

# Facts and Figuring: an Experimental Investigation of Network Structure and Performance in Information and Solution Spaces

JESSE SHORE,<sup>a†</sup> ETHAN BERNSTEIN<sup>b</sup>, AND DAVID LAZER<sup>c</sup>

<sup>†</sup>[jccs@bu.edu](mailto:jccs@bu.edu)

<sup>a</sup>Boston University School of Management

<sup>b</sup>Harvard Business School

<sup>c</sup>Northeastern University

July 22, 2013

## Abstract

Drawing on results from a large, novel laboratory experiment, we examine the impact of organizational communication structure on performance in complex problem-solving tasks. Leaders of organizations are increasingly paying attention to the design of communication structures to optimize performance in complex problem-solving tasks. Prior literature, however, has produced contradictory answers with respect to the benefits of one of the key variables: clustering of network ties. We reconcile these contradictions theoretically by separating searching for information from searching for solutions in the process of problem solving, and experimentally by measuring information separately from solutions. We show that clustering promotes exploration in information space, but decreases exploration in solution space. In other words, the role of network structure is different for raw information and interpretations of that information.

## I. INTRODUCTION

**K**nowledge-centered organizations, by definition, leverage contributions of many individuals to solve problems too complex for a single individual to tackle alone. In such organizations, problem solving occurs not at the individual level but instead at the organizational level, expertise is the core organizational asset, and superior answers (rather than, for example, superior production) define organizational performance. The proliferation of knowledge-centered organizations has produced a managerial mandate to understand the conditions under which

successful integration of intra-organizational knowledge will best enable advanced problem solving and higher performance. However, although much has been written about knowledge and problem solving, our understanding of problem solving in knowledge-centered organizations remains limited.

A number of variables may affect such performance in knowledge-centered organizations, including capabilities in knowledge management (e.g., Alavi and Leidner, 2001; Nordenflycht, 2010); the capacity for knowledge creation (e.g., von Krogh, Nonaka, and Rechsteiner, 2012), knowledge retention (e.g., Edmondson, Pisano, Bohmer, and Winslow, 2003; Marsh and Stock, 2006), knowledge transfer (e.g., Argote and Ingram, 2000; Reagans and McEvily, 2003; Levine and Prietula, 2012;) and organizational learning more broadly (Leavitt and March, 1988; Argote, 1999; Argote, McEvily, and Reagans, 2003); the location and design of knowledge boundaries (Carlile, 2004); the ability to learn from failure (Madsen and Desai, 2010; Edmondson, 2011) and success (KC, Staats, & Gino 2013); the degree of team familiarity (Huckman, Staats, and Upton, 2009); and the presence of dynamic capabilities (Teece, Pisano, and Shuen, 1997; Zollo and Winter, 2002) and ambidexterity (Tushman and O'Reilly, 1996; Raisch and Birkinshaw, 2008). The focus, and contribution, of this work relates to a structural question which spans across many of those variables: for a particular decision, what communication structure should an organization construct to optimize organizational performance in complex problem-solving tasks? This question has only become more important with increases in pervasive technology-mediated communication (e.g., Bailey et al, 2010), permitting organizations to directly manipulate the structure of their internal communication networks or external crowdsourcing platforms to pursue different performance outcomes. Opportunities for improvement abound, waiting for theory to guide the way.

Unfortunately, a review of the connection between communication structure and performance in knowledge-centered organizations reveals incongruent, seemingly contradictory answers for a basic lever to manage communication structure: dense clustering of network ties. Empirical field research and case studies in the literature on network structure and problem solving (e.g., Tushman, 1979; Baldwin, Bedell and Johnson, 1997; Sparrowe et al, 2001; Reagans and Zuckerman, 2001; Reagans and McEvily, 2003; Cummings and Cross, 2003) suggest that dense clusters of ties without structural holes are associated with success in complex tasks. In other words, the more sociometrically 'cohesive' we are within the bounds of some subset of an organization, the better we can combine our individual knowledge to collectively solve complex problems. On the other hand, an emerging literature on networked communication, primarily based on lab research and complemented by several simulation studies, draws on theories of ambidexterity (e.g., Raisch and Birkinshaw, 2008; Raisch et al, 2009) and the tradeoff between exploration and exploitation (March, 1991) to conclude that the best problem-solving network structures are

neither globally dense nor locally clustered, because these characteristics are associated with too much copying of each other's solutions, prematurely stamping out beneficial diversity (Mason, Jones, and Goldstone, 2005; Lazer and Friedman, 2007; Mason and Watts, 2012).

In each study, authors grapple with the same larger theoretical question as we do here: how do we structurally connect knowledge workers within teams, groups, and organizations to most effectively solve problems? Rather than the self-reinforcing, cumulative growth of a body of understanding characteristic of academic theory building (Kuhn, 1962), the state of the literature on this question is more reminiscent of a constellation of findings, some of which are contradictory.

