Open access • Journal Article • DOI:10.1037/A0012798

Propagation of innovations in networked groups. — Source link

Winter Mason, Andy Jones, Robert L. Goldstone

Institutions: Yahoo!, University of North Carolina at Chapel Hill, Indiana University

Published on: 01 Aug 2008 - Journal of Experimental Psychology: General (American Psychological Association)

Propagation of Innovations in Networked Groups

Winter Mason

Indiana University

Andy Jones

University of North Carolina

Robert L. Goldstone

Indiana University

Abstract

A novel paradigm was developed to study the behavior of groups of networked people searching a problem space. We 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 alternate 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.

Diffusion of Information and Innovations

Humans are uniquely adept at adopting one others' innovations. While imitation is commonly thought to be the last resort for dull and dim-witted individuals, cases of true imitation are rare among non-human animals (Blackmore, 1999), requiring complex cognitive processes of perception, analogical reasoning, and action preparation. This is because true imitation requires understanding the intentions of the model, rather than simply the behaviors. Gergely, Bekkering, and Király (2002), replicating a study by Melztoff (1988), showed that infants can infer the intentions of an adult model rather than simply copying the behavior. This capacity for imitation has been termed "no-trial learning" by Bandura (1965), who stressed that, by imitating one another, people perform behaviors that they would not have otherwise considered. When combined with variation and adaptation based on reinforcement, imitation is one of the most powerful methods for quick and effective learning. Cultural identity is largely due to the dissemination of concepts, beliefs, and artifacts across people. The tendency for humans to imitate is so ubiquitous that Meltzoff (1988) has even suggested that humans be called "Homo imitans."

This tendency is also evident in conformity in social groups. One famous experiment on conformity is Asch's (1956) line judgment experiment, in which a series of confederates announced that a pair of lines was the same length, even though they differed by several inches. When the participants had to make their judgment, roughly one-third conformed to the group and declared the same pair of lines to be equal. There are several reasons why the participants may have conformed to the objectively incorrect judgment. Influences can be classified into two types (Deutsch & Gerard, 1955): normative influence, 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. Cialdini and Goldstein (2001) note that having an affiliation goal, making one susceptible to normative influence, and having an accuracy goal, making one susceptible to informational influence, are both subsumed by a general goal of maintaining one's self-concept. The different kinds of influence affect the stability of the influence. With normative influence, people are less likely to conform to the group's opinion when responding privately than when responding publicly (Deutsch & Gerard, 1955). However, when the influence is informational (e.g., Sherif, 1935), people are more likely to continue to conform to the group's opinion privately. At a broader level, the "I'll do it if you do it" mentality can lead to situations in which a small number of people can initiate a positive feedback cycle, an effect that has been popularized as the notion of "tipping points," (Gladwell, 2000). This effect is viscerally apparent if you have ever witnessed a single person start a chant at a stadium that catches on until everyone in the stadium is participating in the chant.

The spread of a chant in a stadium, though operating through a process of conformity, behaves in a manner very similar to disease contagion. In fact, the spread of affect in a population has been described as "mood contagion," (Neumann & Strack, 2000), and the spread of attitudes and beliefs through a population has been modeled as a similar process. Nowak, Szamrej, and Latané (1990) used a computational model to study the dynamic impact of social influence. Friedkin and Johnsen (1999) developed a comparable mathematical model of influence with two components, the social process and the social influence structure. The first is the way in which people modify their opinions in response to influence, and the second is the structure of communications channels and the strength of influence between people in a group. These influence models provide a more specific and formal description of conformity processes within a group.

Laughlin and Ellis (1986) describe a continuum between intellective and judgmental issues. Where an issue falls on this continuum depends on the demonstrability of the correctness of solutions. A math problem, which has accepted means for showing the correctness of a solution would be considered an intellective task, whereas which flavor of ice cream tastes better would be judgmental. However, many problems fall in between this continuum, such as who to hire as a job candidate. While there are objective and accepted means of showing the advantage of one candidate over another on certain dimensions, such as the number of publications, there can be differences of opinion as to the quality of the publications or even which dimensions are most important for hiring a candidate. Kaplan and Miller (1987) looked at the difference in conformity for intellective and judgmental issues with normative and informational influence. Their somewhat intuitive conclusion was that intellective issues led to more informational influence while judgmental issues led to more normative influence.

The dissemination of innovations in a population is also the result of social influence and conformity and so can manifest the same bandwagon behavior. An innovation may spread because it is clearly better than the alternatives, so the conformity is due to the informational influence of others. Innovations may also function more like opinions when it is more difficult to determine if the innovation is better than current practice, or in other words, when the advantage of the innovation is less demonstrable and therefore is less of an intellective issue. If there is no inherent difference, or when the benefits of adopting an innovation are largely due to others using it (e.g., Macintosh vs. IBM or BetaMax vs. VHS), the innovation is even more like a judgmental issue.

So innovations can spread in the same manner as crowd chants, affect, and attitudes. This spread of innovations has been studied from many different perspectives beginning with Gabriel Tarde (1903), a sociologist who noted that the number of people adopting an innovation over time produces a sigmoidal curve for many different kinds of innovation, with slow adoption at first, then a sudden increase that levels off once the innovation has almost fully saturated the population. Ryan and Gross (1943) studied the adoption of hybrid-corn use by Iowa farmers, focusing on individual differences in time of adoption. Bass (1969) borrowed differential equations from physics to model the population-level change in the number of people choosing to use an innovation, and reproduced the S-shaped curve suggested by Tarde. This diffusion model is so popular that it is commonly referred to as the "Bass model," and the term "diffusion of innovation" has become ubiquitous in the field. For a good review of the literataure on the adoption of innovations, see Rogers (1962, 1995), and more recently, Wejnert (2002).

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 (Holland, 1975), and March (1991) presented a detailed analysis of the tradeoff between the two with respect to organizations. In those cases, organizations that rely mostly on exploitation of competitor's innovations benefit in the short run by saving costs on research and development, but lose out in the long run because they never lead the pack.

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, and found that the people who were early in adopting a solution (those with a low threshold) were most influential in causing bandwagoning in a population. Michael Chwe (1999) extended this threshold model and found that the network position of an individual could be more important than their threshold with respect to causing other people to follow the crowd. Valente and Davis (1999) also noted that opinion leaders, those individuals who are most central in a communication network, are the most influential in starting a bandwagoning process. This highlights the importance of another factor in the diffusion of information and innovations -- the social network structure.

Social Network Analysis

The properties of network topologies have been studied in many different arenas, including neural networks, actor collaboration networks, power grids (Watts & Strogatz, 1998), scholarly citation links (Newman, 2001), metabolic networks (Jeong, Tombor, Albert, Oltvai, & Barabasi, 2000), Web links (Albert, Jeong, & Barabasi, 1999), and many more. A wide range of statistics has been developed to describe the global properties of these networks. These properties are usually defined in terms of the nodes, which are the units or actors in a network, and edges, the connections between them.

First, 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. Second, the *geodesic path length* is the smallest number of nodes a message needs to go through to link two nodes. Average shortest path lengths in networks, even large, randomly connected networks, are often times 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). The *clustering coefficient* is the proportion of directly connected neighbors of a node that are themselves directly connected with each other (in other words, forming a triangle), which can be thought of as the “cliquishness” of a network. The actual values of these measures for the networks in Study 1 are shown in Table 1.

