

Propagation of Innovations in Networked Groups

Winter A. Mason
Yahoo! Research

Andy Jones
University of North Carolina

Robert L. Goldstone
Indiana University Bloomington

A novel paradigm was developed to study the behavior of groups of networked people searching a problem space. The authors examined how different network structures affect the propagation of information in laboratory-created groups. Participants made numerical guesses and received scores that were also made available to their neighbors in the network. The networks were compared on speed of discovery and convergence on the optimal solution. One experiment showed that individuals within a group tend to converge on similar solutions even when there is an equally valid alternative solution. Two additional studies demonstrated that the optimal network structure depends on the problem space being explored, with networks that incorporate spatially based cliques having an advantage for problems that benefit from broad exploration, and networks with greater long-range connectivity having an advantage for problems requiring less exploration.

Keywords: innovation, social network, groups, diffusion

Human problem solving can often be seen as a search for a solution within a large problem space. For important problems, many individuals may be simultaneously trying to find the solution to a problem and often are not working alone. Good, novel solutions to problems—innovations—spread in waves of adoption powered by imitation and influence. This means that every individual must choose between adopting other persons' solutions and continuing his or her search for a better solution. The research described in this article suggests that the pattern of communication within a group—the channels an innovation can spread through—can promote either exploitation of good solutions or exploration for novel ones, and that this subsequently affects the efficiency of the group taken as a whole in finding the best solution to a problem. Our interest is in developing laboratory-controlled methods for exploring the propagation of innovations in social groups and for studying how this propagation is affected by the social network structure and nature of the problem to be solved.

A large portion of the work on innovations has focused on a population-level view of product innovation rather than the processes at the individual level (Bass, 1969; Tarde, 1903). At the individual level there are three factors in the process of spreading innovations: the properties of the innovation, the characteristics of the innovator/adopter, and the social and environmental context (Rogers, 1962, 1995). Of particular interest to psychologists are the decision making that goes into choosing to adopt an innovation and the psychological features of the individual that affect that decision. For instance, innovativeness has been treated as a personality variable for some time (Rogers & Shoemaker, 1971), although its validity has been questioned (Midgley & Dowling, 1978). The central social psychological components of innovation diffusion are imitation and influence.

Imitation is so central to the dissemination of innovations that the first seminal article on the spread of innovations appeared in a book titled *The Laws of Imitation* (Tarde, 1903). This capacity for imitation was termed *no-trial learning* by Bandura (1965), who stressed that, by imitating one another, people perform behaviors that they would not have otherwise considered. It has been suggested that cultural identity exists only because of our tendency to imitate results in the dissemination of concepts, beliefs, and innovations (Bikhchandani, Hirshleifer, & Welsh, 1992). The adoption of an innovation can be viewed as the adopter choosing to imitate the innovator, or it can be viewed as the innovator influencing the adopter. Influences have been classified into two types (Deutsch & Gerard, 1955): normative influences, when people are influenced because they desire to obtain social approval from others, and informational influences, when people are influenced because they believe others possess additional or more accurate information. Both types can drive the propagation of innovations, but they have different effects.

Winter A. Mason, Yahoo! Research; Andy Jones, Department of Computer Science, University of North Carolina; Robert L. Goldstone, Psychological and Brain Sciences and Cognitive Science, Indiana University Bloomington.

This research was funded by the Department of Education's Institute of Education Sciences Grant R305H050116 and National Science Foundation Research Evaluation and Communication Grant 0527920 to Andy Jones. We thank Jason Dawson, who helped run the experiments, and Katy Borner, Todd Gureckis, Peter Todd, Alessandro Vespignani, and Stanley Wasserman for helpful suggestions. More information about the laboratory and online versions of experiments related to the current work can be found at <http://cognitrn.psych.indiana.edu>

Correspondence concerning this article should be addressed to Winter A. Mason, Yahoo! Research, 111 West 40th Street, New York, NY 10018. E-mail: winteram@yahoo-inc.com

Laughlin and Ellis (1986) describe a continuum between intellectual and judgmental issues. Where an issue falls on this continuum depends on the demonstrability of the correctness of solutions. A math problem has accepted means for showing the correctness of a solution, so it would be considered an intellectual task. Deciding which flavor of ice cream tastes better would be a judgmental task. Kaplan and Miller (1987) found that intellectual tasks led to more informational influence, while judgmental tasks led to more normative influence. Developing an innovation is often an intellectual task, because the superiority of the innovation is often demonstrable, and so the spread of the innovation is primarily due to the informational influence of others. However, if there is no discernable difference, or if the benefits of adopting an innovation are largely due to others using it (e.g., IBM vs. Macintosh or VHS vs. BetaMax), choosing to adopt the innovation is more like a judgmental task.

For intellectual issues, if there is a problem with one solution that is clearly and demonstrably better than all other alternatives, the search for that solution essentially becomes an optimization problem. Kennedy, Eberhart, and Shi (2001) developed a model of group search for optimization problems. In this model, a group of agents searches a multidimensional space with a mathematically defined fitness function. Each agent has information about the best location it has encountered as determined by the fitness function, as well as the best location any agent in the swarm has encountered. Using this information, the agents quickly find and converge on the single best solution.

For most intellectual tasks, and particularly for innovations, the choice between relying on information from others and obtaining information on one's own involves a tradeoff of costs and benefits. Seeking out information on one's own requires time and energy but is often more trustworthy and individually tailored than information learned by word of mouth. On the other hand, choosing to use information provided by others can be cost-effective, especially if past experience suggests that the source is reliable. These two choices have been characterized as exploration and exploitation, respectively (Holland, 1975), or the choice between searching out new information and using the best information currently available. March (1991) presented a detailed analysis of the tradeoff between the two strategies with respect to organizations. Organizations that rely mostly on exploitation of competitors' innovations benefit in the short run by saving costs on research and development but may lose out in the long run because they never lead the pack.

This individual decision, whether to imitate or explore, is also affected by the immediate social influences. Granovetter (1978) suggested that people act as though they have a threshold number of friends (or neighbors) that must adopt a solution before they will also adopt the solution themselves. In agreement with Midgley and Dowling (1978), he found that the people who were early in adopting a solution (those with a low threshold) were most influential in causing bandwagons in a population. However, if a group of innovators are relatively isolated from the rest of the population, they may have no global impact at all. Therefore it is important to also consider the paths through which an innovation can spread (Chwe, 1999).

Social Network Analysis

The structure of social networks has been shown to be critically important in determining innovation spread and is central to our

experiments. By considering the social network structure in the propagation of innovations, not only can one be more specific about the social context, but one also benefits from the tools developed for network analysis. The properties of network structures have been studied in many different arenas, including neural networks, actor collaboration networks, power grids (Watts & Strogatz, 1998), scholarly citation practices (Newman, 2001), metabolic networks (Jeong, Tombor, Albert, Oltvai, & Barabási, 2000), Web links (Albert, Jeong, & Barabási, 1999), and many more. A wide range of descriptive statistics has been developed to describe the global properties of these networks. These properties are usually defined in terms of *nodes*, which are the units or actors in a network, and *edges* or *ties*, the connections between them. The *degree* of a node is the number of edges connecting that node to other nodes. The degree of a network is the average degree of all nodes. The *geodesic path length* is the smallest number of nodes a message needs to go through to reach another node. Average shortest path lengths in networks—even large, randomly connected networks—are oftentimes surprisingly short. This property has been popularized as the notion of “six degrees of separation” connecting any two people in the world and has been experimentally supported (Milgram, 1967).

Erdős and Rényi (1959) were the first to thoroughly describe the properties of random networks, in which edges between nodes are generated such that nodes i and j are connected with some probability p . When family members draw names from a hat to decide who will exchange gifts with whom, they create a random “giving” network. Random networks tend to have a small average geodesic path length. More formally, the average path length connecting two randomly selected nodes in a random network is $\ln(N)/\ln(K)$ where N is the number of nodes and K is the degree of each node. A random network lacks spatial structure.

