

The spread of innovations in social networks

Andrea Montanari^a and Amin Saberi^{b,1}

^aDepartment of Electrical Engineering and Department of Statistics, and ^bDepartment of Management Science and Engineering, Stanford University, Stanford, CA 94305

Edited* by Brian Skyrms, University of California, Irvine, CA, and approved August 19, 2010 (received for review March 27, 2010)

Which network structures favor the rapid spread of new ideas, behaviors, or technologies? This question has been studied extensively using epidemic models. Here we consider a complementary point of view and consider scenarios where the individuals' behavior is the result of a strategic choice among competing alternatives. In particular, we study models that are based on the dynamics of coordination games. Classical results in game theory studying this model provide a simple condition for a new action or innovation to become widespread in the network. The present paper characterizes the rate of convergence as a function of the structure of the interaction network. The resulting predictions differ strongly from the ones provided by epidemic models. In particular, it appears that innovation spreads much more slowly on well-connected network structures dominated by long-range links than in low-dimensional ones dominated, for example, by geographic proximity.

Markov chain | convergence times | game dynamics

A great variety of social or technological innovations spread in a population through the network of individual interactions. The dynamics of this process results in the formation of new norms and institutions and has been an important subject of study in sociology and economics (1, 2).

More recently, there has been a surge of interest in this subject because of the rapid growth and popularity of online social interaction. By providing new means for communication and interaction, the Internet has become a unique environment for the emergence and spread of innovations. At the same time, the Internet has changed the structure of the underlying network of interactions by allowing individuals to interact independently of their physical proximity.

Does the structure of online social networks favor the spread of all innovations? What is the impact of the structure of a social network on the spread of innovations? The present paper tries to address these questions.

Epidemic vs. Game-Theoretic Models

In the last few years, a considerable effort has been devoted to the study epidemic or independent cascade models on networks (3). The underlying assumption of these models is that people adopt a new behavior when they come in contact with others who have already adopted it. In other words, innovations spread much like epidemics.

The key predictions of such models are easy to grasp intuitively: (i) Innovations spread quickly in highly connected networks; (ii) long-range links are highly beneficial and facilitate such spreading; (iii) high-degree nodes (hubs) are the gateway for successful spreading.

The focus of this paper is on game-theoretic models that are based on the notion of utility maximization rather than exposure (2). The basic hypothesis here is that, when adopting a new behavior, each individual makes a rational choice to maximize his or her payoff in a coordination game. In these models, players adopt a new behavior when enough of their neighbors in the social network have adopted it; that is, innovations spread because there is an incentive to conform.

In this paper, we will show that the key predictions of game-theoretic models are antithetic to the ones of epidemic models:

(i) Innovations spread quickly in locally connected networks; (ii) geographic (or more general finite-dimensional) structures favor such spreading; (iii) high-degree nodes slow down the diffusion process.

In summary, game-theoretic models lead to strikingly different conclusions from epidemic ones. This difference strongly suggests that assuming that diffusion of viruses, technologies, or new political or social beliefs have the same “viral” behavior may be misleading (4).

It also provides a rigorous evidence that the aggregate behavior of the diffusion is indeed very sensitive to the mechanism of interaction between individuals.[†] This intuition can be quite important when it comes to making predictions about the success of new technologies or behaviors or developing algorithms for spreading or containing them.

Model and Description of the Results

We represent the social network by a graph in which each node represents an agent in the system. Each agent or player has to make a choice between two alternatives. The payoff of each of the two choices for the agent increases with the number of neighbors who are adopting the same choice.

The above model captures situations in which there is an incentive for individuals to make the same choices as their immediate friends or neighbors. This may happen when making a decision between two alternative operating systems (e.g., Windows versus Linux), choosing cell phone providers (AT&T versus Verizon), or even political parties (Republican versus Democratic).

We use a very simple dynamics for the evolution of play. Agents revise their strategies asynchronously. Each time they choose, with probability close to 1, the strategy with the best payoff, given the current behavior of their neighbors. Such noisy best-response dynamics have been studied extensively as a simple model for the emergence of technologies and social norms (2, 5–7). The main result in this line of work can be summarized as follows: The combination of random experimentation (noise) and the myopic attempts of players to increase their utility (best response) drives the system toward a particular equilibrium in which all players take the same action. The analysis also offers a simple condition (known as risk dominance) that determines whether an innovation introduced in the network will eventually become widespread.