In this work, we adopt a hypothesis that the fragmentation is more methodological than theoretical (Kuhn & Jackson, 2008), and that the design of prior studies in the lab, field, and simulation environments has unintentionally placed theoretical coherence farther out of reach. Empirical field work on communication structure is typically based on self-reported ties in sociometric surveys, providing external validity but making it hard to verify exactly what is being measured and which way or ways the causal arrows point. For example, it is possible that an expressed sense of solidarity with another individual or individuals (i.e., affective ties) are a result of problem solving success, rather than the other way around. Meanwhile, while lab experiments and simulations do not suffer from the ambiguities of the field, they trade the verisimilitude of complex problem-solving situated in real-world organizations for statistical identification of causality. Moreover, within the experimental literature, "knowledge" has been operationalized in very different ways. Past experimental and simulation studies have included either a single, self-evident solution that must be achieved through recombination of diverse information already present in the network (Bavelas, 1950; Leavitt, 1951; Guetzkow and Simon, 1955; Kearns, Judd, Tan, Wortman, 2008; Judd, Kearns and Vorobeychick, 2010), or a landscape of solutions of varied quality but no other exchange of information about the problem than subjects' observation of each-other's solutions (Lazer and Friedman, 2007; Mason, Jones, and Goldstone, 2005; Mason and Watts, 2012). Additionally, even in prior work that did involve information recombination, the effective information space was very small, such as a neighbor's choice of color from a set of 3 possibilities (as in Judd, Kearns and Vorobeychick, 2010).

While such methodological tradeoffs exist in many fields, we argue that it has been particularly problematic in this one. Problem solving requires both the act of searching for information, or the facts that may be important pieces of the puzzle, and the act of searching for solutions, or the theories that combine puzzle pieces into an answer. Effective problem solving in knowledge-centered organizations requires both (Cabrera and Cabrera, 2002), but neither simple lab experiments nor self-reported field surveys have effectively captured the effects, individually or jointly, of both.

We hypothesize that these differences in domain — whether previous studies instrumented a search of information space, solution space or, as in the real world, both — have been responsible for the inconsistent findings on the connection between structure and problem solving. We therefore adopt a novel, data-rich experimental platform with greater verisimilitude than any in the past, which was specifically designed to include search and sharing of both information and solutions.

In adopting a richer experimental platform, we are able to reconcile the seemingly contradictory findings of prior work on network structure and problem solving. We find that clustering inhibits exploration through solution space, but it promotes exploration through information space. Through the active communication of information, individuals in a connected cluster tend to be in possession of the same knowledge and be aware of each other’s theories. On the one hand, they therefore do not tend to search for information that another individual in the cluster has already found, and the cluster collectively explores information space more efficiently, and extensively, than they would have if they were not all connected to one another. On the other hand, because they are also all aware of each other’s theories, there is an added tendency to interpret that information in the same way, which results in less exploration of theory space.

After theoretically disaggregating information and solution spaces, we use our experimental results to demonstrate these two effects, namely: (1) when I tell you all of the facts that I know, it tends to encourage you to search for novel facts somewhere else; (2) but when I tell you my interpretation of those facts, it tends to influence you to adopt a solution in the direction of my interpretation. Rich data behind those results, and their underlying mechanisms, provide a foundation to construct an integrative model of how clustering of network ties can be used to improve problem-solving performance in knowledge-centered organizations, with implications for both theory and practice.

## II. KNOWLEDGE-CENTERED PROBLEM SOLVING: SEARCHING INFORMATION AND SOLUTION SPACES

Knowledge-centered organizations work with assorted potential solutions to problems, as well as the information that supports those various solutions. When problem-solving is conceived of as a search of solution spaces (Lazer and Bernstein, 2012), depicted as rugged, unobserved landscapes with many local optima (Rivkin, 2000; Rivkin and Siggelkow, 2003; Siggelkow and Rivkin, 2005) that mask a single, global optimal answer, the literature says that exploration (March, 1991) is decreased when networks have a more clustered structure. However, when one focuses on the other half of a knowledge-workers job — the search for and communication of information rather than solutions — prior research supports the oppo-

site conclusion: organizations are able to search more extensively when networks are more clustered.

In recognition of the fact that knowledge workers balance both tasks in real organizational life, we bring the two literatures together here. First, we briefly summarize recent work on exploration and exploitation in knowledge-centered problem solving as it relates to searching for solutions. Second, we explore ties between that literature and the literatures on searching for information in knowledge-centered organizations.

## **II.1 Searching Solution Spaces: Exploration vs. Exploitation**

March (1991:71) argued that members of organizations continually face the choice to “explore” new solutions or “exploit” existing ones, where exploration includes “things captured by terms such as search, variation, risk taking, experimentation, play, flexibility, discovery, innovation” while exploitation includes “such things as refinement, choice, production, efficiency, selection, implementation, execution.” At the organizational level of analysis, seeking novel solutions to problems previously unavailable within organizational boundaries would be a case of exploration, adopting a solution already available from another member of the organization would be a case of exploitation.

The notion of “exploration” implies some space to be explored by members of an organization with interdependencies across solution components. Consider the task of designing a new consumer electronic product. Choices made about the user interface can determine, in part, the set of possible choices in other dimensions, such as functionality, appearance, and electronic hardware. Because the pieces of the problem interact with each other, moving from a good design to a better one may require making changes in multiple dimensions. In other words, the better solution is not “close” to the good solution in the sense of just tuning a single feature, just as the iPod was not a simple step from the first MP3 players introduced prior to it, despite sharing many of the same component parts (Kahney, 2005:10-11). This suggests solution spaces, in more abstract terms, with multiple local optima – a “rugged” landscape to be explored by an individual or set of individuals. Since exploration is costly and uncertain, there can be incentives for individuals to favor exploitation wherever myopic self-interest or intra-organizational competition is present (March, 1991:85). At the collective level, however, the greater the level of exploration, the higher quality the organization’s eventual solution would be expected to be in the long run.