The network structure with the most spatial structure is a completely regular network, such as a lattice or a ring, in which the arrangement of edges form a pattern like a crystal, with distinct local neighborhoods. Messages passed in the game of telephone travel through a ring network, which is a kind of regular network. In regular lattices, the average shortest path required to connect two individuals requires going through $N/2K$ other individuals. Thus, the paths connecting people are much longer, on average, for lattice than random networks.

Watts and Strogatz (1998) demonstrated that by starting with a regular structure such as a lattice and randomly rewiring a small number of connections, the resulting *small-world* network has a low average path length but still maintains a mostly regular structure. This is because nodes that are connected to the same node tend to be spatially close themselves, but the rewired connections act as shortcuts. From an information-processing perspective, then, these are attractive networks because the spatial structure of the networks allows information search to proceed systematically, and the shortcut paths allow the search to proceed quickly (Kleinberg, 2000).

Wilhite (2001) used a computer simulation to compare market trading over various network structures. In one condition, all agents were allowed to negotiate trade with any other agent. In another, most agents could only trade locally, but a few could trade globally (i.e., outside of the local clique). In this latter small-world network, the market reached Pareto equilibrium (i.e., the state where no more trades that mutually benefit both traders can be

made) even faster than in the condition where everyone could trade with everyone. The local structure of the network constrained the search space, but the shortcuts allowed good trades to pass quickly through the group. This illustrates the benefit of small-world networks for the dissemination of information.

How these different network structures affect the propagation of innovations is the focus of our research. As noted, the structural properties of these networks constrain the manner and speed with which information can spread through them. Kennedy et al., (2001) created a variation of their particle swarm model where the agents were connected in a network, and the information about the best location encountered by each agent was shared only over the network ties. They compared a *circle* topology, in which the agents were connected in a ring, to a *wheel* topology, in which one agent was a hub and all other agents were spokes, connected only to the hub. In their model the hub in the wheel actually slows down the transmission of information, because if multiple spokes have roughly similar solutions, only one of those is chosen by the hub and shared with the other spokes. They found that in most problem spaces the circle topology found the best solution the fastest, but in a problem space with many local maxima and only one global maximum (the Rastrigin function), the wheel performed the best. Their explanation was that the wheel configuration increased exploration because of the slow communication of solutions between agents. They also noted that the fully connected topology performed best or nearly best on all of the functions, although the agents did not explore multiple regions simultaneously when fully connected.

There are few studies that use actual human behavior in groups while manipulating the communication network. Latané and colleagues (Latané & Bourgeois, 1996; Latané & L'Herrou, 1996) have studied the spread of influence through social networks passing information through e-mails. Latané and colleagues' work focuses exclusively on purely judgmental issues. Bavelas (1950) and Leavitt (1951) were two of the first researchers to study group performance in networks, noting that the communication structure of a group could aid or inhibit the ability of the group to find a solution to a problem. In the tasks they studied, the group was working cooperatively on a problem. With innovations, however, each individual is typically trying to find his or her own best solution to a problem, and then subsequently individuals can choose to imitate good previous solutions.

The studies reported in this article tie together research on the diffusion of innovation in real groups (e.g., Burt, 1987; Ryan & Gross, 1943), social psychological research on imitation and influence (Cialdini & Goldstein, 2004; Sherif, 1935), computational models of group behavior (Axelrod, 1997; Kennedy et al., 2001; Nowak, Szamrej, & Latané, 1990), and research on the flow of information through networks (Kleinberg, 2000; Latané & L'Herrou, 1996). By comparing the effect of the communication network structure on the propagation of information when humans are faced with the decision to explore the problem space or exploit good solutions to different problems, we hope to integrate and extend these diverse areas of research.

Our Paradigm

In choosing a paradigm for studying information dissemination, we sought to find a case with (a) a problem to solve with answers

that varied continuously on a quantitative measure of quality, (b) a problem search space that was sufficiently large that no individual could cover it all in a reasonable amount of time, and (c) simple communications between participants who would be amenable to computational modeling. We settled upon a minimal search task in which participants guess numbers between 0 and 100 and the computer tells them how many points were obtained from the guess. There was a continuous function that related the guesses to the points earned, but this function was not revealed to the participants. Additionally, random noise was added to the points earned, so that repeated sampling was necessary to accurately determine the underlying points obtainable from a guess. The participants received information on their own guesses and earned points, as well as obtained information on their neighbors' guesses and outcomes. In this way guesses are like solutions to a problem, and good solutions—innovations—can be roughly imitated for roughly similar results.

Examples for a group of 10 participants in each of the network structures that we compared are shown in Figure 1. Circles indicate participants, and lines connect participants who directly exchange information. Notice that three of the networks have a total of 12 connections between participants. Thus, if there is a difference in information dissemination in these networks, then it must be due to the topology, not density, of the connections. We also tested, in addition to these three network structures, a fully connected network (also called a "complete graph" in graph theory), in which everyone had access to the guesses and scores of everyone else.

In a series of three experiments, we compared the different network structures' performance on different kinds of problem spaces. In the first study, we present participants with two equally good solutions and examine the amount of bandwagoning—the degree to which participants imitate each other's solutions—in the different networks. We predict the fully connected network, with immediate transfer of information and a greater amount of information than in the other networks, will have more bandwagoning

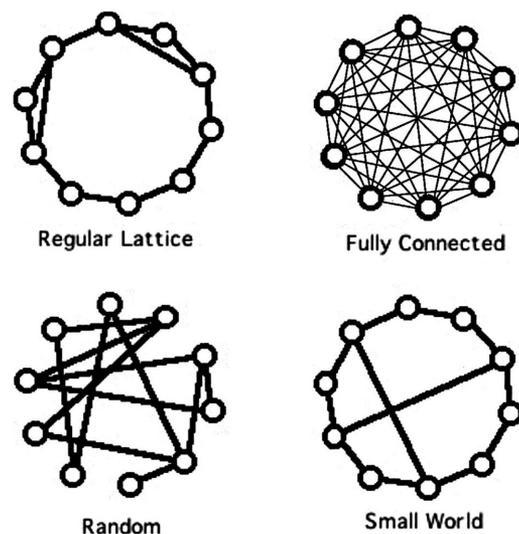


Figure 1. Examples of the different network structures for groups of 10 participants. Circles represent participants, and lines indicate communication channels.

than the other networks. In the following two studies, we focus on the ability of the groups to find the best solution in problem spaces that varied in difficulty. We hypothesize that the most efficient network structure will depend on the difficulty of finding the best solution because of the different information diffusion properties of the different networks. We expect groups in fully connected networks, because of the immediate spread of information, to outperform groups in the other networks on problems with a single, easy-to-find solution. For harder problems, with locally good but globally suboptimal solutions, we predict a need for a balance between the fast spread of information and local structure, so we expect the small-world network will outperform the other networks. For very difficult problems, with particularly hard-to-find solutions, we expect continued exploration to be necessary, and therefore the lattice network (similar to the circle topology used in Kennedy et al., 2001), with its slow spread of information, may be the best network for the problem solvers.

We examined several measures of search performance to compare the different network structures on the different payout functions. The functions we used in all studies were based on a multimodal Gaussian function, so that approximations to the best solution earned close to the best payouts. A person guessing within 0.5 *SD* of the best solution was considered sufficiently close to be “within” the maximum. To illustrate, in a unimodal payout function with a maximum of 40 and standard deviation of 12, a participant can be said to have reached the maximum if he or she guessed between 37 and 43. On the basis of this criterion of success, we could then look at how quickly group members found the best solution, the average proportion of participants guessing in the maximum, and (for Study 1, described later) the amount of bandwagoning in the various networks. We also were interested in how tightly clustered the group members’ guesses were, as an indication of whether participants were making similar guesses, and how volatile the group members’ guesses were, as an indication of how much participants were exploring the problem space. The measures we used and their purposes are listed in Table 1.