The present paper characterizes the rate of convergence for such dynamics in terms of explicit graph-theoretic quantities. Suppose a superior (risk-dominant) technology is introduced as a new alternative. How long does it take until it becomes widespread in the network?

Author contributions: A.M. and A.S. designed research; A.M. and A.S. performed research; and A.M. and A.S. wrote the paper.

The authors declare no conflict of interest.

*This Direct Submission article had a prearranged editor.

[†]As another example, see Liggett (15) for the analysis of the voter models, where again it is observed that greater connectivity tends to slow down the rate at which consensus is reached.

[‡]To whom correspondence should be addressed. E-mail: saberi@stanford.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1004098107/-DCSupplemental.

Our characterization is expressed in terms of quantities that we name tilted cutwidth and tilted cut of the graph. We refer the reader to the following sections for exact definition of these quantities. Roughly speaking, the two quantities are duals of each other: The former characterization is derived by calculating the most likely path to the equilibrium and implies an upper bound on the convergence time; the latter corresponds to a bottleneck in the space of configurations and provides a lower bound.

The proof uses an argument similar to (8–10) to relate hitting time to the spectrum of a suitable transition matrix. The convergence time is then estimated in terms of the most likely path from the worst-case initial configuration. A key contribution of this paper consists in proving that there exists a monotone increasing path with this property. This indicates that the risk-dominant strategy indeed spreads through the network, i.e., an increasing subset of players adopt it over time.

The above results allow us to estimate the convergence time for well-known models for the structure of social networks. This is done by relating the convergence time to the underlying dimension of the graph. For example, if the interaction graph can be embedded on a low-dimensional space, the dynamics converges in a very short time. On the other hand, random graphs or power-law networks are inherently high-dimensional objects, and the convergence may take a time as large as exponential in the number of nodes. As we mentioned earlier, epidemics have the opposite behavior. They spread much more rapidly on random or power-law networks.

Related Work

Kandori et al. (5) studied the noisy best-response dynamics and showed that it converges to an equilibrium in which every agent takes the same strategy. The strategy adopted by all the players was named by Harsanyi and Selten (7) as risk-dominant (see next section for a definition).

The role of graph structure and its interplay with convergence times was first emphasized by Ellison (11). In his pioneering work, Ellison considered two types of structures for the interaction network: a complete graph and a one-dimensional network obtained by placing individuals on a cycle and connecting all pairs of distance smaller than some given constant. Ellison proved that, on the first type of graph structure, convergence to the risk-dominant equilibrium is extremely slow (exponential in the number of players) and for practical purposes, not observable. On the contrary, convergence is relatively fast on linear network and the risk-dominant equilibrium is an important predictive concept in this case. Based on this observation, Ellison concludes that when the interaction is global, the outcome is determined by historic factors. In contrast, when players “interact with small sets of neighbors,” evolutionary forces may determine the outcome.

Even though this result has received a lot of attention in the economic theory [for example, see detailed expositions in books by Fudenberg and Levine (12) and by Young (2)], the conclusion of ref. 11 has remained rather imprecise. The contribution of the current paper is to precisely derive the graph quantity that captures the rate of convergence. Our results make a different prediction on models of social networks that are well-connected but sparse. We also show how to interpret Ellison’s result by defining a geometric embedding of graphs.

Most of our results are based on a reversible Markov chain model for the dynamics. Blume (13) already studied the same model, rederiving the results by Kandori et al. (5). In the last part of this paper, we will also consider generalizations to a broad family of nonreversible dynamics.

There is an extensive literature that motivates games and evolutionary dynamics as appropriate models for the formation of norms and social institutions. For example, see Young (6) for historic evidence about the evolution of the rules of the road

| | |
|---|---|
| a | c |
| d | b |

Fig. 1. Payoff matrix of the coordination game.

as the dynamics of coordination games and to Skyrms (14) and Young (2) on the formation and evolution of social structures.

Finally, we refer to the next two sections for a comparison with related work within mathematical physics and Markov chain Monte Carlo (MCMC) theory.