In simulation models, high-dimensional “rugged” landscapes are simulated (cf. Kauffman and Levin (1987)) to investigate how solution spaces are explored most effectively. Using agent-based modeling, Lazer and Friedman (2007) illustrate how network structure shapes collective problem solving success given the tradeoff between exploration, which they define

as the investigation of novel solutions on the landscape, and exploitation, or the utilization of already known solutions. They show that individual agents in highly connected or efficient networks converge rapidly on a relatively good solution and thus outperform inefficient networks in the short run. However, the inefficiency of less-connected networks that prevents rapid convergence also maintains diversity in the system resulting in more exploration, bringing more potential solutions into the organization. Inefficient networks, they find, eventually converged on superior solutions to those in efficient networks, because a greater region of the problem space was explored before agents converged on the best available solution.

In a large experiment on human subjects, Mason and Watts (2012) also find that networks that collectively explored more did better in the long run. However, they found very different structural correlates of exploration than those of Lazer and Friedman (2007). While Lazer and Friedman (2007) found that inefficient networks with a long average path length between any two individuals explored more, Mason and Watts (2012) found that inefficient networks – with, importantly, a high level of local clustering – explored less. The differences in these findings can be explained by reference to work by Centola and Macy (2007; Centola, 2010), who suggest that adopting a neighbor’s solution is more likely to occur within, rather than between, clusters of ties. Novel, exploratory solutions are both uncertain in advance and have material consequences. Given these conditions, the act of observing that multiple other people have adopted the same existing solution – as is more likely within clusters than between them – makes further diffusion even more likely, because of the apparent emergence of a consensus opinion. In the language of March (1991), this is to say that exploitation in solution space should be more common within clusters, and Mason and Watts (2012) showed strong evidence of this. Results from a similar earlier experiment (Mason, Jones, and Goldstone, 2005) with fewer participants were more equivocal, but could be understood to further support the idea that some structural separation between individuals maintained salutary diversity and thus better performance.

## **II.2 Searching Information Spaces: Information sharing and coordination**

Recent research on problem solving in information space supports the opposite conclusion about the effect of clustering: more clustered networks perform better. We examine this literature here, beginning with the related experimental literature and then addressing caveats and complications introduced by empirical research.

### II.2.1 Recent graph coloring experiments

There is an expanding recent experimental literature on collective problem solving in networks, primarily using variants of the “distributed graph coloring problem” as an experimental task (e.g., Kuhn and Wattenhofer, 2006; Kearns, Suri and Montfort, 2006). These studies can be seen as information-based (not theory-based) problem solving. In graph coloring tasks subjects must choose from among a discrete set of colors such that they do or do not match the choice of their neighbors. There is no subjective interpretation required for these tasks: each subject takes in information about their neighbors and selects their own color according to the instructions. In the graph coloring experimental literature, although there was no voluntary communication, each individual could observe the color choices of all their neighbors simultaneously (Enemark, McCubbins and Weller (2012) study the effect of extending visibility beyond the immediate neighbors).

In general, greater density of ties improves performance in these tasks (McCubbins, et al., 2009). A caveat, however, is that the effect of more ties depends on the specifics of the graph coloring task (Judd, Kearns and Vorobeychick, 2010). When there are sufficiently few sub-tasks to execute (colors to choose from in these experiments), adding extra edges can eliminate some or even all of the possible fully coordinated divisions of labor (valid graph colorings) (Enemark et al., 2011). To illustrate, consider a typical fire-fighting ‘bucket brigade’ in which there are only two subtasks: passing a bucket full of water forward to the next person and reaching back for the next bucket. A bucket brigade would be disrupted if it was carried out not in a line but in a cluster of individuals. Likewise, coloring a graph with only 2 colors such that no adjacent nodes have the same color is impossible if that graph contains any closed triads. However, for tasks that have a large number of sub-tasks, such as a search through information space, Enemark et al. (2011) show that a greater density of ties results in better performance.

A major mechanism in this increased performance within dense clusters appears to be greater mutual knowledge. Knowing what your neighbors’ neighbors are doing dramatically eases the distributed graph coloring problem (Enemark, McCubbins, and Weller, 2012). Clustering of ties means that (many of) one’s neighbors’ neighbors are also one’s own neighbors, and therefore that there is extensive mutual information in one’s neighborhood that can increase performance in search of information spaces.

### II.2.2 Earlier experiments

The first laboratory-based social network experiments were from MIT’s Small Group Network Laboratory (SGNL) by Bavelas (1950) and Leavitt (1951). In these experiments, five individuals were seated in carrels with slots to pass messages between carrels. The carrels were arranged to

impose one of four potential communication structures. Each participant was given a card printed with five out of a master set of six symbols, and by passing messages through their communication slots the group was tasked with finding the single symbol that was printed on all five of their cards. Each carrel was also fitted with six switches, one for each symbol for the participants to indicate their answer.

These experiments show that centralized networks were more effective in cooperative problem solving than clustered networks (Bavelas, 1950; Leavitt, 1951), apparently contradicting the graph coloring literature. In similar experiments, Guetzkow and Simon (1955) subsequently showed that the reason is because centralized networks encourage a more efficient form of task organization because the structure all but guarantees that at least one person receives all the information available in the network, which he or she can then disseminate to the rest of the organization members. Shaw (1954) went on to argue that “complex” (meaning more difficult) problems were better solved in “decentralized structures” (clusters), but Mulder (1960) later established that ultimately centralized networks performed better for both simple and difficult problems, once a centralized and coordinated decision structure evolves within the experimentally imposed communication structure.