Study 1

Before comparing the effect of network structure on group performance, it seems reasonable to consider how much conformity would be evidenced independent of the quality of solutions. Completely rational beings could engage in bandwagon behavior even where there is no objective difference between solutions (Banerjee, 1992), but it is not known how network structure could affect this behavior. To this end, we created a bimodal payout function with two equal maxima (see Figure 2) to see if and when participants would largely converge on the same solution even though there is no advantage for either maximum over the other.

We expect the most bandwagoning to happen with networks that allow rapid dissemination of information. These are the networks with the shortest average path length—the fully connected network—followed by the small-world and random networks.

Method

One hundred sixty-four Indiana University undergraduate students participated for partial course credit. There were 13 groups that ranged in size from 5 to 17 people (depending on the number of students signing up for each session) with a median of 10 people per group. Due to a programming error, when a participant accidentally exited the program in the middle of a session, the generation of the remaining networks was corrupted, making all of the data for these groups unusable. The number of participants and groups reported in each study do not include these groups with corrupted data. Each session was run in a computer laboratory with 20 client computers used by the participants and one server operated by the experimenter. Participants signed on to their computer and gave themselves a handle or were assigned an ID. Once all participants had signed on, the experimenter started the session and the following instructions appeared to each of the participants:

Thank you for participating in this experiment on how ideas move from person to person in a social group. Your task is to try to accumulate as many points as possible. On each trial, you will type in a number between 0 and 100, and the computer will tell you how many points your number receives. There is a systematic relationship between the number you put in, and the points you receive, but the relationship will often be difficult for you to understand. Every time you type in the same number, it will be worth about the same number of points, but there may also be a bit of randomness added in to the earned points. Usually, numbers that are close to each other will receive similar points. At the end of each block of trials, you will be told how many points you earned, and how many points people earned in general.

In addition to telling you how many points your guess was worth, after each round of guesses, the computer will show you what numbers other people guessed, and how many points those guesses earned. You can use this information to help you decide what number to guess on the next round. Other people will also see the number that you entered, and how many points you received.

After participants read this, the controlling program created the network structure for the first of eight problems. Each problem consisted of 15 rounds in which participants had 20 s to guess a number between 0 and 100. When a round ended, the guesses were sent to the server, which would calculate each participant’s score (which was always between 0 and 50), add normally distributed noise with a mean of 0 and standard deviation of 5, and return the feedback. This began the next round, and participants now had

Table 1
Dependent Measures and Their Meanings

Dependent measure	Meaning
Average steps to guess in global maximum	Average speed of finding best solution
Average proportion of group guessing in global maximum	Overall convergence on best solution
Relative entropy (Kullback–Leibler)	Clustering of guesses over range
Volatility	Amount of exploration

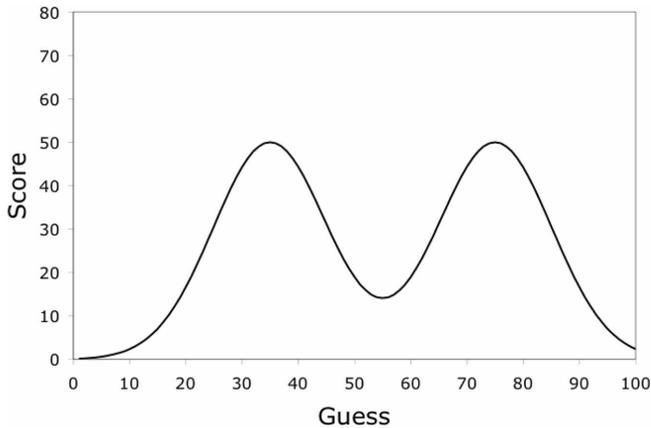


Figure 2. Study 1: An example of the equal bimodal payout function.

available their guess and score from the previous round as well as a list of their neighbors' IDs, guesses, and scores while they decided on their next guess (see Figure 3). At the end of the 15th round, participants were given feedback on their score and a message indicating the next problem would begin shortly. After 15 s, the server created a new network structure and began the next problem. We generated a fixed random ordering of the payout functions and network structures (see Table 2) that was counterbalanced with the reverse order between groups. Thus, the ordering of the payouts and networks was partially balanced, but due to the small number of groups we were not able to adequately assess the effects of the ordering. Participants searched, in addition to the bimodal payout function, a problem space with an easy-to-find local maximum and a hard-to-find global maximum (see Table 2). The results from this exponential problem space are not reported here because the global maximum was too hard to find—no one found it in any of the groups. The payout functions were all based on a multimodal Gaussian function¹, and the positions of the maxima were different for each problem. The network structure for each problem was either full, lattice, small world, or random, similar to those in Figure 1 but constructed for the different sized groups as explained below.

To create a network, the server takes all of the client computers and treats each as a node. For the random network, the server creates a number of edges equal to 1.3 times the number of nodes. These edges connect randomly selected nodes under the constraint that a path exists between every node (i.e., that the graph is connected). This is conceptually equivalent to the algorithm proposed by Molloy and Reed (1995) for generating random networks with a predefined degree distribution. For the lattice network, the server connects the clients in a ring and then randomly picks 30% of the nodes and connects each of these nodes to a neighbor two steps away. For the small-world network, the server begins by placing the clients in a ring and then picks 30% of the nodes randomly and adds a connection to another random node under the restriction that the connected nodes are at least three nodes apart following the lattice path. These probabilities ensure that the average degree is equivalent for all of these network structures. For the fully connected network, the server created edges to connect each of the nodes to all other nodes, for a total of $N(N-1)/2$ edges.

Therefore each participant had access to more information in the fully connected network than in the other three networks.

Our small-world networks are comparable to those generated in Ahmed and Abdusalam's (2000) study of percolation in networks. Our method for generating the small-world networks—unlike those used to create traditional small-world networks (e.g., Watts & Strogatz, 1998)—resulted in less clustering because neighbors of a node were not more likely to be neighbors of each other. However, they still had a small average geodesic path length and maintained the regularity of the lattice network. This regularity is evidenced by the similarly small variance in the degree (the count of each node's neighbors) in the small-world ($SD = 0.68$) and lattice networks ($SD = 0.82$) relative to the variance in the random networks ($SD = 1.14$).

Analysis

In these studies, the unit of analysis is the group, not the participants within the group. Therefore we had 13 "subjects," and the analyses were four-level (network type) repeated measures analyses of variance (ANOVAs). We did not include between-subjects variables such as group size or counterbalance order in the analysis because there were not enough groups to make the tests meaningful.

We compared the amount of imitation within a group by using three measures. First, to measure how clustered the guesses were on each round, we compared the spread of guesses to a uniform distribution by using the relative entropy (or Kullback–Leibler²) statistic. The less the guesses are uniformly spread across the total possible range of guesses, the higher the relative entropy. If there is more convergence to a single peak, the guesses will be more clustered around each other and thus will have higher relative entropy.

To see if participants were remaining in one peak or were flipping between the two peaks, we looked at the *volatility* of the guesses, which we define to be the average difference in guesses between rounds for each participant. Higher volatility indicates

¹ The formula we used to create the payout functions had the following form:

$$f(x) = a_1 e^{-[b_1(x-c_1)]^2} + a_2 e^{-[b_2(x-c_2)]^2} + a_3 e^{-[b_3(x-c_3)]^2}$$

where a represents the height of a maximum, b represents the inverse of the variance around a maximum, and c represents the location of a maximum in the range of guesses between 0 and 100. For the unimodal function, a_2 and $a_3 = 0$, and for the needle and equal functions, $a_3 = 0$.