Definitions

A game is played in periods $t = 1, 2, 3, \dots$ among a set V of players, with $|V| = n$. The players interact on an undirected graph $G = (V, E)$. Each player $i \in V$ has two alternative strategies denoted by $x_i \in \{+1, -1\}$. The payoff matrix A is a 2×2 -matrix illustrated in Fig. 1. Note that the game is symmetric. The payoff of player i is $\sum_{j \in N(i)} A(x_i, x_j)$, where $N(i)$ is the set of neighbors of vertex i .

We assume that the game defined by matrix A is a coordination game, i.e., the players obtain a higher payoff from adopting the same strategy as their opponents. More precisely, we have $a > d$ and $b > c$. Let $N_+(i)$ and $N_-(i)$ be the set of neighbors of i adopting strategy $+1$ and -1 , respectively. The best strategy for a node i is $+1$ if $(a - d)N_+(i) \geq (b - c)N_-(i)$ and it is -1 otherwise. For the convenience of notation, let us define $h = \frac{a-d-b+c}{a-d+b-c}$ and $h_i = h|N(i)|$ where $N(i)$ is the set of neighbors of i . In that case, every node i has a threshold value h_i such that the best strategy given the actions of others or the best-response strategy can be written as $\text{sign}(h_i + \sum_{j \in N(i)} x_j)$.

We assume that $a - b > d - c$, so that $h_i > 0$ for all $i \in V$ with nonzero degree. In other words, when the number of neighbors of node i taking action $+1$ is equal to the number of its neighbors taking action -1 , the best response for i is $+1$. Harsanyi and Selten (7) named $+1$ the “risk-dominant” action because it seems to be the best strategy for a node that does not have any information about its neighbors. Notice that it is possible for h to be larger than 0 even though $b > a$. In other words, the risk-dominant equilibrium is in general distinct from the “payoff-dominant equilibrium,” the equilibrium in which all the players have the maximum possible payoff.

We study noisy best-response dynamics in this environment. In this dynamics, when the players revise their strategy they choose the best response action with probability close to 1. Still, there is a small chance that they choose the alternative strategy with inferior payoff.

More formally, a noisy best-response dynamics is specified by a one-parameter family of Markov chains $P_\beta\{\dots\}$ indexed by β . The parameter $\beta \in R_+$ determines how noisy is the dynamics, with $\beta = \infty$ corresponding to the noise-free or best-response dynamics.

We assume that each node i updates its value at the arrival times of an independent Poisson clock of rate 1. The probability that node i takes action y_i is proportional to $e^{\beta y_i (h_i + \sum_{j \in N(i)} x_j)}$. More precisely, the conditional distribution of the new strategy is

$$P_{i,\beta}(y_i | \Sigma_{N(i)}) = \frac{e^{\beta y_i (h_i + \sum_{j \in N(i)} x_j)}}{e^{\beta (h_i + \sum_{j \in N(i)} x_j)} + e^{-\beta (h_i + \sum_{j \in N(i)} x_j)}} \quad [1]$$

Note that this is equivalent to the best-response dynamics $y_i = \text{sign}(h_i + \sum_{j \in N(i)} x_j)$ for $\beta = \infty$. The above chain is called heat bath or Glauber dynamics for the Ising model. It is also known as

logit update rule, which is the standard model in the discrete choice literature (16).

Let $\underline{x} = \{x_i; i \in V\}$. The corresponding Markov chain is reversible with the stationary distribution $\mu_\beta(\underline{x}) \propto \exp(-\beta H(\underline{x}))$ where

$$H(\underline{x}) = - \sum_{(ij) \in E} x_i x_j - \sum_{i \in V} h_i x_i. \quad [2]$$

For large β , the stationary distribution concentrates around the all-(+1) configuration. In other words, this dynamics predicts that the +1 equilibrium or the Harsanyi–Selten’s risk-dominant equilibrium is the likely outcome of the play in the long run.

The above was observed in by Kandori et al. (5) and Young (6) for a slightly different definition of noisy-best-response dynamics (see sections below and in *SI Appendix*). Their result has been studied and extended as a method for refining Nash equilibria in games. Also, it has been used as a simple model for studying formation of social norms and institutions and diffusion of technologies. See ref. 12 for the former and ref. 2 for an exposition of the latter.