In the protocol deployed in these experiments, each act of communication could only be directed to a single individual. Therefore, an individual placed in the clustered network could not “speak” to all of the other subjects, despite the fact that they were all connected to each other. Unlike recent experiments in which network ties could be used simultaneously, in these early experiments ties could only be used asynchronously. Clustering was therefore not associated with everybody having the same information. Indeed, the more paths information could take through the network, the less certain participants could be that they were communicating in an efficient way to complete the task. Although it is surely possible to construct a real-world collective problem solving task like the Bavelas experiments, these seem much more the exception than the norm. As long as communications can be addressed to multiple individuals at a time – for example, because they are all sitting around a table in a meeting, or because they can broadcast their status to multiple individuals on a social media platform – then clustering would ensure that those individuals had the full access to shared information, which is argued to be associated with better performance in both recent experiments and the early experimental work.

### **II.2.3 Knowledge Sharing, Information Processing, and Social Cohesion**

All of the above benefits of clustering for search of information spaces depend on a crucial assumption: that knowledge is indeed transferred between connected individuals. A relevant but distinct body of empirical



research does not take that for granted, and instead foregrounds the notion that if one of the jobs of an organization is to leverage the knowledge of multiple individuals (Kogut and Zander, 1992; Argote and Ingram, 2000), then the level of communication among those individuals is an important variable for organizational effectiveness. Communication allows privately held information to enter the collective knowledge base, but communication takes effort and therefore is costly for the communicator. Solving the social dilemma by balancing the individual and the group perspectives is thus an important managerial problem (Cabrera and Cabrera, 2002). Clustering is associated with easing this tension and therefore with better organizational performance.

Central actors, who bridge between multiple organizational divisions and thus have less clustering in their local neighborhood, face particular information processing requirements (Schneider, 1987) in part because switching between unrelated streams of information coming from different areas of the network is difficult (Speier et al., 1999). Furthermore, when information processing requirements are too high and individuals experience information overload, they show decreased problem solving performance, and ability to identify important information (O'Reilly III, 1980). We extrapolate that individuals who are not in a cluster of ties may therefore share fewer pieces of important information with those around them. At the organization level, therefore, structural holes around central actors could create bottlenecks, amplifying the separation of organizational sub-units.

Social cohesion – often associated with clustering – can promote knowledge sharing (O'Reilly, Caldwell, and Barnett, 1989; Hansen, 1999; Reagans and Zuckerman, 2001; Tsai, 2002; Reagans and McEvily, 2003). Shared norms, which can be more homogenous within than between clusters (Burt, 2004), can shape the level of knowledge sharing (Haas and Park, 2010; Quigly et al. 2007). On the other hand, knowledge sharing can be hampered by hierarchical (implying not clustered) structure, formal ties and competition for resources (Tsai, 2002). Other research has found social capital to be a major driver of knowledge sharing without reference to clustering. Anonymity depresses communication (Ma and Agarwal, 2007) and in computer-mediated communities, those who share knowledge have significant pro-social tendencies (Wasko and Faraj, 2000). Without significant social capital to facilitate knowledge integration, electronic communication was more likely to be misunderstood and negatively affect decision quality (Robert, Dennis and Ahuja, 2008). Additionally, knowledge diffuses more readily when individuals hold related knowledge (Hansen, 2002).

According to the above research, the level of knowledge sharing is higher within clusters, contributing to superior performance. In this empirical literature, however, it is far from clear that structure per se is the driving force in these effects, but experimental results do point in the same direction: clustering promotes more extensive collective search of information space.

### II.3 Reconciling the contradictions: information vs. solutions

Some studies find that clustering improves performance, while others find that it harms performance. Above, we have argued that the dividing line is whether the studies were mostly concerned with raw information or interpretations of that information (theories and solutions). The research focused on information finds better performance with greater clustering, while the research emphasizing solutions finds worse performance. The reason for the divergence in results lies in the difference between information and its interpretation.

When communication can be directed to more than one network neighbor at a time, clustering promotes mutually held information. Mutually held information in turn allows individuals to coordinate with each other: when everybody knows what everybody else is doing, it allows them to act in a way that is complementary to the actions of the rest of the group. Clustering does not guarantee coordination, but by promoting mutual awareness, it does allow individuals to act in a complementary way to each other if they choose to do so. When it comes to gathering information, the same principle applies: the more people are aware of what information their neighbors are searching for, the easier it will be to avoid doing redundant work.

On the other hand, a theory is an interpretation of a collection of information, or alternatively, the meaning that an actor attributes to that information. Information and theories can therefore be seen as isomorphic in many ways to words and their meanings (e.g., Lakoff, 1987). We use words to signify meanings, but the process depends on interpretation by the recipient of the communication. To take a trivial example, in a grocery store, the word “apple” is likely to refer to the fruit; in an IT department, it is likely to refer to a computer. The example is trivial because the mental models necessary to interpret the word “apple” are simple, widely shared, and built from much experience. Interpreting information to generate a theory is not always so simple. Economists may have access to the same information about the economy – the unemployment rate, GDP growth, the availability of credit and so on – but they must supply a great deal of interpretation to produce a forecast. Theories have elements of complex contagions in that they are uncertain and can be consequential if they justify one action over another (Centola and Macy, 2007). Under these circumstances, we should expect that clusters would decrease exploration of theory space, because of the pressure to copy theories of one’s neighbors.