Erdös & Rényi (1959) were the first to thoroughly describe the properties of random networks, in which edges between nodes are generated such that Node i and Node j are connected with some probability p. When a family draws 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. With a large number of nodes random networks tend to have a small clustering coefficient, although with fewer nodes the probability of three nodes forming a triangle is higher, and so the clustering coefficient tends to be higher.

Another useful network structure is a completely regular network, such as a lattice or a ring, in which the arrangement of edges can be construed as following a spatial structure solely made up of local connections. 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. Additionally, in lattice networks nodes that are spatially close tend to be connected to each other, so the clustering coefficient tends to be high.

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 short-cut paths allow the search to proceed quickly (Kleinberg, 2000).

Allen Wilhite (2001) compared market trading over various network structures. In one condition, all agents were allowed to negotiate trade with any other agent. In another, agents could only trade locally, in small cliques. In a third condition, 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 (the state where no more trades that mutually benefit both traders can be made) even faster than the condition where everyone could trade with everyone. Although the agents made more trades in the small-world network, their search space was constrained, minimizing the time it took to find the optimal trading partner. This is further evidence that small-world networks have advantageous features for the dissemination of information.

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 groups passing information over email. Like Friedkin & Johnsen's (1999) work, the influence in these studies flow through a social network. Unlike their work, however, Latané and colleagues' work focuses almost exclusively on judgmental issues, and none of their work uses networks that have been analyzed using traditional network measures such as clustering and path length.

Research on group problem solving dates back to Shaw (1932), but Leavitt (1951) and Bavelas (1950) were some of the first 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 trying to find their own best solution to a problem, and then subsequent individuals imitate good solutions. The group can then be evaluated with respect to how quickly the innovation spreads. As in March's (1991) and Wilhite's (2001) studies, there is a tradeoff with respect to the amount of individual exploration versus the exploitation of good solutions by imitation.

There is excellent work studying the diffusion of innovation in real groups (e.g., Ryan & Gross, 1943; Rogers, 1962, 1995), social psychological research on how individuals use information provided by others (Sherif, 1935; Cialdini & Goldstein, 2004), as well as computational models of information transmission (Nowak, Szamrej, & Latané, 1990; Axelrod, 1997; Kennedy, Eberhart, & Shi, 2001). The studies reported in this paper tie together these diverse areas by exploring the diffusion of innovative ideas among a group of networked participants, each of whom is trying to individually find the best solution to a search problem. This provides a unique and novel method for studying the effect of network structure on group performance with respect to innovation diffusion in different formally defined problem spaces using actual human behavior.

Our Paradigm

In choosing a paradigm for studying information dissemination, we sought to find a case with: 1) a problem to solve with answers that varied continuously on a quantitative measure of quality, 2) a problem search space that was sufficiently large that no individual could cover it all in a reasonable amount of time, and 3) simple communications between participants that would be amenable to computational modeling. We settled upon a minimal search task in which participants guess numbers between 0 – 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 manner, participants could choose to imitate high-scoring guesses from their peers.

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 that 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. In addition to these three network structures, we also used 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 in the different 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 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 normal functions or combinations of normal functions, so that approximations to the best solution earned close to the best payouts. We decided that a person guessing within half a standard deviation of the best solution could be considered close enough 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 they guessed between 37 and 43. Based on 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 guessing close together, and how volatile the group members' guesses were, as an indication of how much participants were exploring. The measures we used and their purposes are listed in Table 2.

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. As previously mentioned, bandwagon behavior has been observed even where there is no objective difference between solutions. In this study we created a bimodal payout function with two equal maxima (see Figure 2) to see if and when participants would 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. Naturally, as group size increases we expect the effects of the network structure to increase, as the topological differences between the networks become greater.

Method

Fifty-six groups of Indiana University undergraduate students ranging in size from 5 – 18 people with a median of 12 people per group participated for partial course credit, for a total of 679 participants. Five groups had to be dropped due to data logging error, but this did not affect the distribution of group sizes. Each session was run in a computer lab with 20 client computers used by the participants and one server operated by the experimenter. Participants signed onto the computer and gave themselves a handle or were assigned an ID. Once all participants had signed onto the computer, 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 8 problems. Each problem consisted of 15 rounds in which participants had 20 seconds each round to guess a number between 0 - 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 variance of 25, and return the feedback. This began the next round, and participants now had available their guess and score 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 seconds, the server created a new network structure and began the next problem. There were two fixed random orderings of 8 payout functions (including the functions used in Study 2) and network structures that were counterbalanced between groups. There were no significant effects of order, so this factor will not be included in any future analyses. The positions of the maxima were different for each of the 8 problems (for information on how the payout functions were generated, see Appendix). 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 Malloy and Reed (1995) for generating random networks with a pre-defined 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 3 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 the other three networks.

Our small-world networks are comparable to those generated in Ahmed and Abdusalam's (2000) study of percolation in networks. Unlike traditional small-world networks (e.g., Watts & Strogatz, 1998) our method for generating the small-world networks caused the clustering coefficient to be low 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 number of neighbors each node has) in the small-world (SD = 0.68) and lattice networks (SD = 0.74) relative to the variance in the random networks (SD = 1.12).

The size of the groups varied depending on how many participants signed up for the various sessions, so we could not strictly control group size. As a result, the distribution of group sizes was very inhomogenous and non-normal, so for our analysis we used quartile splits to break group size into four approximately equally sized categories. In these studies, the unit of analysis is the group, not the participants within the group. Therefore each newly constructed network constitutes a new group. This makes the experiment a 4 (network structure) x 4 (group size) between-subjects (or rather, between-groups) design. All analyses, unless otherwise reported, were full factorial univariate ANOVAs.

Results

We examined several measures of search performance to compare the bandwagoning of different network structures. First, we compared the relative entropies to see how clustered the guesses were. To measure how clustered the guesses were, we used the relative entropy statistic (or Kullback-Leibler[1]) to compare the spread of guesses to a uniform distribution. The less the guesses are uniformly spread across the total possible range of guesses, the higher the relative entropy. Presumably if there were more bandwagoning, the guesses would be more clustered around one of the two global maximum, and thus would have higher relative entropy. All of the network structures had a high relative entropy, although the small-world network was significantly smaller (M = 1.67, SD = 0.57), indicating less clustering in the guesses than the other three networks (Full = 1.84, SD = 0.55; Lattice = 1.85, SD = 0.57; Random = 1.88, SD = 0.64), $F(3, 749) = 8.816$, MSE = 0.121, $p < 0.001$. The variances in entropy between the networks were inhomogeneous ($p < 0.05$), so Games-Howell post-hoc tests were used. They confirmed that the small-world networks were significantly less clustered than the other three networks ($p < 0.001$), and there were no significant differences between the other three. Interestingly, as group size increased, the clustering decreased significantly ($F(3,749) = 99.674$, $p < 0.001$), indicating less bandwagoning. The interaction with network type was also significant ($F(9,749) = 6.978$, $p < 0.001$), as the amount of clustering stayed relatively high in the fully-connected network regardless of group size.

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 more exploration, so as bandwagoning increases, volatility is expected to decrease. Supporting the analysis of relative entropy, the small world network had the highest volatility (M = 5.85, SD = 5.48) compared to the other three networks (Full = 5.22, SD = 5.90; Lattice = 4.15, SD = 4.26; Random = 4.40, SD = 4.68), and the difference between networks was significant, $F(3, 749) = 5.172$, MSE = 26.037, $p < 0.005$. The only significant pair-wise comparison revealed by post-hoc tests was between the small-world and lattice networks, $p < 0.001$. As group size increased, volatility decreased, indicating less exploration. This main effect was significant ($F(3, 749) = 3.550$, $p < 0.05$), and so was the interaction with network type, $F(9,749) = 2.740$, $p < 0.005$. The fully-connected networks saw a much greater decrease in volatility as group size increased than did the other three networks.