Our aim is to determine whether the convergence to this equilibrium is fast. For example, suppose the behavior or technology corresponding to action -1 is the widespread action in the network. Now the technology or behavior $+1$ is offered as an alternative. Suppose $a > b$ and $c = d = 0$ so the innovation corresponding to $+1$ is clearly superior. The above dynamics predict that the innovation corresponding to action $+1$ will eventually become widespread in the network. We are interested in characterizing the networks on which this innovation spreads in a reasonable time.

To this end, we let T_+ denote the hitting time or convergence time to the all-(+1) configuration and define the typical hitting time for $+1$ as

$$\tau_+(G; \underline{h}) = \sup_{\underline{x}} \inf_{\underline{y}} \{t \geq 0: P_{\beta}^{\underline{x}}\{T_+ \geq t\} \leq e^{-1}\}. \quad [3]$$

For the sake of brevity, we will often refer to this as the hitting time and drop its arguments.

Notice that $\tau_+(G; \underline{h})$ conveys to a very strong notion of convergence: It is the typical time at which all the agents adopt the new strategy. In practice, the behavior of a vanishing fraction of agents is often unobservable. It is therefore important to consider weaker notions as well and ask what is the typical time $\tau_\delta(G; \underline{h})$ such a fraction $(1 - \delta)$ of the agents (say, 90% of them) adopts the new strategy. Most of our results prove to be robust with respect to these modifications.

Relations with Markov Chain Monte Carlo Theory

The reversible Markov chain studied in this paper coincides with the Glauber dynamics for the Ising model and is arguably one of the most studied Markov chains of the same type. Among the few general results, Berger et al. (17) proved an upper bound on the mixing time for $h = 0$ in terms of the cutwidth of the graph. Their proof is based on a simple but elegant canonical path argument. Because of our very motivation, we must consider $h > 0$. It is important to stress that this seemingly innocuous modification leads to a dramatically different behavior. As an example, for a graph with a d -dimensional embedding (see below for definitions and analysis), the mixing time is $\exp\{\Theta(n^{(d-1)/d})\}$ for $h = 0$, while for any $h > 0$ is expected to be polynomial. This difference is not captured by the approach of ref. 17: Adapting the canonical path argument to the case $h > 0$ leads to an upper bound of order $\exp\{\Theta(n^{(d-1)/d})\}$. We will see below that the correct behavior is instead captured by our approach.

There is another important difference with respect to the MCMC literature. As in ref. 5 and subsequent papers in the social science literature, we focus on the small-noise ($\beta \rightarrow \infty$) limit. In

particular, our estimates of the convergence time will take the form $\tau_+(G; \underline{h}) = \exp\{\beta \Gamma_*(G) + o(\beta)\}$. The constant $\Gamma_*(G)$ will then be estimated for large graphs. In particular, even if $\Gamma_*(G)$ is upper bounded by a constant, the $o(\beta)$ term can hide n -dependent factors.

The same point of view (namely studying Glauber dynamics in the $\beta \rightarrow \infty$ limit) has been explored within mathematical physics to understand “metastability.” This line of research has led to sharp estimates of the convergence time (more precisely, of the constant $\Gamma_*(G)$) when the graph is a two- or three-dimensional grid (18–20).

Finally, Ising models with heterogeneous magnetic fields h_i have been intensively studied in statistical physics. However, in that context it is common to assume random h_i s with symmetric distribution. This changes dramatically the model behavior with respect to our model, whereby $h_i \geq 0$ for all i .

Main Results: Specific Graph Families

The main result of this paper is the derivation of graph theoretical quantities that capture the low-noise behavior of the hitting time. In order to build intuition, we will start with some familiar and natural models of social networks:

Random graphs. A random k -regular graph is a graph chosen uniformly at random from the set of graphs in which the degree of every vertex is k . A random graph with a fixed degree sequence is defined similarly. Random graphs with a given degree distribution (and in particular power-law) are standard models for the structure of social networks, World Wide Web, and the Internet.

Another popular class of models for the structure of social networks are preferential-attachment models. In this context, a graph is generated by adding the nodes sequentially. Every new node attaches to d existing vertices, and the probability of attaching to an existing vertex is proportional to its current degree. It has been shown that the degree distribution of graphs in this model converges to power-law.

d -dimensional networks. We say that the graph G is embeddable in d dimensions or is a d -dimensional range- K graph if one can associate to each of its vertices $i \in V$ a position $\xi_i \in \mathbb{R}^d$ such that, (i) $(i, j) \in E$ implies $d_{\text{Eucl}}(\xi_i, \xi_j) \leq K$ (here $d_{\text{Eucl}}(\dots)$ denotes Euclidean distance); (ii) any cube of volume v contains at most $2v$ vertices.