In sum, our careful reading of the literature on clustering suggests the following possible interpretation: most extensive aggregate exploration in solution space occurs when actors are independent, while most extensive aggregate exploration in information space occurs when actors are interdependent. The experiment we describe below allows us to ask and test if clustering affects performance in information space and solution space

differentially. We expect to see clustering associated with more exploration in information space, but less exploration in solution space.

### III. DATA AND METHODS

#### III.1 The experimental platform

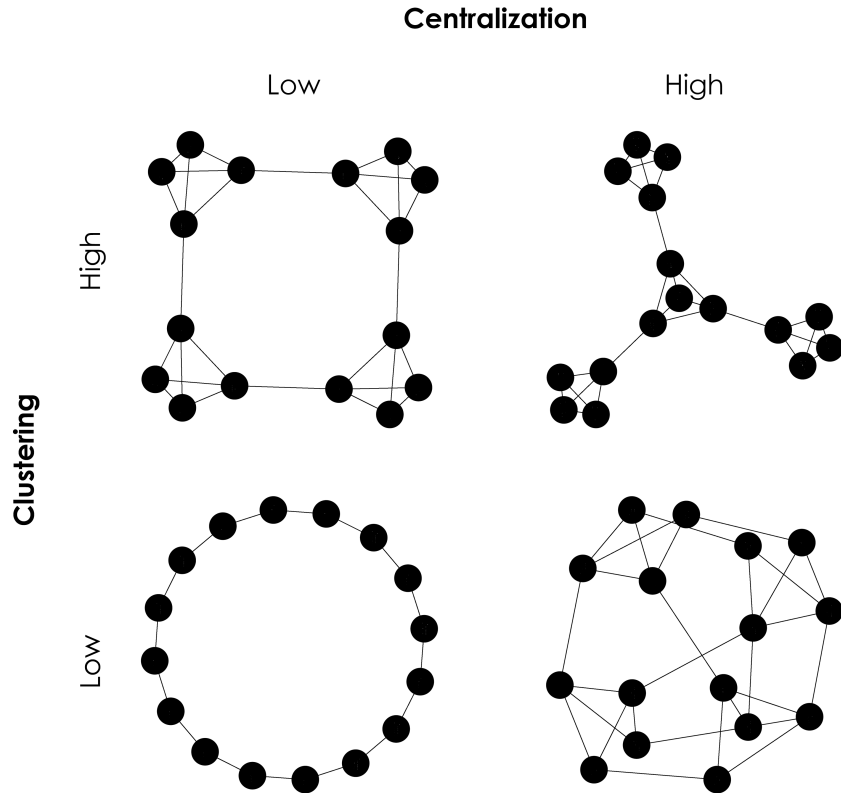
##### III.1.1 The task

To instrument the connection between solution search, information search, and problem-solving performance, we aimed to develop an experimental platform with several key characteristics: (1) maximum verisimilitude, which meant both that the task was similar to real problem-solving work and that the means for accomplishing the task within the platform had real-world analogues; (2) maximum accessibility, which required the task to be easily understandable and solvable with expertise commonly available in our subject pool; and (3) maximum instrumentation, which required that actions taken by the participants to be captured as richly as possible in subsequently analyzable data.

Based on those criteria, we selected a whodunit protocol, much like a game of Clue© or Cluedo©, in which the task involved piecing together clues to "connect the dots." Rather than creating a platform entirely from scratch, we were invited to customize a platform developed by the United States Department of Defense's Command and Control Research Program called ELICIT (Experimental Laboratory for Investigating Collaboration, Information-sharing, and Trust), which had many of the characteristics we sought. While we modified much of the platform, we agreed to keep the nature of the Department of Defense's whodunit task, which involved predicting the who, what, where, and when of an impending terrorist attack (in place of, for example, the who, what, and where of the murder in Clue©).

Specifically, there were four logically-independent sub-problems to be solved: (1) who would carry out the attack (group involved); (2) what the target of the attack was (e.g. an embassy, a church, etc); (3) where the attack would take place (country); and (4) when the attack would take place (with four interdependent components—month, day, hour, and AM/PM). Each question and sub-question had a dedicated text box in which to register an answer. Participants were rewarded 15 cents per minute per sub-problem (3.75 cents per minute for each component of "when"), for a maximum of 60 cents (a penny per second) for each minute that they had the correct answer registered. Participants therefore had a strong incentive to record and adjust their theories as soon as they had developed a theory about the answer. At no point during the experiment did they know for certain whether their answers were correct, just as would be the case in real life.

There was a discrete set of actions available to experimental subjects.



**Figure 1:** Network treatments used in the experiments. Top to bottom, left to right, we call them the “caveman network” (abbreviation: CAVE), “hierarchy” (HIER), “ring lattice” (LATT), and “rewired caveman” (RCAVE).

They could search for new clues by entering a keyword into a search text box and clicking a button. They could share these clues with one or more of their neighbors and, if they wished, add free text annotations to these shared clues. They could register their theories by typing them into the separate spaces given for the who, what, where and when sub-problems. Finally, although their neighbors’ theories were not usually visible, they could choose to check them by clicking a button.

Since the information space was large, participants would have been extremely unlikely to find all of the necessary clues to solve all of the problems within the time limit on their own. However, because participants could also choose to share clues with their network neighbors, annotate those clues, and view their neighbors’ registered solutions, the task was a collective problem-solving situation.