Another good measure of bandwagoning is the difference in the number of participants within one-half standard deviation of each of the maxima. This difference was normalized by the percentage of people guessing within either of the peaks, because it was significantly different across networks, $F(3,761) = 4.141$, $MSE = 0.094$, $p < 0.01$. 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. As might be expected, the fully connected network had a much higher degree of bandwagoning ($M= 0.76$, $SD = 0.29$) than the other networks (Lattice = 0.59, $SD = 0.34$; Random = 0.56, $SD = 0.38$; Small-World = 0.51, $SD = 0.33$), and the differences between networks was significant, $F(3, 746) = 14.413$, $p < 0.001$. Post-hoc tests revealed that only the comparisons between the fully-connected networks and the other networks were significant ($p < 0.001$). In support of the anlysis of relative entropy, the proportion of bandwagoning decreased significantly as group size increased ($F(3,746) = 22.439$, $p < 0.001$), and the interaction with network type was also significant ($F(9,746) = 7.712$, $p < 0.001$), again because the fully-connected network did not show much of a decrease in bandwagoning as the group size increased.

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 only received feedback from 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 full network, when people converged rapidly, this meant a steady pull toward whatever solution was found first. For the small-world network, the highly regular spatial structure and short path lengths more often led to at least one subgroup 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 because the occasional higher score from another subgroup could quickly pass to another subgroup.

As group size increased, the amount of bandwagoning in the lattice, small-world, and random networks decreased. This is likely a result of the greater possibility for two clusters to form as more participants are guessing. It is possible that if there were more rounds of guessing, these networks would also end up converging on a single maximum. For the fully connected network, however, the informational influence was greater (as there was more information for each participant), so group size had little to no effect on the amount of bandwagoning behavior observed.

Study 2

In this experiment, 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 multi-modal function increased the difficulty of the search and introduced 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 multi-modal functions used, which has three peaks, but one of the peaks is somewhat higher than the other two.

The basic prediction is that the 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, so those networks that have the fastest dissemination of information will perform best. These are the ones with the shortest path length, and are the ones that would show the most bandwagoning – the fully-connected network, followed by the small-world and random networks. However, in the multimodal problem space, bandwagoning behavior 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 the bandwagoning behavior observed in the fully-connected networks in Study 1.

Method

The procedure was the same as study 1, only using the unimodal and multimodal problem spaces. These were included in the two fixed random orders of problems spaces described in Study 1, so the participants in this study were the same as those in Study 1. As mentioned previously, there were no significant order effects. Three groups in the unimodal problem space and two groups in the multimodal problem space had to be dropped due to data logging error, but this did not affect the distribution of group sizes. Again, the distribution of group sizes was very inhomogenous and non-normal, so we used quartile splits to break group size into four approximately equally sized categories. As in Study 1, the experiment was a 4 (network structure) x 4 (group size) between-subjects design. All analyses were full factorial univariate ANOVAs unless otherwise reported.

Results

Unimodal

In this section we analyze the groups' performances in the unimodal payout function. A full factorial between-groups ANOVA on each group's speed of convergence in the unimodal function revealed a significant difference between the network types, $F(3, 37) = 3.405$, $MSE = 3.341$, $p < 0.05$. The fastest convergence occurred in the fully-connected network, in which participants on average took 2.95 rounds ($SD = 0.93$) to reach the maximum. The second fastest was the small-world network ($M = 4.0$, $SD = 1.46$), followed by the lattice network ($M = 4.92$, $SD = 2.19$) and the random network ($M = 5.21$, $SD = 2.07$). The variances of the rate of convergence for the networks were significantly inhomogeneous ($p < 0.05$), so a Games-Howell post-hoc comparison was used, which revealed that the fully-connected network was significantly faster than the random network ($p < 0.001$) and marginally faster than the lattice network ($p < 0.1$). No other comparisons were significant. There were no significant effects of, or interactions with, group size.

An analysis of the percent of participants in each group guessing within the global maximum over all 15 rounds showed a significant main effect for network type, $F(3,779) = 38.284$, $MSE = 0.083$, $p < 0.001$. Averaged over all groups and rounds, the fully-connected networks had 73.31 percent of participants guessing in the maximum ($SD = 20.94$), compared to 68.78 ($SD = 26.7$) for the small-world network, 54.89 ($SD = 29.53$) for the lattice network, and 45.04 ($SD = 35.29$) for the random network. Post-hoc comparisons using the Games-Howell test revealed that the fully-connected and small-world networks had significantly more convergence than the lattice and random networks ($p < 0.001$), and the lattice networks had significantly more than the random networks ($p < 0.001$). Group size also had a significant main effect, such that larger groups had a higher proportion of participants guessing in the maximum ($F(3, 779) = 10.729$, $p < 0.001$). There was also a significant interaction ($F(9, 779) = 3.933$, $p < 0.001$), such that the lattice and random networks show a stronger effect of group size, with low convergence in the smaller group sizes and more convergence in the larger networks.

A helpful visualization of these results is the average proportion of participants guessing within the maximum over the 15 rounds. As can be seen in Figure 5a, the fully connected networks quickly find the maximum and roughly 80 percent of group members continue to guess in the maximum for the remaining rounds. The small-world network does not find the maximum as quickly as the fully-connected network, but attains the same amount of convergence. The lattice and random networks converge more slowly on the maximum and never reach this level of convergence, with the lattice having at most 71.86 percent of participants guessing in the maximum and the random having at most 53.21 percent.

The lattice network had a low average relative entropy ($M = 1.43$, $SD = 0.42$) compared to all of the other networks (full: $M = 1.71$, $SD = 0.36$; small: $M = 1.61$, $SD = 0.45$; random: $M = 1.71$, $SD = 0.41$) indicating that the distribution of guesses in the lattice network was typically less clustered than for the other networks ($F(3, 791) = 17.259$, $MSE = 0.169$, $p < 0.001$). Additionally, the lattice network had the highest volatility, with an average difference in guesses between rounds of 6.59 ($SD = 7.78$), while the fully-connected networks had the lowest volatility ($M = 3.71$, $SD = 3.71$). Volatility was significantly different between networks ($F(3, 791) = 7.367$, $MSE = 33.642$, $p < 0.001$). Games-Howell post hoc tests revealed that the lattice had significantly more volatility than the fully connected networks ($p < 0.01$) and the random networks ($M = 4.71$, $SD = 5.41$), $p < 0.05$. The small-world networks ($M = 5.64$, $SD = 6.12$) also had significantly more volatility than the fully connected networks, $p < 0.01$.