² The Kullback–Liebler is

$$\sum_{i=0}^N p_i \log \left(\frac{p_i}{q_i} \right)$$

where p_i is the actual frequency and q_i is the expected frequency of guesses in each "bin" summed from $i = 0$ to N , the number of bins that segment the range of guesses. For our purposes we divided the range of guesses from 0 to 100 into 20 bins of 5 points each. Thus, if one participant guesses in each of the 20 bins, the relative entropy will be minimized. If all of the participants guess in one bin, the relative entropy will be maximized.

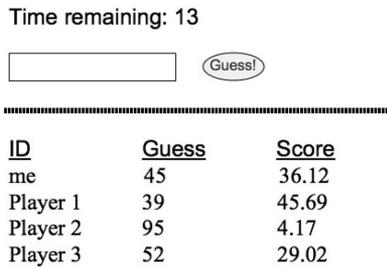


Figure 3. Participant’s view of the experiment after making a guess.

more exploration, so as participants converge on a solution, volatility is expected to decrease.

Our third measure of imitation is the difference in the number of participants guessing within 0.5 SD of each of the maxima. As it was possible for some group members to be guessing outside of either maximum, this difference was normalized by the total number of people guessing within both of the peaks. With this measure of bandwagoning, if there were an equal number of participants in each of the maxima, the difference would be zero. However, if participants were following the crowd and mostly guessing in one of the maxima, this statistic would be close to one. When there were no participants guessing in either peak, we set the value to zero (this occurred in only five cases).

Results

The first measure of conformity we report is relative entropy. There was a significant effect of the network structure, $F(3, 36) = 3.18, p < .05$, on entropy. The fully connected network was the most clustered ($M = 1.71, SD = 0.42$), followed by the lattice network ($M = 1.67, SD = 0.41$), with the random ($M = 1.53, SD = 0.43$) and small-world ($M = 1.48, SD = 0.44$) networks being the least clustered. Our hypothesis was that the fully connected network would be most clustered because of the instantaneous transmission of information and because of the additional available information. It is unclear why the lattice network also had a high degree of clustering.

Contrary to expectations, the volatility did not differ significantly between networks. As a reminder, the volatility is the

average difference between individuals’ guesses across all rounds for each group. The small-world network had the greatest amount of volatility ($M = 6.46, SD = 2.31$), and the lattice network had the least ($M = 4.26, SD = 1.56$), with the fully connected and random networks having roughly equivalent volatility (full: $M = 5.32, SD = 3.33$; random: $M = 5.34, SD = 1.53$).

Finally, we looked at our measure of conformity, the bandwagoning measure. The differences between networks was very reliable, $F(3, 36) = 4.62, p < .01$. As predicted, the fully connected network had a much higher degree of bandwagoning ($M = 0.76, SD = 0.30$) than did the other networks (lattice: $M = 0.59, SD = 0.34$; random: $M = 0.53, SD = 0.34$; small world: $M = 0.50, SD = 0.33$). We submitted these differences to post hoc comparisons with Tukey’s HSD adjustment and found the fully connected network was significantly different from the small-world network, $t(13) = 3.81, p < .05$, and none of the other networks were different from each other.

Discussion

There are many reasons why participants would converge on the same maximum when there are other equivalent solutions. People could be conforming due to normative pressures (Deutsch & Gerard, 1955), although this experimental paradigm minimizes their influence, especially in the sparse network structures in which participants received feedback from only a few neighbors. In this case, participants most likely latched onto each other’s solutions because of the perceived advantage. Once one of the two maxima had a number of participants in it, the probability increased that one of the participants would get a higher score due to noise, so other participants were more likely to see that high score and imitate the solution. In the fully connected network, because everyone had the same information, people tended to converge rapidly to whatever solution was found to have the highest score first. For the small-world network, the highly regular spatial structure and short path lengths more often led to different subgroups guessing in each of the maxima. The random noise added to the scores on each guess led to more switching between the solutions for the small-world network (as evidenced by the higher volatility and low bandwagoning) because the occasional higher score from one subgroup could quickly pass to another subgroup.

Table 2
Order of Network and Payout Functions in Studies 1–3

Order	Study 1		Study 2		Study 3	
	Network	Function	Network	Function	Network	Function
1	Lattice	Exponential	Lattice	Unimodal	Small world	High noise
2	Small world	Equal	Small world	Multimodal	Lattice	Needle
3	Random	Equal	Random	Multimodal	Small world	Low noise
4	Full	Exponential	Full	Unimodal	Full	High noise
5	Random	Exponential	Random	Unimodal	Full	Needle
6	Lattice	Equal	Lattice	Multimodal	Random	Needle
7	Small world	Exponential	Small world	Unimodal	Full	Low noise
8	Full	Equal	Full	Multimodal	Small world	Needle

Note. The order for half of all groups was the reverse of the order shown here. Cells with bold text are problem spaces not reported in this article.

Study 2

In this study, we compared two payout functions. The *unimodal* function has a single best solution that could be found with a hill-climbing method (for example, see Figure 4a). This is like searching for the best guitar to buy in a town with only one guitar shop. The *multimodal* function (having multiple local maxima, or *modes*) increased the difficulty of the search by introducing local maxima. A local maximum is a solution that is better than all of its immediate neighboring solutions yet is not the best solution possible. Thus, a simple hill-climbing method might not find the best possible solution. In a town with more than one guitar shop, one might find the best guitar in one of the stores, but there might be another shop that has an even better guitar. Figure 4b shows one of the multimodal functions used; it has three peaks, but one of the peaks is somewhat higher than the other two.

The participants were reminded that they were to maximize their points over the rounds. So while exploration might lead to a higher payoff, the search might also lead to lower points, imposing a cost in deviating from the best solution found so far. The basic prediction is that this tradeoff in exploration and exploitation will predict which network will be optimal for which problem space. In the unimodal problem space there is no benefit to increased exploration after one has found the locally best solution, because the

locally best solution is also the globally best solution. Therefore, those networks that have the fastest dissemination of information are expected to perform best. These are the ones with the shortest path length—the fully connected network, followed by the small-world and random networks. However, in the multimodal problem space, too much reliance on the currently best available information could cause premature convergence on a local maximum. For this reason we predict the small-world networks will be best fit to this type of problem space, as they have fast transmission of information but also local structure that encourages hill-climbing search and prevents bandwagoning behavior, as shown in Study 1.

Method

One hundred fifty Indiana University undergraduate students participated for partial course credit in nine groups ranging in size from 7 to 18 people, with a median of 10 people per group. The procedure was the same as in Study 1, except that participants searched both unimodal and multimodal problem spaces. Again, a fixed random ordering of the payout functions and network structures was counterbalanced with the reverse ordering between groups (see Table 2). In this study we directly compared performance between networks and problem spaces. The design of this experiment was a 4 (network structure; within-subjects) \times 2 (problem space; within-subjects) repeated measures design.

Results

Our first performance metric is the *speed of convergence*, defined as the average number of rounds that group members took to guess within 0.5 *SD* of the global maximum (this was set to 15, the number of rounds, if the participant never guessed within the maximum). The repeated-measures ANOVA revealed that there was a main effect of the problem space, $F(1, 24) = 30.4, p < .001$, such that individuals were much faster at finding the global maximum in the unimodal problem space ($M = 3.89, SD = 1.62$) than in the multimodal problem space ($M = 6.44, SD = 3.04$). This is to be expected, as the unimodal function was designed to be an easier search problem than was the multimodal function. There was also a main effect of the network, $F(3, 24) = 3.91, p < .05$, indicating that the structure of the network affected the speed of convergence. Post hoc comparisons with Bonferroni correction showed that individuals took marginally significantly longer on average to find the global maximum in the lattice network ($M = 6.65, SD = 3.88$) than they did in the fully connected network ($M = 4.46, SD = 2.25$), $t(8) = 2.95, p < .06$, and the small-world network ($M = 4.43, SD = 1.6$), $t(8) = 2.98, p < .06$.