This is a simple model for networks in which social interaction is restricted by some underlying geometry such as physical proximity or closeness of interests.

Small-World Networks. These network models are studied extensively, especially for understanding the small-world phenomenon. The vertices of this graph are those of a d -dimensional grid of side $n^{1/d}$. Two vertices i and j are connected by a “short-range” edge if they are nearest neighbors. Further, each vertex i is connected to k other vertices $j(1), \dots, j(k)$ drawn independently with distribution $P_i(j) = C(n)|i - j|^{-r}$. We call these edges long-range links.

These networks can be seen as hybrid models with short and long-range links. Qualitatively, they are more similar to d -dimensional graphs when r is large. On the other hand, when r is close to zero the long-range edges form a random graph by themselves. The following theorem gives the rate of convergence to the equilibrium for the above network models:

Theorem 1. *As $\beta \rightarrow \infty$, the convergence time is $\tau_+(G) = \exp\{2\beta \Gamma_*(G) + o(\beta)\}$ where*

1. *If G is a random k -regular graph with $k \geq 3$, a random graphs with a fixed degree sequence with minimum degree 3 or a preferential-attachment graph with minimum degree 2, then for h small enough, $\Gamma_*(G) = \Omega(n)$.*

$$\Gamma(G; \underline{h}) \equiv \min_{S: \emptyset \rightarrow V} \max_{t \leq n} [\text{cut}(S_t, V \setminus S_t) - |S_t|_h]. \quad [7]$$

Here the min is taken over all linear orderings of the vertices $i(1), \dots, i(n)$, with $S_t \equiv \{i(1), \dots, i(t)\}$. Note that if for all i , $h_i = 0$, the above is equal to the cutwidth of the graph.

Given a collection of subsets of V , $\Omega \subseteq 2^V$ such that $\emptyset \in \Omega$, $V \notin \Omega$, we let $\partial\Omega$ be the collection of couples $(S, S \cup \{i\})$ such that $S \in \Omega$ and $S \cup \{i\} \notin \Omega$. We then define the tilted cut of G as

$$\Delta(G; \underline{h}) \equiv \max_{\Omega} \min_{(S_1, S_2) \in \partial\Omega} \max_{i=1,2} [\text{cut}(S_i, V \setminus S_i) - |S_i|_h], \quad [8]$$

the maximum being taken over monotone sets Ω (i.e., such that $S \in \Omega$ implies $S' \in \Omega$ for all $S' \subseteq S$).

Theorem 3. Given an induced subgraph $F \subseteq G$, let \underline{h}^F be defined by $h_i^F = h_i + |N(i)|_{G \setminus F}$, where $|N(i)|_{G \setminus F}$ is the degree of i in $G \setminus F$. For reversible asynchronous dynamics we have $\tau_+(G; \underline{h}) = \exp\{2\beta\Gamma_*(G; \underline{h}) + o(\beta)\}$, where

$$\Gamma_*(G; \underline{h}) = \max_{F \subseteq G} \Gamma(F; \underline{h}^F) = \max_{F \subseteq G} \Delta(F; \underline{h}^F). \quad [9]$$

Note that tilted cutwidth and tilted cut are dual quantities. The former corresponds the maximum increase in the potential function $H(\underline{x})$ along the lowest path to the $+1$ equilibrium. The latter is the lowest value of potential function along the highest separating set in the space of configurations. The above theorem shows that tilted cut and cutwidth coincide for the “slowest” subgraph of G provided that the h_i s are nonnegative. This identity is highly nontrivial: For instance, the two expressions in Eq. 9 do not coincide for all subgraphs F . The hitting time is exponential in this graph parameter.

Monotonicity of the Optimal Path The linear ordering in Eq. 7 corresponds to an evolution path leading to the risk-dominant (all $+1$) equilibrium from a worst-case starting point. Characterizing the optimal path provides insight on the typical process by which the network converges to the $+1$ equilibrium (20). For instance, if all optimal paths include a certain configuration \underline{x} , then the network will pass through the state \underline{x} on its way to the new equilibrium, with probability converging to 1 as $\beta \rightarrow \infty$.