Multimodal

In the multimodal landscape we again expect shorter path lengths to correspond with faster convergence on the global maximum, but we anticipate that lack of spatial structure could lead to less exploration, and thus early convergence on a local maxima and a slower convergence on the global maximum. As predicted, the average number of steps for the first person to reach the global maximum was less in the small-world network ($M = 5.07$, $SD = 1.43$) than even the fully-connected network (M =

6.12, SD = 2.09). The main effect of network type was significant ($F(3,38) = 4.787$, $MSE = 5.292$, $p < 0.01$), and while Games-Howell post-hoc tests did not show that the difference between the small-world and fully-connect networks was significant, only the small-world networks were significantly faster than the lattice (M = 8.77, SD = 3.75; $p < 0.05$) and the random (M = 7.92, SD = 2.53; $p < 0.005$) networks. There was a main effect for group size ($F(3,38) = 4.207$, $p < 0.05$), and post hoc analyses show that the largest groups found the global maximum marginally faster than the smallest groups, $p < 0.07$. There was no significant interaction between network type and group size.

The small-world network also had the greatest convergence on the global maximum across all rounds (M = 0.519, SD = 0.27), again closely followed by the fully-connected network (M = 0.481, SD = 0.3), with the lattice (M = 0.341, SD = 0.3) and the random networks (M = 0.29, SD = 0.26) showing very little convergence on the global maximum. The differences between the networks were significant, $F(3, 794) = 28.239$, $MSE = 0.068$, $p < 0.001$, and post-hoc tests showed that the fully-connected and small-world networks differed significantly from the lattice and random networks ($p < 0.001$) and not from each other. Group size also had a significant main effect ($F(3,794) = 24.656$, $p < 0.001$), with increasing group size leading to greater overall convergence on the global maximum. The interaction between group size and network was also significant ($F(9,794) = 4.246$, $p < 0.001$), although the effect appears to mostly driven by the very large difference in convergence between the small sized lattice networks (M = 0.1, SD = 0.15) and the larger sized lattice networks (M = 0.48, SD = 0.29).

An examination of the percentage of participants within the global maximum on each round highlights the advantage of the small-world network. As can be seen in Figure 5b, the small-world network consistently dominates the other network structures until round 11, when the network with full information finally catches up, while the other networks never reach the small-world networks' convergence. The hypothesis that the fully connected networks prematurely converge on the local maximum is supported by comparing the percentage of group members guessing in either of the two local maxima. This was significantly different between network types ($F(3,806) = 35.75$, $MSE = 0.055$, $p < 0.001$), and post-hoc tests revealed that the fully-connected networks did indeed have significantly more participants guessing in the local maxima than the small-world networks, $p < 0.001$.

Unlike the unimodal payout function, with the multimodal function the fully-connected networks had the highest volatility ($M = 6.1$, $SD = 5.49$), and the difference between network types was significant ($F(3, 806) = 2.66$, $MSE = 25.201$, $p < 0.05$. Since the fully-connected networks did not tend to find the global maximum early, this high volatility could be due to fluctuations of exploration and convergence. Similar to the unimodal function, the lattice network had the lowest entropy ($M = 1.28$, $SD = 0.41$), indicating the least amount of clustering in the distribution of guesses. The difference between networks was significant ($F(3, 806) = 3.06$, $MSE = 0.167$, $p < 0.05$), and Games-Howell post-hoc tests confirmed the only significant pairwise comparison was the lattice network to the fully-connected network ($M = 1.39$, $SD = 0.41$), $p < 0.05$.

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 faster than 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 on average. Unexpectedly, the random networks found the maximum the slowest on average, and had the least amount of convergence overall. Despite the short average path lengths in random networks, the fact that information was traveling randomly through the network could have caused more instances of conflicting information reaching participants simultaneously, causing uncertainty in the decision making process. In contrast, the other networks had a spatial structure that is more conducive to a systematic search of a problem space.

Group size generally had the expected effects. As groups became larger, more participants were available to search the problem space, so convergence on the maxima was faster and the proportion converging was higher as group sizes increased. The expected interaction, such that greater group sizes led to increased differences between the network types, was not consistently found, and the differences found typically occurred in the lattice networks. This makes sense, as the biggest difference in topological measures such as average path length for different sized networks will be in lattice networks. It is possible that the interaction between network type and group size would be evident with much larger groups.

When the problem space had good solutions that were nonetheless sub-optimal, as with the multimodal payout function, the story was different. In this case the small-world network groups found the best solution faster and converged on the global maximum more than every other network, even the fully connected network in which everyone had complete information about every other participants' guesses and scores. We expect that the decision making processes of the individuals in each group would not differ between the network types, as the information presented to them was essentially the same. Therefore, differences between the network types must solely be due to the information transmission properties of the networks. The advantage of the small-world over the fully connected networks is like a novel group-based form of the "less-is-more" effect reported in individual decision making literature (Gigerenzer & Todd, 1999).

This somewhat counter-intuitive 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 only the best solution one third of the time. When the group converged prematurely on a local maximum, it took them 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 into different spatial regions, but the information could travel quickly through “short-cuts,” 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 decreases the probability that everyone will follow their lead before another sub-group finds 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 rapid convergence on an easy-to-find but suboptimal solution. Like the multimodal payout function, networks with a well-defined spatial structure will allow continued exploration. However, with a very 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 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

Forty-eight groups of Indiana University undergraduate students ranging in size from 7 – 19 people with a median of 12.5 people per group participated for partial course credit, for a total of 628 participants. The procedure was the same as Study 2, only using the needle payout function instead of the unimodal or multimodal functions.

The distribution of group sizes was again very inhomogenous and non-normal, so we used quartile splits to break group size into four approximately equally sized categories. As in Study 1 and 2, the experiment was a 4 (network structure) x 4 (group size) between-subjects design. All analyses were full factorial univariate ANOVAs unless otherwise reported.

Results

In the needle payout function, the global maximum has a small range, so the longer participants explore, the more likely one of them is to find it. Thus we predicted that the more spatially segregated networks would be more likely to find the "needle." There were no significant differences in the average number of steps it took for a participant to guess the global maximum (i.e., the "needle"). Nonetheless, the average fraction of participants in the global maximum over all rounds was significantly higher with the lattice network (M = 0.241, SD = 0.346) than the other three network structures (Full = 0.155 (0.267); Small-world = 0.118 (0.203); Random = 0.114 (0.243), $F(3,704) = 8.027$, MSE = 0.06, $p < 0.005$. A Games-Howell post-hoc analysis confirms this, as the lattice network was significantly greater than the fully-connected ($p < 0.05$) and the other two networks ($p < 0.001$). As can be seen in Figure 7a, the average number of participants in the lattice networks guessing in the global maximum increases at a faster rate than any of the other network structures. This is in contrast to the average fraction of participants that guessed within the local maximum. After round 5, the fraction of participants in the local maximum drops for the lattice network but plateaus or continues to increase for the other network structures (see Figure 7b).

Group size also had a significant main effect ($F(3,704) = 16.14$, $p < 0.001$) on the proportion guessing in the needle, and this had an interesting inverse U-shaped curve, such that the smallest and largest group sizes did not typically find the needle, while the middle-sized groups did. This could be because the smallest sized groups were unable to find the maximum at all while the largest groups did less exploring. This is supported by the fact that while the type of network did not have an effect on the average volatility over all rounds, group size did ($F(3,716) = 3.418$, MSE = 30.544, $p < 0.05$). Games-

Howell post-hoc tests confirmed that the largest groups had significantly less volatility than the two smallest group sizes, $p < 0.05$. There was also an interaction effect of group size x network for overall convergence, ($F(9, 704) = 5.047$, $p < 0.001$). The small-world and random networks never found the needle in the smallest and largest groups, and the lattice networks had the inverse U-shaped curve, but the fully-connected network converged on the needle in the largest groups more than the lattice networks. The advantage of large, fully-connected networks is likely a result of greater sampling of the problem space at the outset.

