiting time is less than 25 periods in both situations. In other words, almost everyone will be playing A after about 25 revision opportunities per individual.

Note that this waiting time is substantially shorter than it takes for a given individual to switch to A when his neighbors are playing B . Indeed, the probability of such a switch is $e^0/(e^0 + e^{4\beta}) \approx e^{-4\beta} \approx 0.018$. Hence, in expectation, it takes about $1/0.018 = 54$ periods for any given agent to adopt A when none of his neighbors has adopted, yet it takes only about half as much time for nearly everyone to adopt. The reason, of course, is that the process is




Fig. 3. Simulated waiting times to reach 99% A starting from all B for circles (solid curve) and 2D lattices (dashed curve). Time periods represent the average number of updates per individual.

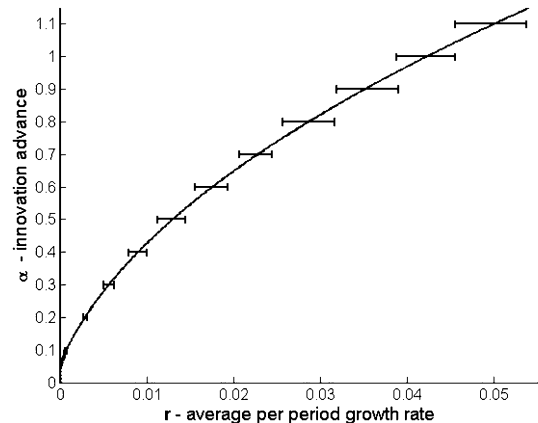


Fig. 5. The level of innovative advance $\alpha = f(r)$ required to achieve an average growth rate of r per period. Simulation results for log-linear learning with $\beta = 1$ on a 30×30 2D lattice are shown.

speeded up by contagion. The rate at which the innovation spreads results from a combination of autonomy and contagion.

Bursts of Innovation

We have seen that the speed with which innovations spread in a social network depends crucially on the interaction between three features: the size of the advance α , the degree of rationality β , and the existence of autonomous groups that allow the innovation to gain a secure foothold. The greater the advance is from the innovation relative to the status quo, the more rapidly it spreads for any given topology, and the more that people are clustered in small autonomous groups the more rapidly the innovation spreads for any given size of advance. Furthermore the degree of rationality must be at an intermediate level for the rate of spread to be reasonably fast: If β is too high it will take a very long time before anyone even tries out the innovation, whereas if β is too low there will be so much random behavior that even approximate convergence will not take place.

In this section we examine how the size of the advance α affects the rate of diffusion for a given family of networks. The results suggest that social innovations tend to occur in large bursts rather than through small incremental improvements. The reason is that a small improvement takes a much longer time to gain an initial foothold than does an innovation that results in a major gain.

To illustrate this point concretely, let \mathcal{G} be the family of 2D lattices where each agent has degree 4. Fig. 4 shows the simulated waiting time to reach the target $p = 0.99$ as a function of α . For small values of α the waiting time grows as a power of α and is many thousands of periods long, whereas for large values of α (e.g., $\alpha > 1$) the waiting time is less than 20 periods.

To understand the implications of this relationship, suppose that each successive innovation leads to an advance of size α . Let $T(\alpha)$ be the waiting time for such an innovation to spread (for a given level of penetration p and level of rationality β). Assume that each new advance starts as soon as the previous one has

reached the target p . Then the *rate of advance* per period is $r(\alpha)$, where $(1 + r(\alpha))^{T(\alpha)} = 1 + \alpha$; that is,

$$r(\alpha) = (1 + \alpha)^{1/T(\alpha)} - 1. \quad [19]$$

The inverse function $f(r)$ mapping r to α is shown in Fig. 5 for a 2D lattice and log-linear learning with $\beta = 1$. Fig. 5 shows that to achieve a 1% average growth rate per period requires innovative bursts of size at least 40%. A 3% growth rate requires innovative bursts of at least 80%, and so forth.