### III.1.2 Treatments

We tested four 16-person network treatments (see Figure 1 for visualizations and Table 1 for descriptive statistics). At the top left of Figure 1 is the so-called “caveman” network (Watts, 1999), containing four four-person cliques. The “hierarchy” is likewise composed of four such cliques, but arranged in a conventional centralized structure. The “rewired caveman” is a small world network, constructed by removing links from the caveman, then adding links that create shortcuts through the network. As a result, individuals in the rewired caveman are “closer together” topologically: the most distant pair of individuals is only three hops away, and the average distance between all pairs is shorter than the caveman and the hierarchy. The rewired caveman is also more centralized and less clustered than the caveman. Finally, there is the ring “lattice,” which is neither clustered nor centralized.

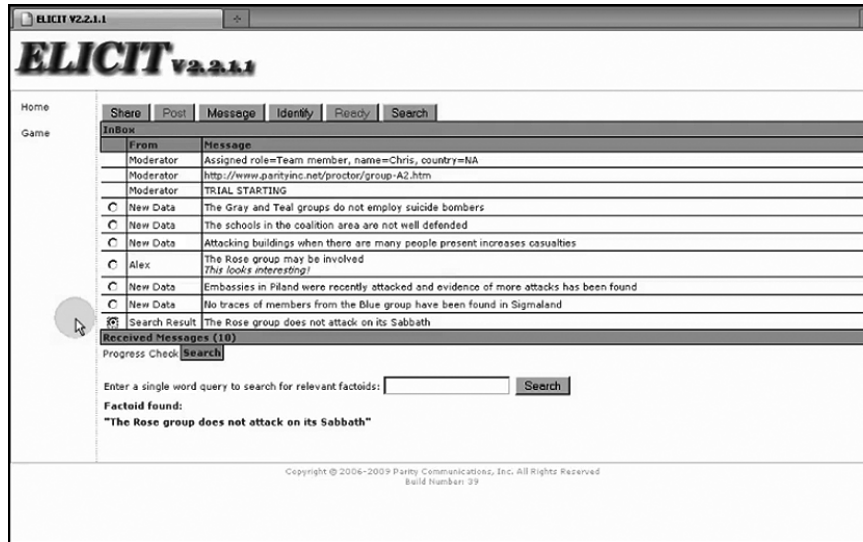
**Table 1:** *Descriptive statistics for networks*

	CAVE	RCAVE	HIER	LATT
average degree	3.5	3.625	3.375	2
density	.467	.483	.45	.267
average path length	2.47	1.99	2.81	4
diameter	5	3	5	8
mean clustering coefficient	.667	.304	.727	0
centralization (eigenvector)	.033	.115	.161	0

In some of our statistical analyses, we also test the effects of nodal degree and clustering coefficient. Both are individual-level structural metrics: degree is simply the number of connections a node has; clustering coefficient measures the extent to which a node’s neighbors are also neighbors with each other, or the number of connections among neighbors divided by the number of possible connections among neighbors (Watts and Strogatz, 1998).

### III.1.3 Execution of the experiment

The experiment was carried out in a laboratory setting, with each participant seated at a computer in a private carrel. All experimental activities were executed through a web browser interface (see Figure 2), with the exception of scratch paper, which was collected and scanned at the experiment’s conclusion. Each experimental run lasted 25 minutes. Participants were given two clues at the start of each round, and were allowed to search for clues once per minute. Each clue was only relevant to one sub-problem. Some clues contained useless information. Some clues contained misleading information. Subjects had to combine multiple clues to conclusively arrive at the correct answer. The number of clues required to conclusively arrive at the correct answer ranged from 2 to 10, with a median of 5 and a



**Figure 2:** A screenshot of the web browser interface. Pictured is the “inbox” of an experimental subject.

mean of 5.3.

Each experimental session began with an instructional video that explained the platform and the task in uniform fashion to every subject across all sessions. To control for individual aptitude, each person then took a pre-test in the same format as the experimental task, without interaction with other experimental subjects.

At any given time, lab space allowed us to run a maximum of two concurrent 16-person experimental runs. Prior to a run beginning, a network treatment was chosen at random, and study subjects were randomly assigned to a position in that network which was uncorrelated to their physical location within the laboratory. Participants were assigned a pseudonym in order to further obscure their identities from other subjects.

In observational research, the effects of network structure on effective knowledge sharing have been intertwined with the effects of other, non-structural types of social capital (Reagans and McEvily, 2003). Indeed it is essentially impossible to identify the independent effects of structure in field work (Shalizi and Thomas, 2011), requiring experimental methods to untangle its consequences. Given the difficulties in achieving a high level of knowledge sharing in real organizations, we suggest that experimental protocols should allow the level of knowledge sharing to vary by making it voluntary. In adopting elegant and parsimonious protocols, recent experiments have not done this, but we believe it is a variable that should not be excluded and thus include it in our protocol.

There were 417 unique individuals, who played a total of 1120 person-rounds. Participants were recruited through the subject pool of an elite

university. The mean self-reported math SAT score was 716, and the mean verbal score was 701, consistent with reported data from the university. The ratio of self-reported gender was approximately equal, with 49.5% male and 50.5% female. The present paper reports results from a subset of the collected data, consisting of 816 person-rounds, played by 352 unique individuals. The remaining data included additional treatment variables intended to test other phenomena and are not comparable to data we analyze here.

Multiple problem sets were used to limit contamination between sessions of the experiment. Additionally, within each problem set, proper nouns (names of terrorist groups, country names, place names) were randomly permuted to reduce the risk of contamination between sessions.