The Kullback-Leibler measure of relative entropy supports the hypothesis that increased exploration improved the groups' ability to find the needle. As before, we divided the range of guesses into 20 "bins." In this case the four network types differed significantly from one another, $F(3, 716) = 34.198$, $MSE = 0.259$, $p < 0.001$. As can be seen from Figure 8, the full and lattice networks were both higher than the small-world and random networks, and post-hoc tests revealed that the full and lattice networks did not differ from each other and both differed significantly from the small-world and random networks ($p < 0.001$). However, as can be seen by the difference in where the groups were guessing, the participants in the lattice network tended to cluster in the range of the "needle," while the full network typically clustered in the local maximum.

Discussion

The payout function used in Study 3 is meant to represent situations in which a problem has a precise best solution that is not easily approximated, and a lesser solution that is easy to roughly imitate. In our study, the group of people connected in a regular, lattice network were best able to find the optimal solution. Given that there were no significant differences in how quickly participants found the needle on average, it seems that the advantage the lattice had over the other networks was based on how often the needle was found. This is most likely due to the increased exploration engendered by the long path lengths and local, spatial neighborhoods preserved in the lattice networks.

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 showed 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 is no objectively correct solution. This bandwagoning behavior occurs even when normative influence is minimized. This effect was moderated by the network types, as 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 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 group-wide solution and therefore will continue to have group members exploring the problem space. In our studies, the lattice network was most likely to converge on the needle and therefore outperformed the other networks with this payout function.

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 similar solutions, but only one was globally the best. In this case, very 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 was found by the small-world network which possesses both preserved spatial neighborhoods and long-distances connections.

From our results so far, 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 the 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,. Since there is a tradeoff between the exploration of a problem space and the 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.

Recently, 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 0, 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, 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 neighbor's score was better, or mutated their string by a digit if the mutation resulted 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 faster than the lattice networks. However, 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 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 of the range of any of these sources of influence, we can say the participant is exploring. Unfortunately, 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 between 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 vs. imitation. Broadly speaking, these results have implications for many different areas of study. For research on group performance and organizational psychology, this highlights the importance of the communication patterns within a group with respect to the type of problem being approached by the group. 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 could be tested, such as scale-free networks (which are increasingly observed in a wide range of real networks; Barabasi & Albert, 1999), or hierarchies, which are interesting because they are a typical organizational structure. Additionally, different problem spaces remain to be explored, including multidimensional and dynamically evolving problem spaces. 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). The diameter of the world wide web. *Nature, 401*, 130-131.

Asch, S. E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs, 70(9).*

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.

Barabási, A.-L., & Albert, R. (1999). 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 patternsi n task-oriented groups. *Journal of the Acoustical Society of America, 22(6*), 725-730.

Blackmore, S. J. (1999). *The meme machine.* Oxford: Oxford University Press.

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.

Erdös, P. & Rényi, A. (1959). On random graphs. *Publicationes Mathematicae, 6*, 290-297.

Friedkin, N. E. & Johnsen, E. C. (1999). Social influence networks and opinion change. *Advances in Group Processes, 16*, 1-29.

Gergely, G. Bekkering, H. & Király, I. (2002). Rational imitation in preverbal infants. *Nature, 415*, 755.

Gladwell, M. (2000). *The tipping point.* London: Little Brown and Company.

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: University of Michigan Press.

Jeong, H., Tombor, B., Albert, R., Oltvai, Z. N., & Barabási, A.-L. (2000). 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 Publishers.

Kleinberg, J. (2000). Navigation in a small world. *Nature, 406*, 845.

Latané, B. & Bourgeois, M. J. (1996). Dynamic social impact in electronic groups: Consolidation and clustering in computer-mediated communication. *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 intellective tasks. *Journal of Experimental Social Psychology, 22*, 177-189.

Lazer, D. & Friedman, A. (2005). The parable of the hare and the tortoise: small worlds, diversity, and system performance. KSG Working Paper No. RWP05-058. Available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=832627.

Leavitt, H. J. (1951). Some effects of certain communication patterns on group performance. *Journal of Abnormal Psychology, 46*, 38-50.

Malloy, M. & Reed, B. (1995). A critical point for random graphs with a given degree sequence. *Random Structures & Algorithms, 6*, 161-179.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science, 2*, 71-87.

Melztoff, A. N. (1988), Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology, 24*, 470-476.

Milgram, S. (1967). The small world problem. *Psychology Today, 2*, 60-67.

Neumann, R. & Strack, F. (2000). "Mood contagion"" The automatic transfer of mood between persons. *Journal of Personality and Social Psychology, 79,* 211-223.

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.

Ryan, B. & Gross, N. C. (1943). The diffusion of hybrid seed corn in two Iowa communities. *Rural Sociology, 8*, 15-24.

Shaw, M. E. (1932). Comparison of individuals and small groups in the rational solution of complex problems. *American Journal of Psychology*, 44, 491-504.

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.

Valente, T. W., & Davis, R. L. (1999). Accelerating the diffusion of innovations using opinion leaders. *Annals of the American Academy of Political and Social Science, 566*, 55-67.

Watts, D. J. & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature, 393*, 440-442.

Wejnert, B. (2002). Integrating models of diffusion of innovations: A conceptual framework. *Annual Review of Sociology, 28*, 297-326.

Wilhite, A. (2001). Bilateral trade and ‘small-world’ networks. *Computational Economics, 18*, 49-64.

Author Note

Winter A. Mason, Psychological and Brain Sciences, Indiana University; Andy Jones, Department of Computer Science, University of North Carolina; Robert L. Goldstone, Psychological and Brain Sciences and program in cognitive science, Indiana University.

This research was funded by Department of Education, Institute of Education Sciences grant R305H050116 and NSF REC grant 0527920 to the second author. 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 on-line 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 Mason, Psychological and Brain Sciences, Indiana University, Bloomington, IN, 47405. Email: wimason@indiana.edu.

Footnotes

[1]The Kullback-Liebler is

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

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 – 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.

Table 1

Average Geodesic Path Length and Clustering Coefficient for the Networks Used in Study 1

Network	Path Length	Clustering
Full	1.00	1.00
Lattice	3.08	0.36
Small-World	2.61	0.09
Random	2.57	0.37

Table 2

Dependent measures used and meaning

Dependent measure	Meaning
Average steps to guess in global maximum	Average speed of finding best solution
Average percent 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 Caption

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

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

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

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

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

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

Figure 7. Study 3: Percent of participants within ½ SD of a) the global maximum (the "needle") and b) the local maximum

Figure 8. Study 3: Relative entropy (Kullback-Liebler) of the participants' guesses for the different network types

QuickTime™ and a
TIFF (LZW) decompressor
are needed to see this picture.