Remarkably in Eq. 7 it is sufficient to optimize over linear orderings instead of generic paths in $\{+1, -1\}^V$. This is suggestive of the fact that the convergence to the risk-dominant equilibrium is realized by a monotone process: The new $+1$ strategy effectively spreads as a new behavior is expected to spread. A similar phenomenon was indeed proved in the case of two- and three-dimensional grids (18, 24). Here we provide rigorous evidence that it is indeed generic.

Figs. 4 and 5 illustrate the spread of the risk-dominant strategy in two-dimensional small-world networks. The evolution of the

number of players adopting the new strategy is clearly monotone (although nonmonotonicities on small time scales are observable due to the modest value of β). Also, the qualitative behavior changes significantly with the density of long-range connections.

Nonreversible and Synchronous Dynamics

So far, we have been focusing on asynchronous dynamics, and the noise was introduced in such a way that the resulting chain was reversible. Our next step is to demonstrate that our observations are applicable to a wide range of noisy best-response dynamics.

We consider a general class of Markovian dynamics over $\underline{x} \in \{+1, -1\}^V$. An element in this class is specified by $p_{i,\beta}(y_i | \underline{x}_{N(i)})$, with $p_{i,\beta}(+1 | \underline{x}_{N(i)})$ a nondecreasing function of the number $\sum_{j \in N(i)} x_j$. Further we assume that $p_i(+1 | \underline{x}_{N(i)}) \leq e^{-2\beta}$ when $h_i + \sum_{j \in N(i)} x_j < 0$. Note that the synchronous Markov chain studied in Kandori et al. (5) and Ellison (11) is a special case in this class.

Denote the hitting time of all $(+1)$ -configuration in graph G with $\tau_+(G; \underline{h})$ as before.

Proposition 4. Let $G = (V, E)$ be a k -regular graph of size n such that for $\lambda, \delta > 0$, every $S \subset V$, $|S| \leq \delta n$ has vertex expansion at least λ . Then for any noisy-best-response dynamics defined above, there exists a constant $c = c(\lambda, \delta, k)$ such that $\tau_+(G; \underline{h}) \geq \exp\{\beta cn\}$ as long as $\lambda > (3k/4) + (\max_i h_i/2)$.

Note that random regular graphs satisfy the condition of the above proposition as long as h_i s are small enough. This proposition can be proved by considering the evolution of one-dimensional chain tracking the number of $+1$ vertices.

Proposition 5. Let G be a d -dimensional grid of size n and constant $d \geq 1$. For any synchronous or asynchronous noisy-best-response dynamics defined above, there exists constant c such that $\tau_+(G; \underline{h}) \leq \exp\{\beta c\}$.

The above proposition can be proved by a simple coupling argument very similar to that of Young (25).

Together, these two propositions show that for a large class of noisy-best-response dynamics including the one considered in (11), the isoperimetric constant of the network and more precisely its tilted cutwidth capture the rate of convergence.

Concluding Remarks

One of the weaknesses of the Nash equilibrium for predicting the outcome of a play is that even very simple games often have several equilibria. The coordination game studied in this paper is a very good example: Indeed, it is not hard to construct graphs with an exponential number of equilibria.

One of the standard ways for dealing with this problem is to use stability with respect to noise as a way to determine whether one equilibrium is more likely than another (7). In our opinion, the




Fig. 4. Diffusion of the risk-dominant strategy in a small-world network with $d = 2$, $k = 3$, $n = 50^2$ and mostly short-range connections, namely $r = 5$. We use $\beta = 0.75$ and $h = 0.5$. (Upper) Evolution of the number of nodes adopting the new strategy (in the inset, same curve with time axis starting at $t = 0$). (Lower) Configurations at times such that the number of nodes adopting the risk-dominant strategy is, from left to right, 125, 375, 625, 875, 1125, 1375, 1625, 1875 (indicated by squares in Upper).




Fig. 5. As in Fig. 4 but for a small-world network with a larger number of long-range connections, namely $r = 2$.

techniques developed for analyzing Markov chains in the last two decades are quite potent for advancing this area and achieving a deeper understanding of game dynamics and equilibrium selection.