## III.2 Statistical framework

Wherever possible, we consider both individual and collective-level correlations. At the collective level, we have 51 independent data points, each corresponding to a single run of the experiment. For these models, we use ordinary least squares (OLS). At the individual level, we have 816 observations with two types of interdependence. First, they are, nested into the 51 runs mentioned above. Second, since individuals played multiple runs of the experiment, the 816 observations were generated by 352 unique individuals. We therefore include random effects for both run and unique individual, and estimate linear mixed models (abbreviated LMM in the tables of results).

For discrete outcome variables at the individual level, we employed mixed-effects Poisson (GLMM-Poisson) regression (Bates, Maechler, and Bolker, 2012). The number of times a participant checked the answers of their neighbors exhibited zero-inflation, and thus a zero-inflated Poisson mixed-effects model (GLMM-ZIP) was estimated in a Markov Chain Monte Carlo framework (Hadfield, 2010). Statistical analysis was carried out in R (R Core Team, 2012).

The experimental treatments we impose are the relational structures into which individuals are placed: the networks as a whole. Aggregate results at the collective level can therefore be inferred to be causal results of these treatments. The tables of results that we present in the following section contain both causal and non-causal results, as follows. Results from the models that use entire network treatments as the independent variables can be interpreted as causal inferences, while those that use individual level variables – even exogenously imposed structural variables, like degree – should be interpreted as correlations. Outcomes at the individual level depend not only on the local structure, but also the structure of the remainder of the network, and node-level metrics do not account for this. Nonetheless, they can help us to understand the causal results at the collective level by providing more detail.

### **III.3 Outcome variables**

#### **III.3.1 Exploration and exploitation in information space**

We treat exploration and exploitation as mutually exclusive classifications of a single action (March, 1991; Lazer and Friedman, 2007; Mason and Watts, 2012). That is, we assume that if an action is an example of exploration, then it is not also an example of exploitation. This is important in our measurement of these constructs, because in different cases one or the other is easier to observe. For example, it is easier to measure the amount or extent of information gathering (exploration) than mental processing of information already held (exploitation).

We measure exploration in information space in terms of a greater extent and lower redundancy of clues found and shared by the subjects. At the individual level, we can measure the level of redundancy of facts received from a focal subject's partners and the total number of facts found by subject's own search (the search interface did not return redundant facts, so the total number of facts represents the extent of exploration in information space by search). At the collective level, we can measure the extent and redundancy of both searches and shares.

Sharing facts (what we call "shares" of facts) do not equal receipts of facts because an individual share can be received by more than one individual. This is true not only in computer-mediated communication such as this experimental platform, but also in face to face interaction in which more than two people are present at the same time. One of the effects of the communication network structure is to influence the relationship between the number of shares and number of receipts of facts. When measuring sharing of information, we use unique facts shared to capture the contribution of an individual to the collective pool of information.

#### **III.3.2 Exploration and exploitation in solution space**

We measure both exploration and exploitation in solution space. Exploitation in solution space takes the form of checking and then copying a neighbor's theory. We have time-stamped records of every action undertaken during the experiment, so this can be directly measured from the data. We define copying to be when an individual checks their neighbors' answers, and then registers an answer they observed the next time they enter their own theories, provided this occurs within 10 minutes of observation of the neighbors. We measure the extent of exploration in solution space at the aggregate level in terms of the total number of unique theories that were registered during the experiment.

Establishing uniqueness of theories required us to be able to consolidate answers like "power plant" and "powerplant" and "electric power plant" into one theory, which we did in two steps. First, automated pre-processing of entries removed punctuation, converted the text to lowercase, and



combined repeated entries where one example had a simple typo (defined as a single insertion, substitution or deletion of a character – provided that it did not involve the first letter of the word – like “power plant” and “powerplant”). Second, we used a human coder to remove more substantial typos (such as transposing letters or whole phonemes) as well as answers where the intent was clearly the same (for example, we considered “power plant” and “electric power plant” to be the same).

### III.3.3 Performance

Given that clustering is expected to have both positive and negative effects, depending on the domain of reference, we also measure overall performance. Performance was measured in pay per minute received by individuals. When measuring performance of an entire network, we simply took the sum of the pay per minute for each individual member of the network.

The result section below reports on each of the three sets of variables above in turn.

## IV. RESULTS

The most clustered networks, the caveman and the hierarchy, delivered the highest collective problem-solving performance in our Clue©-type problem-solving task. The aggregate result from our experiment therefore supports the side of the existing debate that encourages clustering in networks. The more interesting insights, however, lie in the details behind that performance improvement which we explore in the three sections below.

### IV.1 Exploration and exploitation in information space

Higher performance was driven, in part, by extensive collective search through information space. We find that clustering was associated with greater exploration of information space in the sense of finding and sharing not more, but rather more unique, information. By being in a cluster, individuals tended to contribute more to the collective exploration through information space, not from more search, but instead by being more coordinated in their search.

Holding degree constant, increased clustering did not result in an increased number of searches. At the individual level, the clustering coefficient of an individual’s position was not correlated with the number of searches for facts that they perform (Table 2, Model 1). Likewise, at the collective level, the clustered networks did not search at a different rate than the rewired caveman (Model 2).