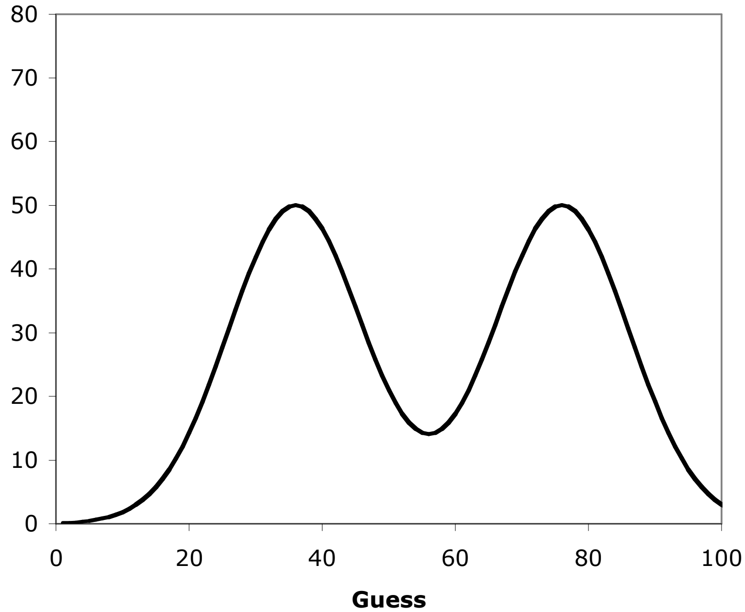
80
70
60
50
40
30
20
10
0
0
20
40
60
80
100
Guess

QuickTime™ and a
TIFF (LZW) decompressor
are needed to see this picture.

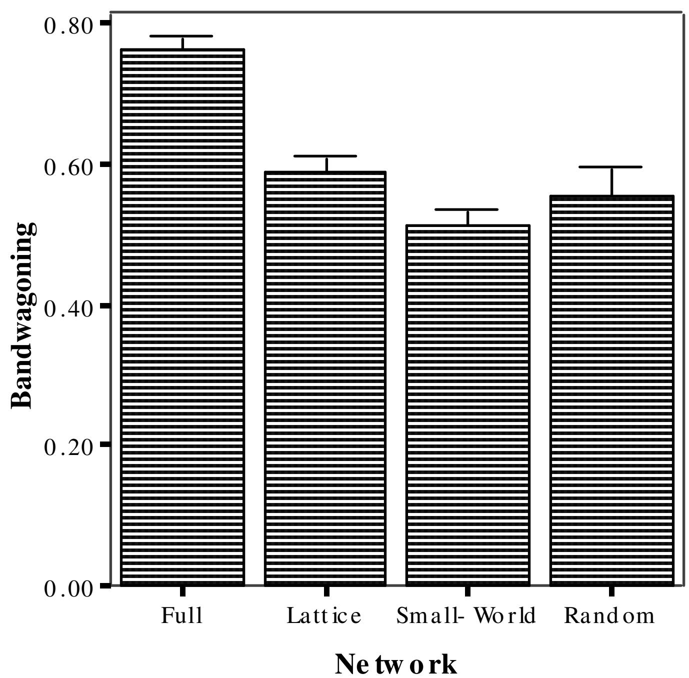
0.80
0.60
0.40
0.20
0.00
Bandwagoning
Full
Lattice
Small-World
Random
Network

a)

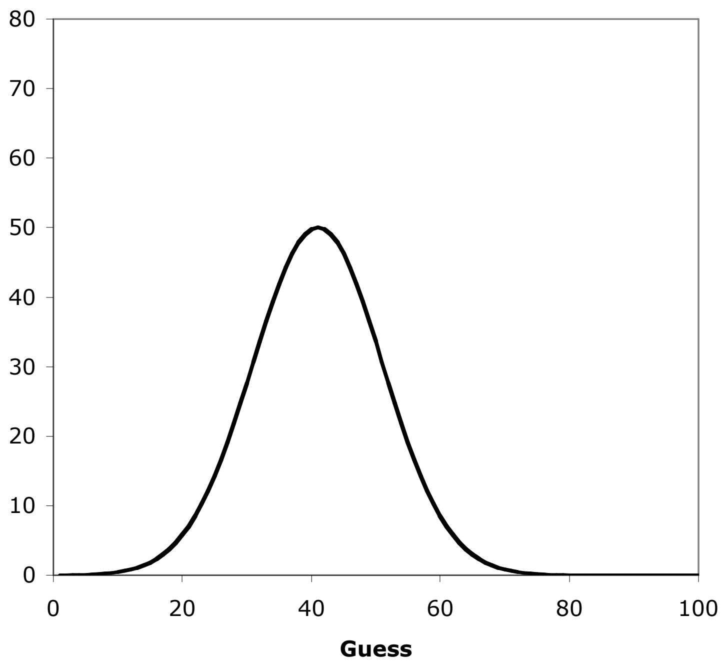
80
70
60
50
40
30
20
10
0
0
20
40
60
80
100
Guess

b)

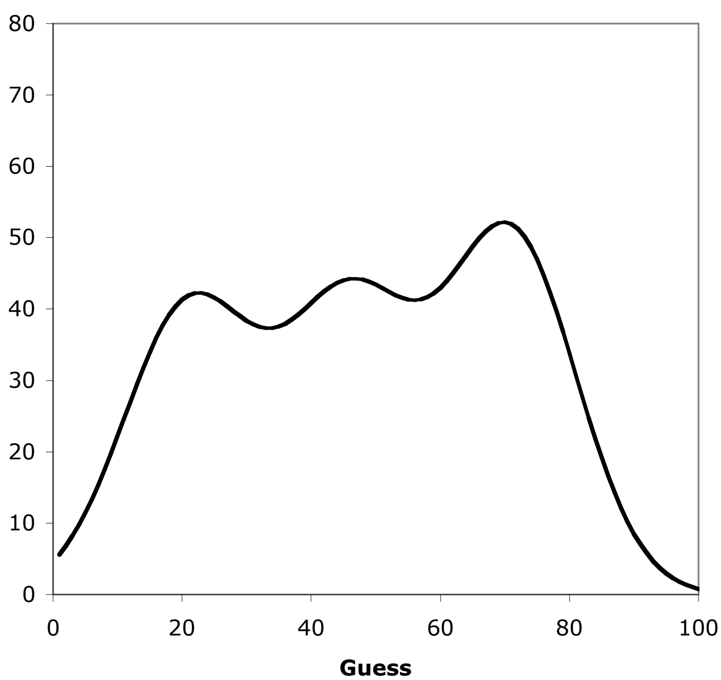
80
70
60
50
40
30
20
10
0
0
20
40
60
80
100
Guess

a)

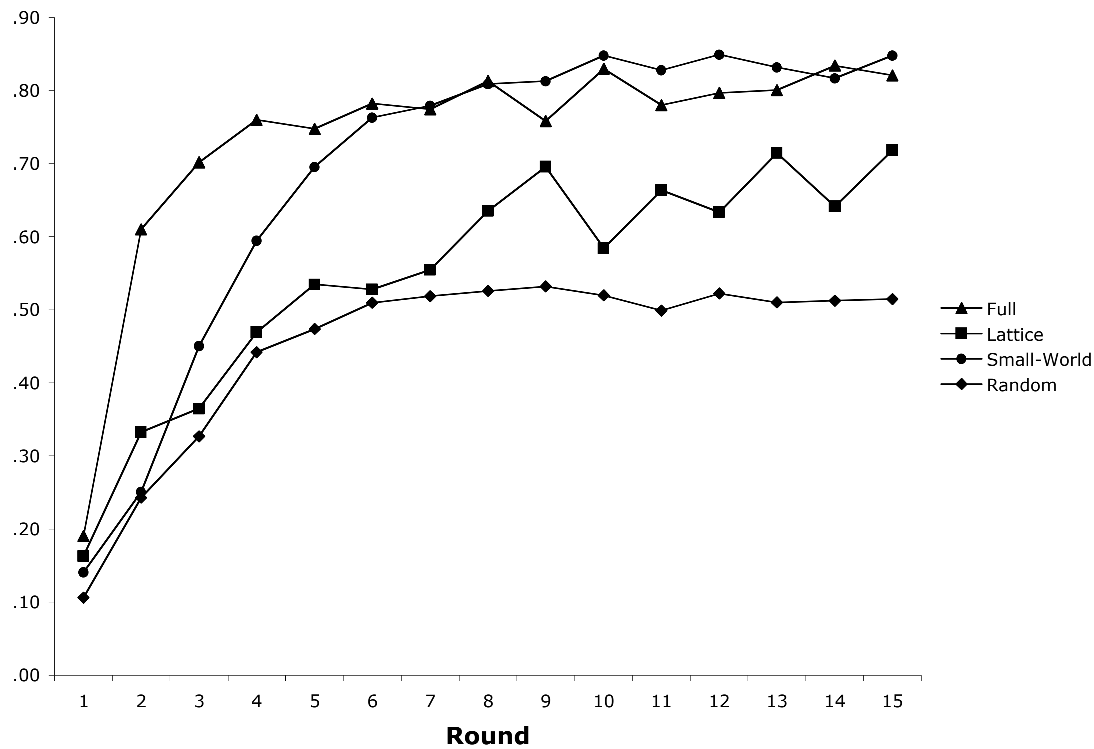
.90
.80
.70
.60
.50
.40
.30
.20
.10
.00
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
Full
Lattice
Small-World
Random