The current paper focused on a particular game but presented a robust characterization of the rate and the likely path of convergence as a function of network structure. We see this as a first step in developing algorithms for making predictions of play for more generic games. Ideally, instead of singling out a particular equilibrium, such predictions would be a stochastic function of evolutionary dynamics, historic factors, and the length of the play.

As a final note, complete proofs of all the theorems and lemmas presented in this paper and a more extensive comparison with results in the economics literature are available in [SI Appendix](#). A conference version of this paper was presented at FOCS 2009.

ACKNOWLEDGMENTS. We would like to thank Daron Acemoglu, Glenn Ellison, Fabio Martinelli, Roberto Schonmann, Eva Tardos, Maria Eulalia Vares, and Peyton Young for helpful discussions and pointers to the literature. The authors acknowledge the support of the National Science Foundation under grants CCF-0743978, DMS-0806211, and CCF-0915145.

- Schelling T (1978) *Micromotives and Macrobehavior* (W.W. Norton, New York).
- Young HP (2001) *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton Univ Press, Cambridge, MA).
- Kleinberg J (2007) Cascading behavior in networks: Algorithmic and economic issues. *Algorithmic Game Theory*, eds N Nisan et al. (Cambridge Univ Press, Cambridge, MA).
- Watts D (2007) Challenging the Influentials Hypothesis. *Measuring Word of Mouth*, 3 pp:201–211.
- Kandori M, Mailath H, Rob F (1993) Learning, mutation, and long run equilibria in games. *Econometrica* 61:29–56.
- Young HP (1993) The evolution of conventions. *Econometrica* 61:57–84.
- Harsanyi J, Selten R (1998) *A General Theory of Equilibrium Selection in Games* (MIT Press, Cambridge, MA).
- Donsker M, Varadhan S (1976) On the principal eigenvalue of second-order elliptic differential operators. *Commun Pur Appl Math* 29:595–621.
- Diaconis P, Saloff-Coste L (1993) Comparison theorems for reversible Markov chains. *Ann Appl Probab* 3:696–730.
- Sinclair A, Jerrum M (1989) Approximate counting, uniform generation and rapidly mixing Markov chains. *Inf Comput (IANDC)* 82:93–133.
- Ellison G (1993) Learning, local interaction, and coordination. *Econometrica* 61:1047–1071.
- Fudenberg D, Levine D (1998) *The Theory of Learning in Games* (MIT Press, Cambridge, MA).
- Blume L (1995) The statistical mechanics of best-response strategy revision. *Game Econ Behav* 11:111–145.
- Skyrms B (2003) *The Stag Hunt and the Evolution of Social Structure* (Cambridge Univ Press, Cambridge, UK).
- Liggett T (1999) *Stochastic Interacting Systems: Contact, Voter, and Exclusion Processes* (Springer, New York, NY).
- McKelvey P, Palfrey T (1995) Quantal response equilibria for normal form games. *Game Econ Behav* 10:6–38.
- Berger N, Kenyon C, Mossel E, Peres Y (2005) Glauber dynamics on trees and hyperbolic graphs. *Probab Theory Rel* 131:311–340.
- Neves E, Schonman R (1991) Critical droplets and metastability for a Glauber dynamics at very low temperatures. *Commun Math Phys* 137:209–230.
- Bovier A, Manzo F (2002) Metastability in Glauber dynamics in the low-temperature limit: Beyond exponential asymptotics. *J Stat Phys* 107:757–779.
- Olivieri E, Vares ME (2004) *Large Deviations and Metastability* (Cambridge Univ Press, Cambridge, UK).
- Mihail M, Papadimitriou C, Saberi A (2006) On certain connectivity properties of the Internet topology. *J Comput Syst Sci* 72:239–251.
- Gkantsidis C, Mihail M, Saberi A (2003) Throughput and congestion in power-law graphs. *Proc of SIGMETRICS* 148–159.
- Flaxman A (2006) Expansion and lack thereof in randomly perturbed graphs. *Proc of Web Alg Workshop* 4936:24–35.
- Neves EJ, Schonman RH (1992) Behavior of droplets for a class of Glauber dynamics at very low temperatures. *Probab Theory Rel* 91:331–354.
- Young HP (2006) The diffusion of innovation in social networks. *The Economy as an Evolving Complex System, III: Current Perspectives and Future Directions* (Oxford Univ Press, New York).