Of course, these numbers depend on the topological properties of the grid and on our assumption that agents update using a log-linear model. Different network structures and different learning rules may yield different results. Nevertheless, this example illustrates a general phenomenon that we conjecture holds across a range of situations. Institutional change that involves a series of small step-by-step advances may take a very long time to spread compared with a change of comparable total magnitude that occurs all at once. The basic reason is that it takes much longer for a small advance to gain a secure foothold in an autonomous group: The group must be quite large and/or it must be quite interconnected to prevent a small advance from unraveling. Furthermore, under a small advance there are fewer pathways through the network that allow contagion to complete the diffusion process. The key point is that social innovations are technologies that facilitate—and require—coordination with others to be successful. It is this feature that makes social change so difficult and that favors large advances when change does occur.

ACKNOWLEDGMENTS. I am indebted to Gabriel Kreindler for suggesting a number of improvements and for running the simulations reported here. The paper also benefited from comments by Tomas Rodriguez Barraquer and participants at the Sackler Colloquium held at the Beckman Institute, December 3–4, 2010. This research was supported by Grant FA 9550-09-1-0538 from the Air Force Office of Scientific Research.

- North D (1990) *Institutions, Institutional Change and Economic Performance* (Cambridge Univ Press, Cambridge, UK).
- Young HP (2009) Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *Am Econ Rev* 99:1899–1924.
- Arrow KJ (1962) The economic implications of learning by doing. *Rev Econ Stud* 29:155–173.
- Katz M, Shapiro C (1985) Network externalities, competition, and compatibility. *Am Econ Rev* 75:424–440.
- David PA (1985) Clio and the economics of QWERTY. *Am Econ Rev Papers Proc* 75:332–337.
- Arthur WB (1994) *Increasing Returns and Path Dependence in the Economy* (Univ of Michigan Press, Ann Arbor, MI).
- Watts D (1999) *Small Worlds: The Dynamics of Networks Between Order and Randomness* (Princeton Univ Press, Princeton).
- Morris S (2000) Contagion. *Rev Econ Stud* 67:57–78.
- Golub B, Jackson MO (2008) How homophily affects the speed of learning and diffusion in networks. Stanford University arXiv:0811.4013v2 (physics.soc-ph).
- Jackson MO, Yariv L (2007) The diffusion of behaviour and equilibrium structure properties on social networks. *Am Econ Rev Papers Proc* 97:92–98.
- Centola D, Macy M (2007) Complex contagions and the weakness of long ties. *Am J Sociol* 113:702–734.
- Ellison G (1993) Learning, local interaction, and coordination. *Econometrica* 61:1046–1071.
- Blume LE (1995) The statistical mechanics of best-response strategy revision. *Games Econ Behav* 11:111–145.
- Young HP (2006) *The Economy as a Complex Evolving System III*, eds Blume LE, Durlauf SN (Oxford Univ Press), pp 267–282.
- Montanari A, Saberi A (2010) The spread of innovations in social networks. *Proc Natl Acad Sci USA* 107:20196–20201.
- Skyrms B, Pemantle R (2000) A dynamic model of social network formation. *Proc Natl Acad Sci USA* 97:9340–9346.
- Jackson MO, Watts A (2002) The evolution of social and economic networks. *J Econ Theory* 106:265–295.
- Goyal S, Vega-Redondo F (2005) Learning, network formation, and coordination. *Games Econ Behav* 50:178–207.
- Jackson MO (2008) *Social and Economic Networks* (Princeton Univ Press, Princeton).
- Ellison G (1997) Learning from personal experience: One rational guy and the justification of myopia. *Games Econ Behav* 19:180–210.
- Lazarsfeld PF, Merton RK (1954) *Freedom and Control in Modern Society*, ed Berger M (Van Nostrand, New York).
- Centola D (2010) The spread of behavior in an online social network experiment. *Science* 329:1194–1197.
- Kandori M, Mailath G, Rob R (1993) Learning, mutation, and long run equilibria in games. *Econometrica* 61:29–56.
- Young HP (1993) The evolution of conventions. *Econometrica* 61:57–84.
- McFadden D (1973) *Frontiers in Econometrics*, ed Zarembka P (Academic, New York), pp 105–142.
- Foster DP, Young HP (1990) Stochastic evolutionary game dynamics. *Theor Popul Biol* 38:219–232.
- Young HP (1998) *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton Univ Press, Princeton).