Most importantly, the within-subjects analysis showed an interaction between the network type and problem space on the time to find the global maximum, $F(3, 24) = 3.98, p < .05$, shown in Table 3. It is interesting that in the unimodal problem space, individuals in the fully connected network were the fastest to discover the global maximum, while in the multimodal problem space, those in the small-world network were the fastest on average to find the global maximum. Post hoc comparisons with Bonferroni corrections showed that in the unimodal problem space, only the fully connected network was marginally significantly faster than the slowest (lattice) network, $t(8) = 2.91, p < .06$, while in the multimodal problem space, only the small-world

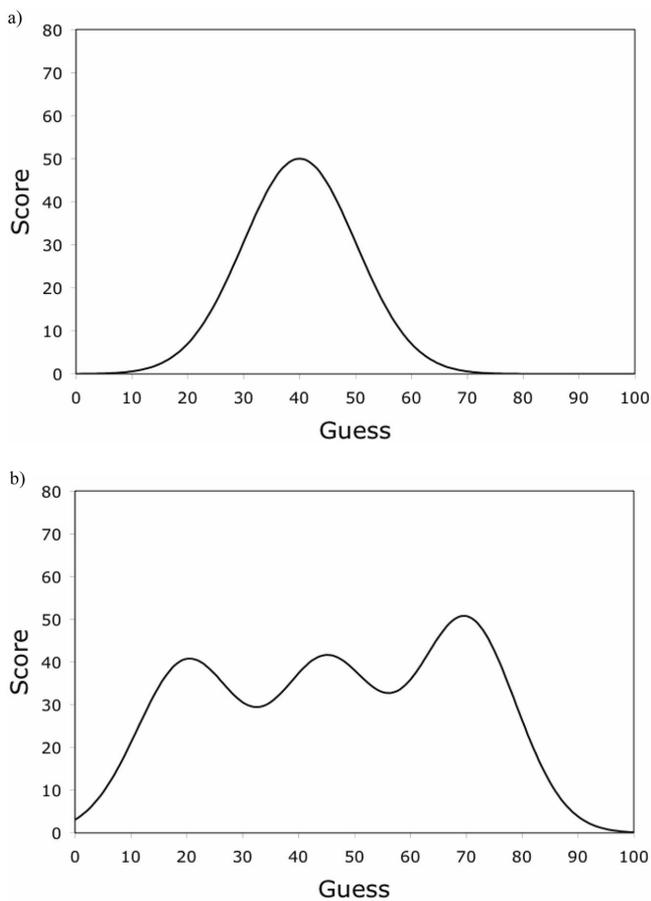


Figure 4. Study 2: Examples of the (a) unimodal and (b) multimodal payout functions.

Table 3
Comparison of Group Performance in Unimodal and Multimodal Problem Spaces

Network	Average rounds to find global maximum (<i>SD</i>)		Average percent in maximum (<i>SD</i>)	
	Unimodal	Multimodal	Unimodal	Multimodal
Full	3.08 (1.07)	5.83 (2.32)	74.9 (10.66)	50.6 (16.51)
Lattice	4.79 (2.26)	8.51 (4.38)	56.77 (22.9)	38.8 (29.24)
Small world	4.18 (1.64)	4.69 (1.61)	70.25 (13.35)	55.94 (17.54)
Random	3.51 (0.81)	6.73 (2.10)	72.05 (9.77)	42.08 (18.83)

network was significantly faster than the slowest (lattice) network, $t(8) = 3.4, p < .05$.

Our second performance metric was the average proportion of individuals guessing in the global maximum over all rounds. This measure is higher when the group members not only find the maximum quickly, but also continue to guess within the global maximum. This requires the participants to diminish or extinguish their exploration of the space after finding the global maximum. The repeated-measures ANOVA of this variable revealed the expected main effect of problem space, $F(1, 24) = 37.3, p < .001$, where the proportion of individuals guessing in the global maximum on average in the unimodal problem space ($M = 0.685, SD = 0.161$) was significantly more than those in the multimodal problem space ($M = 0.469, SD = 0.213$). Additionally, there was a main effect of network on the average percent in the global maximum, $F(3, 24) = 4.55, p < .05$. As with the speed of convergence, post hoc comparisons with Bonferroni correction showed that a smaller proportion of people in the lattice network guessed in the global maximum on average ($M = 0.478, SD = 0.271$) than in the fully connected network ($M = 0.627, SD = 0.184$), $t(8) = 3.15, p < .05$, and the small-world network ($M = 0.63, SD = 0.168$), $t(8) = 3.23, p < .05$.

Most importantly, there was a marginally significant interaction between the network and the problem space, $F(3, 24) = 2.81, p < .07$, also shown in Table 3. The analysis of the proportion guessing in the global maximum corroborates the analysis of the average number of rounds to find the global maximum. In the unimodal problem space, the fully connected network had the highest proportion guessing in the global maximum, while in the multimodal problem space, the small-world network had the highest proportion. Post hoc comparisons reveal that in the unimodal problem space, the lattice network had a significantly smaller proportion of participants guessing in the global maximum than did all three other networks: full: $t(8) = 3.98, p < .02$; random: $t(8) = 3.35, p < .03$; small world: $t(8) = 2.957, p < .06$. In the multimodal problem space, however, only the small-world network had a marginally significantly greater proportion of participants guessing in the global maximum than did the lattice network, $t(8) = 2.67, p < .09$.

The effects of network and problem space on speed of convergence and proportion guessing in the maximum can be visualized with a plot of the proportion guessing in the global maximum over rounds (see Figures 5a and 5b). In the unimodal problem space (see Figure 5a), all networks except the lattice converge quickly on the maximum. In the multimodal problem space (see Figure 5b), it

takes longer for all networks to find the maximum and the overall convergence is less, but the small-world network is faster than all of the other networks.

We were also interested in the degree of clustering in the different networks and problem spaces. The relative entropy largely mirrored the speed of convergence and percent guessing in the maximum. There was a main effect of the problem space on relative entropy, $F(1, 24) = 20.18, p < .005$, such that the clustering was much higher in the unimodal problem space ($M = 1.71, SD = 0.33$) than in the multimodal problem space ($M = 1.49, SD = 0.37$). This was expected, as people had an easier time finding the maximum and converged more in the unimodal problem space than in the multimodal problem space. There was also a marginally significant main effect of network on relative entropy, $F(1, 24) = 2.56, p < .08$. The lattice network had less clustering ($M = 1.51, SD = 0.4$) than did the other networks, with the fully connected network having the highest clustering ($M = 1.65, SD =$