b)

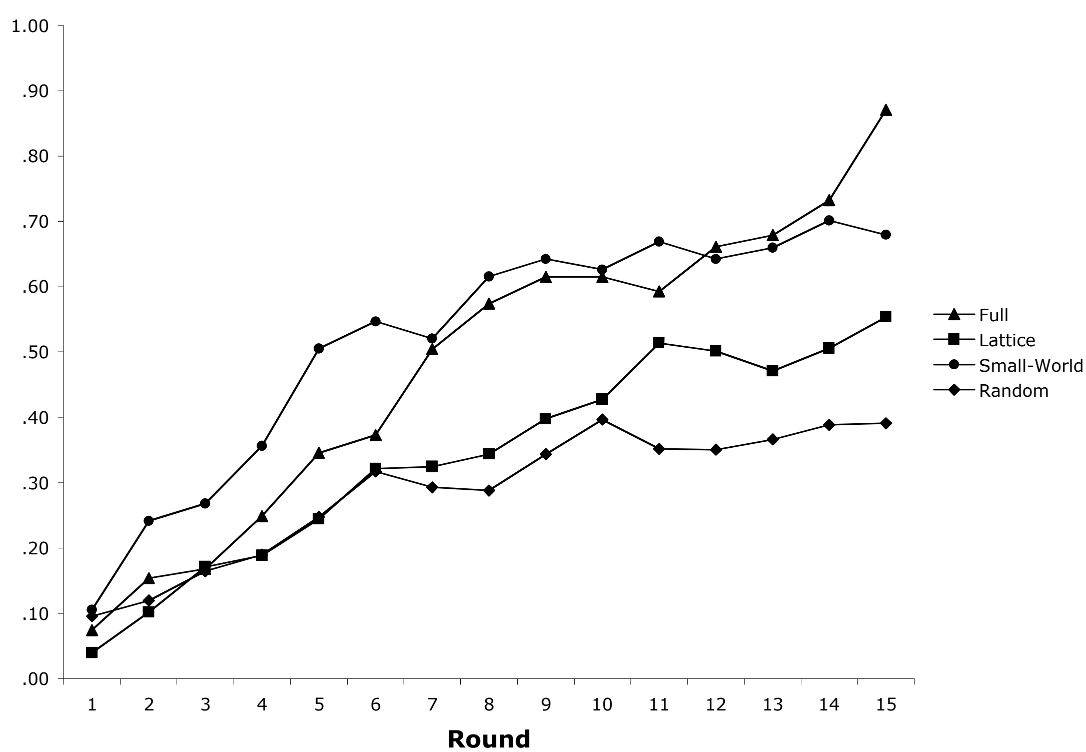
1.00
.90
.80
.70
.60
.50
.40
.30
.20
.10
.00
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
Full
Lattice
Small-World
Random

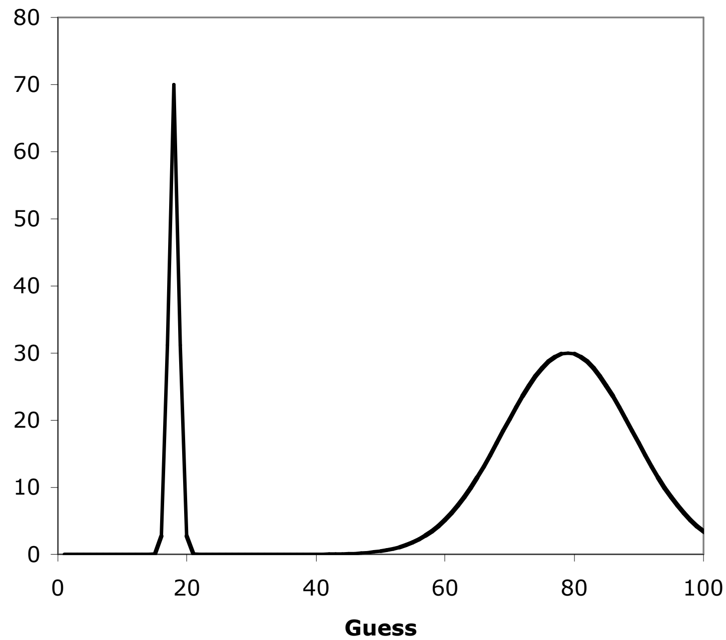
80
70
60
50
40
30
20
10
0
0
20
40
60
80
100
Guess

a)

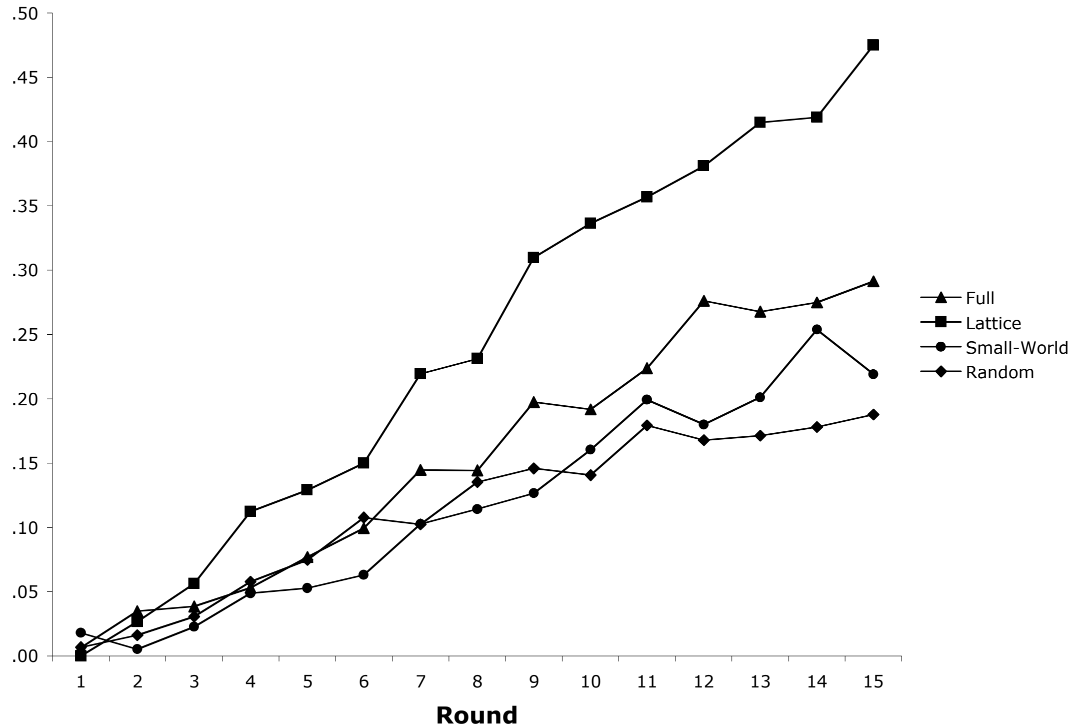
.50
.45
.40
.35
.30
.25
.20
.15
.10
.05
.00
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
Round
Full
Lattice
Small-World
Random

b)

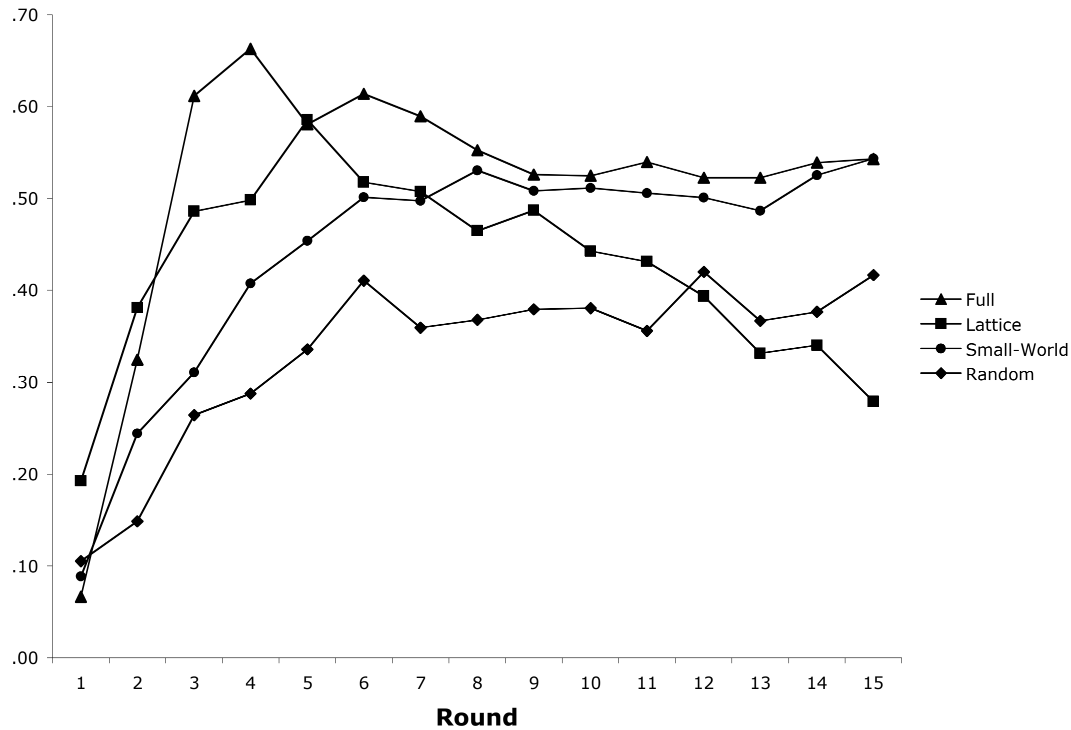
.70
.60
.50
.40
.30
.20
.10
.00
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
Round
Full
Lattice
Small-World
Random

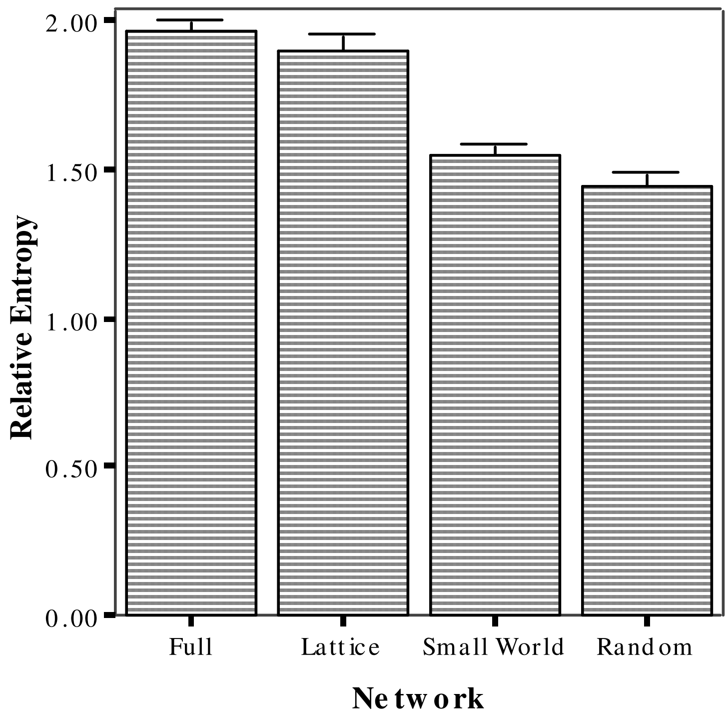
2.00
1.50
1.00
0.50
0.00
Relative Entropy
Full
Lattice
Small World
Random
Network

Appendix

Functions Used to Transform Participants' Guesses Into Corresponding Scores

Payout Function	Network	Equation
Unimodal	Full	$50e^{-(0.07(x-40))^2}$
	Lattice	$50e^{-(0.07(x-65))^2}$
	Small-World	$50e^{-(0.07(x-15))^2}$
	Random	$50e^{-(0.07(x-15))^2}$
Multimodal	Full	$40e^{-(0.07(x-15))^2} + 40e^{-(0.07(x-40))^2} + 50e^{-(0.07(x-70))^2}$
	Lattice	$40e^{-(0.07(x-20))^2} + 40e^{-(0.07(x-45))^2} + 50e^{-(0.07(x-70))^2}$
	Small-World	$40e^{-(0.07(x-35))^2} + 50e^{-(0.07(x-60))^2} + 40e^{-(0.07(x-85))^2}$
	Random	$50e^{-(0.07(x-30))^2} + 40e^{-(0.07(x-55))^2} + 40e^{-(0.07(x-80))^2}$
Needle	Full	$70e^{-(0.9(x-32))^2} + 30e^{-(0.07(x-83))^2}$
	Lattice	$30e^{-(0.07(x-26))^2} + 70e^{-(0.9(x-67))^2}$
	Small-World	$70e^{-(0.9(x-17))^2} + 30e^{-(0.07(x-78))^2}$
	Random	$30e^{-(0.07(x-38))^2} + 70e^{-(0.9(x-72))^2}$
Equal	Full	$50e^{-(0.07(x-15))^2} + 50e^{-(0.07(x-55))^2}$
	Lattice	$50e^{-(0.07(x-25))^2} + 50e^{-(0.07(x-65))^2}$
	Small-World	$50e^{-(0.07(x-45))^2} + 50e^{-(0.07(x-85))^2}$
	Random	$50e^{-(0.07(x-35))^2} + 50e^{-(0.07(x-75))^2}$