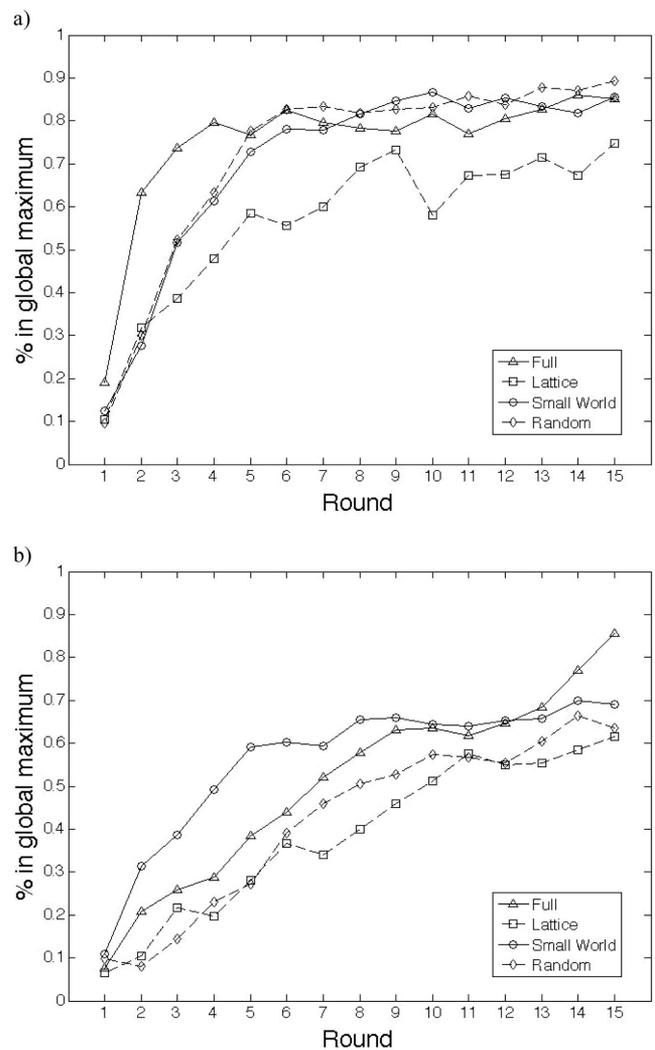


Figure 5. Study 2: Percent of participants within 1 *SD* of the global maximum on each round in the (a) unimodal and (b) multimodal payout function.

0.34), followed by the small-world network ($M = 1.62$, $SD = 0.36$) and the random network ($M = 1.6$, $SD = 0.36$).

There was also a significant interaction between the problem space and the network on relative entropy, $F(3, 24) = 5.72$, $p < .005$. In the unimodal problem space, the fully connected network had the most clustering ($M = 1.81$, $SD = 0.29$) and the lattice network had the least ($M = 1.5$, $SD = 0.39$), with little difference between the small-world ($M = 1.76$, $SD = 0.31$) and the random networks ($M = 1.76$, $SD = 0.27$). In the multimodal problem space, however, there was almost no difference between the networks. The lattice network had the highest relative entropy ($M = 1.53$, $SD = 0.44$), followed by the fully connected ($M = 1.50$, $SD = 0.32$) and the small-world network ($M = 1.48$, $SD = 0.37$), with the random network having the lowest relative entropy ($M = 1.45$, $SD = 0.39$).

There were no main effects of network or problem space on the volatility of guesses, but there was a significant interaction between the two, $F(3, 24) = 6.93$, $p < .005$. In the unimodal problem space, the fully connected network had the lowest volatility ($M = 3.73$, $SD = 1.43$), followed by the random network ($M = 5.19$, $SD = 0.85$) and the small-world network ($M = 6.0$, $SD = 1.44$), with the lattice having the highest volatility ($M = 7.82$, $SD = 4.14$). In the multimodal problem space, the pattern was very different: The lattice had the lowest volatility ($M = 4.43$, $SD = 1.44$), followed by the small-world network ($M = 4.98$, $SD = 1.87$) and random network ($M = 5.12$, $SD = 1.36$), with the fully connected having the highest volatility ($M = 6.28$, $SD = 1.4$).

Discussion

When there was only one good solution—when the payout function was unimodal—there was a direct relationship between the average shortest path length and the speed with which the group converged on the best solution. In this case, the fully connected network converged more quickly than did the other three networks. The lattice network took longer to converge on the best solution because the advantageous innovations had to work their way through longer chains of people, and only about half of the group members converged on the maximum.

When the problem space had multiple good solutions that were nonetheless suboptimal, the story was different. In this case the small-world network groups tended to find the best solution faster and converged on the global maximum more robustly than did every other network—even the fully connected network, in which everyone had complete information about every other participant's guesses and scores. As there is no reason to expect the decision-making processes of the individuals to change from one network type to the next, and the type of information presented to them was the same (neighbors' guesses and scores), the differences between the network types must be due to the information transmission properties of the networks. The advantage of the small-world over the fully connected network is akin to a novel group-based form of the “less is more” effect reported in individual decision-making literature (Gigerenzer & Todd, 1999).

This somewhat counterintuitive result—that limiting the available information might actually improve a group's performance—is a result of the way the groups were searching the problem space. In the fully connected network, participants would often latch onto the first good solution that was found, and this was

the best solution only one third of the time. When the group converged prematurely on a local maximum, it took longer for an adventurous (or bored) participant to explore and find the globally best solution. In the small-world network, however, the participants were segregated by the regular, latticelike connections, but the information could travel quickly through “shortcuts,” allowing for different locally connected groups to explore different regions of the problem space. Thus, while one locally connected group might latch onto a local maximum, the small-world topology decreased the probability that everyone would follow their lead before another subgroup found the global maximum.

Study 3

In Study 2, the global maximum was just as easy to find as either of the two local maxima. However, in some cases the best solution is harder to find than other solutions. For instance, the most exclusive and best restaurant might not be located near any other restaurants or even have a sign outside! In these cases, prolonged exploration of the problem space can result in a higher payoff than does rapid convergence on an easy-to-find but suboptimal solution. As with the multimodal payout function, networks with a well-defined spatial structure will allow continued exploration. However, with a hard-to-find problem, even more exploration may be necessary before groups converge, and so networks with long path lengths (such as the lattice networks) may be more successful in finding the best solution than would the other networks. To examine this situation, we created a bimodal payout function (hereafter referred to as the “needle” function) with one wide local maximum and one thin, hard-to-find global maximum (see Figure 6).

Method

One hundred fifty-seven Indiana University undergraduate students participated for partial course credit in 10 groups ranging in size from 7 to 19 people, with a median of 12.5 people per group. The procedure was the same as in Study 1 and Study 2, except for using the needle payout function instead of the bimodal, unimodal, or multimodal functions. As in Study 1 the experiment was a four-level (network structure), within-subjects design. In addition, these groups searched a bimodal problem space with a global maximum in which we varied the random error on the feedback

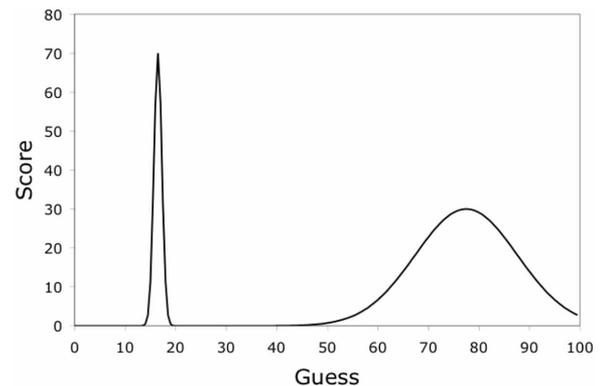


Figure 6. Study 3: An example of the “needle” payout function.

(see Table 2), not reported here. All analyses were repeated-measures ANOVAs unless otherwise reported.

Results

In the needle payout function, the global maximum has a small range (guesses in the neighborhood of the global maximum yielded a positive score only when within two guesses of the maximum), so the longer that participants explore, the more likely one of them is to find it. Unfortunately, this also means that participants may be unlikely to find it at all, which limits the power of the analyses. In fact, the number of groups in which any participants found the maximum was small. In the lattice network, 7 of the 10 groups had at least one participant who found the maximum, while in the fully connected and random networks only 4 of the 10 groups found it, and in the small-world network only 3 of the 10 groups found it. Thus, it is difficult to make strong statistical conclusions on the basis of the limited number of groups that successfully found the needle in the problem space.

We predicted that the more spatially segregated networks would be more likely to find the needle, and that seems to be the case, as the groups in the lattice network were more likely to find the maximum. This can also be seen in the average proportion of participants in the global maximum over all rounds. The lattice network had the highest average ($M = 0.21$, $SD = 0.25$), followed by the fully connected network ($M = 0.17$, $SD = 0.23$), with the small-world ($M = 0.1$, $SD = 0.17$) and random ($M = 0.08$, $SD = 0.18$) networks trailing behind. Because of the small number of participants who found the maximum at all, however, this difference was not significant, $F(3, 27) = 0.85$.

This pattern is even more dramatic when focusing on the final proportion of participants in the global maximum. In the last round, groups in the lattice network had on average 41.9% of participants in the global maximum ($SD = 37.3$), compared with just 27.5% in the fully connected network ($SD = 36.5$), 19.5% in the small-world network ($SD = 31.6$), and a mere 14.8% in the random network ($SD = 31.1$). Unfortunately, even this large difference was not statistically reliable, $F(3, 27) = 1.48$. Looking at the progression over time, we concluded that the lattice network clearly has the predicted advantage. As can be seen in Figure 7a, the proportion of participants in the lattice network guessing in the global maximum is steadily increasing over rounds, and at a faster rate than that of any other network. The picture is even sharper when looking at the proportion guessing in the local maximum (see Figure 7b). Most participants in the lattice network guess in the local maximum initially, but over time the proportion is steadily decreasing, unlike with the other networks. This could suggest that over time the likelihood of an individual finding the global maximum increases and that the information about the location of that maximum slowly percolates through the network.

The difference in the clustering of the guesses was reliably different between networks, $F(3, 27) = 10.41$, $p < .001$. The fully connected and lattice networks had much higher relative entropy ($M = 1.8$, $SD = 0.29$, and $M = 1.72$, $SD = 0.34$, respectively) compared with the small-world network ($M = 1.48$, $SD = 0.18$) or the random network ($M = 1.27$, $SD = 0.4$). This indicates that the fully connected and lattice networks had higher clustering, but when taken together with the results of the proportion guessing in the global versus local maxima, it indicates that the lattice network

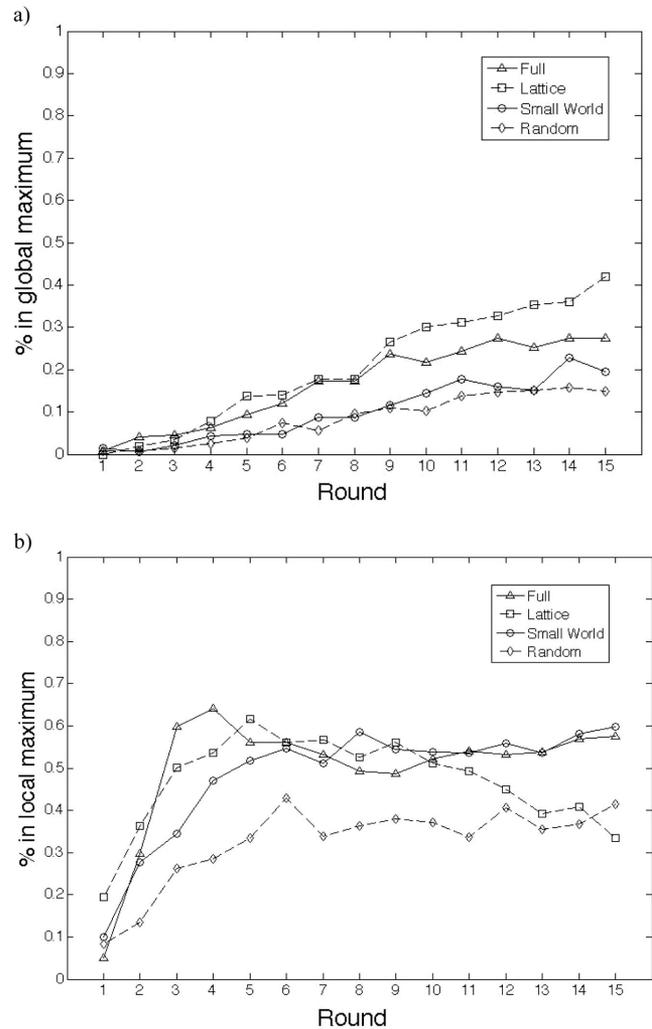


Figure 7. Study 3: Percent of participants within 0.5 SD of (a) the global maximum (the “needle”) and (b) the local maximum.

was more often clustered in the global maximum rather than in the local maximum.

Discussion

The payout function used in Study 3 represents situations in which a problem has (a) a precise best solution that is not easily approximated and (b) a lesser solution that is easy to roughly imitate. In this study, neither the fully connected nor small-world network performed best. In fact, there was a trend for the lattice to be the best performing network. This is surprising because the lattice network was the worst performing network for the easier problem spaces. The pattern of results could suggest this is due to the increased exploration engendered by the long path lengths and local, spatial neighborhoods preserved in the lattice networks, or conversely, the increased probability of the other networks to converge on the suboptimal local maximum.

General Discussion

In these studies, participants searched a problem space as a group, sharing information about solutions by way of various social network structures. The work reported here provides additional evidence that diffusion of innovation in groups is strongly affected by the structure of the communication channels available to members of the group. More importantly, it shows that different network structures are best fit to different problem spaces.

As expected from previous research, participants tended to converge on a single solution in a problem space with two equally good solutions, indicating a tendency to bandwagon even when there are two optimal solutions. This bandwagoning behavior occurs even when normative social influence is minimized by providing anonymity and eliminating communication outside of the guesses and scores. Network type moderated this effect; fully connected networks had the most bandwagoning, and small-world networks had the least.

In a unimodal problem space, where there is a single best solution that is better than all similar solutions, the best network structure is one in which information about good solutions travels as quickly as possible in a systematic way. Among the networks we studied, this was achieved to the greatest extent in the fully connected networks, followed by the small-world networks.

The needle problem space—in which there are two locally optimal solutions, one of which is easy to find but not as good as the other, more difficult-to-find solution—showed a different pattern. In this case, a high degree of exploration of the problem space is beneficial, because it increases the chances that some individual in the group will find the needle. The network structure with a long average geodesic path length and highly regular structure will be slowest to converge on a groupwide solution and therefore will continue to have group members exploring the problem space. Although it wasn't significant, in our study the lattice network tended to outperform the other networks in the needle problem space, which was in stark contrast to its poor performance in the other problem spaces.

In a multimodal problem space, however, neither the lattice nor the fully connected network performed optimally. In these problem spaces, there were three solutions that were better than all other solutions, but only one was globally the best. In this case, rapid convergence as seen in the fully connected network can lead to a locally good but globally suboptimal solution, but prolonged exploration as found in the lattice network only reduces the speed and extent of convergence of the group. In this case, the best performance tended to be in the small-world network, which possesses both preserved spatial neighborhoods and long-distance connections.

It appears as though the fit between a given network structure and a problem space depends on the amount of exploration required by the network. For the network structures we studied, the lattice promotes the most exploration, followed by the small-world and random networks, with the fully connected network producing the least exploration. The needle payout function requires the most exploration to find the global maximum, followed by the multimodal and then the unimodal. Given the tradeoff between exploration of a problem space and exploitation of good solutions (Holland, 1975; March, 1991), this tradeoff seems to be highly relevant to the ability of a group to succeed at our task.

Lazer and Friedman (2005) used an agent-based computational model to compare the performance of various networks when group members are searching different problem spaces for the globally optimal solution. In this case, the agents are searching an "NK" problem space, in which each digit in a string of N numbers is dependent on K other digits for computing the contribution of that digit to the score of the string. In this way, by varying K relative to N , the "ruggedness" of the problem space can be manipulated. When K is zero, the problem space has a single maximum. When $K = N - 1$, the performance of any single solution offers no information about adjacent solutions. In between these values of K , there is a gradient between adjacent solutions, but the entire problem space has local maxima as well as a global maximum. In their simulation, agents started out with a random string and either imitated their neighbors in the network if their neighbors' scores were better or mutated their string by a digit if the mutation would result in a higher score.

Lazer and Friedman (2005) then compared the performance of the group as whole with respect to convergence on the global maximum, varying the communication network structure between agents and the ruggedness of the problem space. As expected, they found that on simple problem spaces with little or no local maxima, networks with smaller average path lengths, such as fully connected networks and small-world networks, converged on the globally optimum solution most quickly. They also found that networks with slow transmission of information, such as lattice networks, engendered more exploration and therefore in the long run ended up outperforming the fully connected and small-world networks in more "rugged" problem spaces by finding and converging on a better solution, supporting our conclusions.

One possible extension of this work is to model the decision strategies that individuals within the groups are using when approaching the task. We expect that participants have essentially four pieces of information that could be influencing their guesses: their last guess, their best guess, their neighbors' last best guess, and their neighbors' best guess. By categorizing participants' guesses as falling within a certain range of each of these sources of information, we can estimate the relative influence of each of these sources, and when a guess falls outside the range of any of these sources of influence, we can say the participant is exploring. However, there is also ambiguity in the data, as a single guess may be categorized as influenced by multiple sources of information, which makes it difficult to differentiate the strategies. Nonetheless, a preliminary analysis shows that, as would be expected, the amount of exploring decreases as time proceeds, and this varies according to the network and problem space. Future computational modeling could have agents using these different strategies in different networks to compare against the observed distribution of guesses and performance of the networks in the different problem spaces.

Research on the benefits of network structure on the flow of information has often focused on the positive properties of small-world networks, such as the spatial structure and short path lengths (Kleinberg, 2000; Wilhite, 2001). The results of our research cast this view in the wider perspective of fit between network structure and problem space, highlighting the importance of exploration versus imitation. Broadly speaking, these results have implications for many different areas of study. For research on group performance and organizational psychology, this highlights the impor-

tance of the communication patterns within a group with respect to the type of problem being approached by the group. For example, when working on solving a problem for which the solution is presumably difficult to find, research and development programs in organizations may benefit by limiting the communication between researchers. For sociology and cultural psychology, the different amount of bandwagoning with respect to network structure and size is important. Additionally, the results speak loosely to the advantages and disadvantages inherent in the increased information transmission afforded by the Internet. While good ideas may spread quickly through the broad Internet network, it may result in too little diversity in ideas, or in the rapid spread of suboptimal ideas.

Ultimately, the paradigm developed here can be used to study the problem-solving abilities of groups under a wide range of conditions. For instance, different communication structures—such as scale-free networks, which are increasingly observed in a wide range of real networks (Barabási & Albert, 1999), or hierarchies, which are interesting because they are a typical organizational structure—could be tested. Additionally, different problem spaces—including multidimensional and dynamically evolving problem spaces—remain to be explored. It seems reasonable to predict that a network structure that permits a group to quickly converge upon a solution may be less fit when the problem space changes.

References

- Ahmed, E., & Abdusalam, H. A. (2000). On social percolation and small world network. *European Physical Journal B*, 15, 569–571.
- Albert, R., Jeong, H., & Barabási, A.-L. (1999, September 9). The diameter of the World Wide Web. *Nature*, 401, 130–131.
- Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *Journal of Conflict Resolution*, 4, 203–226.
- Bandura, A. (1965). Behavioral modification through modeling procedures. In L. Krasner & L. P. Ulmann (Eds.), *Research in behavior modification: New development and implications* (pp. 310–340). New York: Rinehart and Winston.
- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, 107, 797–817.
- Barabási, A.-L., & Albert, R. (1999, October 15). Emergence of scaling in random networks. *Science*, 286, 509–512.
- Bass, F. M. (1969). A new-product growth model for consumer durables. *Management Science*, 15, 215–227.
- Bavelas, A. (1950). Communication patterns in task-oriented groups. *Journal of the Acoustical Society of America*, 22(6), 725–730.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashions, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100, 992–1026.
- Burt, R. (1987). Social contagion and innovation: Cohesion versus structural equivalence. *American Journal of Sociology*, 92, 1287–1335.
- Chwe, M. S.-Y. (1999). Structure and strategy in collective action. *American Journal of Sociology*, 105, 128–156.
- Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55, 591–621.
- Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *Journal of Abnormal and Social Psychology*, 51, 629–636.
- Erdős, P., & Rényi, A. (1959). On random graphs. *Publicationes Mathematicae*, 6, 290–297.
- Gigerenzer, G., & Todd, P. (1999) *Simple heuristics that make us smart*. New York: Oxford University Press.
- Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83, 1420–1443.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor, MI: University of Michigan Press.
- Jeong, H., Tombor, B., Albert, R., Oltvai, Z. N., & Barabási, A.-L. (2000, October 5). The large-scale organization of metabolic networks. *Nature*, 407, 651–654.
- Kaplan, M. F., & Miller, C. E. (1987). Group decision making and normative versus informational influence: Effects of type of issue and assigned decision rule. *Journal of Personality and Social Psychology*, 53(2), 306–313.
- Kennedy, J., Eberhart, R. C., & Shi, Y. (2001). *Swarm intelligence*. San Francisco: Morgan Kaufmann.
- Kleinberg, J. (2000, August 24). Navigation in a small world. *Nature*, 406, 845.
- Latané, B., & Bourgeois, M. J. (1996). Experimental evidence for dynamic social impact: The emergence of subcultures in electronic groups. *Journal of Communication*, 46, 35–47.
- Latané, B., & L'Herrou, T. L. (1996). Spatial clustering in the conformity game: Dynamic social impact in electronic groups. *Journal of Personality and Social Psychology*, 70, 1218–1230.
- Laughlin, P. R., & Ellis, A. L. (1986). Demonstrability and social combination processes on mathematical intellectual tasks. *Journal of Experimental Social Psychology*, 22, 177–189.
- Lazer, D., & Friedman, A. (2005). The hare and the tortoise: The network structure of exploration and exploitation. *ACM International Conference Proceeding Series*, 89, 253–254.
- Leavitt, H. J. (1951). Some effects of certain communication patterns on group performance. *Journal of Abnormal Psychology*, 46, 38–50.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2, 71–87.
- Midgley, D. F., & Dowling, G. R. (1978). Innovativeness: The concept and its measurement. *Journal of Consumer Research*, 4, 229–242.
- Milgram, S. (1967). The small world problem. *Psychology Today*, 2, 60–67.
- Molloy, M., & Reed, B. (1995). A critical point for random graphs with a given degree sequence. *Random Structures & Algorithms*, 6, 161–179.
- Newman, M. E. J. (2001). The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98, 404–409.
- Nowak, A., Szamrej, J., & Latané, B. (1990). From private attitude to public opinion: A dynamic theory of social impact. *Psychological Review*, 97, 362–376.
- Rogers, E. M. (1962). *Diffusion of innovations*. New York: Free Press.
- Rogers, E. M. (1995). *Diffusion of innovations* (4th ed.). New York: Free Press.
- Rogers, E. M., & Shoemaker, F. F. (1971). *Communication of innovations: A cross-cultural approach*. New York: Free Press.
- Ryan, B., & Gross, N. C. (1943). The diffusion of hybrid seed corn in two Iowa communities. *Rural Sociology*, 8, 15–24.
- Sherif, M. (1935). A study of some social factors in perception. *Archives of Psychology*, 187, 60.
- Tarde, G. (1903). *The laws of imitation*. New York: Holt.
- Watts, D. J., & Strogatz, S. H. (1998, June 4). Collective dynamics of “small-world” networks. *Nature*, 393, 440–442.
- Wilhite, A. (2001). Bilateral trade and ‘small-world’ networks. *Computational Economics*, 18, 49–64.

Received January 11, 2008

Revision received March 14, 2008

Accepted March 18, 2008 ■