However, at the collective level, the facts found by the caveman network were significantly less redundant than those found by the ring lattice and

**Table 2: Exploration and Exploitation in Information Space**

	Dependent variable							
	Total facts found by search				Total Searches Unique Facts Found		Total receipts Unique Facts Received	
Model number	1		2		3		4	
Random effects	Variance				Variance			
Individuals	0.012						0.001	
Run	0.000						0.006	
Fixed effects	Estimate	<i>p</i>	Estimate	<i>p</i>	Estimate	<i>p</i>	Estimate	<i>p</i>
(Intercept)	3.085	<0.001***	267.037	<0.001***	4.002	<0.001***	0.512	<0.001***
Degree	-0.085	<0.001***					0.161	<0.001***
Clust. Coef.	-0.019	0.423					-0.131	<0.001***
Hierarchy			2.424	0.722	0.114	0.198		
Rewired Cave			7.231	0.257	0.199	0.017*		
Ring Lattice			46.248	<0.001***	0.761	<0.001***		
Pre-test	0.002	0.025*					-0.000	0.315
Second round	0.061	0.002**	18.668	0.003**	0.137	0.082.	0.064	0.031*
Third round	0.071	<0.001***	21.320	<0.001***	0.190	0.018*	0.112	<0.001***
Factoid set 2	0.119	<0.001***	35.172	<0.001***	-0.236	0.004**	-0.116	<0.001***
Factoid set 3	0.098	<0.001***	28.814	<0.001***	0.075	0.353	-0.047	0.117
<i>n</i>	816		816		51		816	
Model type	GLMM (Poisson)		OLS		OLS		LMM	
Unit of analysis	Individual		Aggregate		Aggregate		Individual	

the rewired caveman treatments (Model 3). In other words, if the objective of information search was exploration, the caveman network was more efficient, in that they collectively covered more new ground with each search. The mean redundancy for the hierarchy was also lower than the ring lattice and the rewired caveman treatments, although the difference was not significant – a fact that we consider further below in the section on centralization.

Holding clustering constant, greater degree was actually associated with decreased number of searches (Model 1). At the collective level, the lattice searched at a higher rate than the other three networks (Model 2), which was likely a consequence of its members having a low degree.

In addition to searching for information, exploration of information space at the collective level requires that information to be effectively transferred throughout the network. We found that individuals in clustered positions received significantly less redundant information from their network neighbors (Model 4).

## IV.2 Exploration and exploitation in solution space

Exploration of theories, or what we call search of solution space, operated quite differently. At the collective level, the total number of unique theories registered was highest in the unclustered networks: the lattice and the rewired caveman both had significantly more unique theories than the caveman network.

Table 3: Exploration and Exploitation in Solutions Space

	Dependent variable							
	Theory checking		Theory copying		Unique theories			
Model number	5		6		7		8	
Random effects	Variance		Variance					
	Individuals	0.102		0.028				
Run	0.000		0.035					
Fixed effects	Estimate	<i>p</i>	Estimate	<i>p</i>	Estimate	<i>p</i>	Estimate	<i>p</i>
	(Intercept)	1.922	<0.001***	-0.021	0.847	20.841	<0.001***	35.097
Degree	0.022	0.473	0.072	0.031*				
Clust Coef	0.126	0.052.	0.069	0.048*				
Hierarchy					-2.426	0.5346	0.440	0.864
Rewired Cave					-2.640	0.4675	6.151	0.013*
Ring Lattice					-7.950	0.0616.	6.174	0.029*
No theory checks			-1.069	<0.001***				
Pre test	0.002	0.223	0.002	0.500				
Second round	0.275	<0.001***	0.370	<0.001***	14.888	<0.001***	-3.218	0.160
Third round	0.294	<0.001***	0.408	<0.001***	7.716	<0.001***	-5.005	0.032*
Factoid set 2	-0.014	0.772	-0.216	0.020*	-5.013	0.147	9.536	<0.001***
Factoid set 3	0.011	0.791	0.128	0.175	4.275	0.229	3.051	0.193
<i>n</i>	816		816		51		51	
Model type	GLMM (Poisson)		GLMM (ZIP)		OLS		OLS	
Unit of analysis	Individual		Individual		Aggregate		Aggregate	

For exploration and exploitation in solution space, we model the effects of our independent variables on the propensity of individuals to check and copy their neighbors' theories, and on the aggregate amount of copying and total number of unique theories registered at the whole organization level.

In the space of solutions, our results agree with those of Mason and Watts (2012): clustering promotes exploitation (Table 3). More specifically, clustering is associated with more checks of neighbors' theories (Model 5: marginally significant finding) and more outright copying of their theories (Model 6). However, consistent with the predictions of the information processing literature (Galbraith, 1974; Schneider, 1987), we also found that degree (Model 6) and total information received (Model 7) were correlated with less exploration (greater copying). At the collective level, the two clustered networks had significantly fewer unique theories registered in the aggregate than did the unclustered networks (Model 8).

### IV.3 Performance

At the individual level, clustering is not correlated with performance (Table 4, Model 9). Interestingly, although there is no evidence for a positive effect of clustering on performance at the individual level, the most clustered networks (the caveman and the hierarchy) did best as a whole (summing pay

**Table 4: Performance**

	Dependent variable			
	Pay per minute			
Model number	9		10	
Random effects	Variance			
Individuals	7.240			
Run	8.231			
Fixed effects	Estimate	<i>p</i>	Estimate	<i>p</i>
(Intercept)	16.199	<0.001***	392.14	<0.001***
Degree	1.441	0.003**		
Clust. Coef.	1.364	0.216		
Hierarchy			-26.24	0.172
Rewired Cave			-36.25	0.046*
Ring Lattice			-65.73	0.002**
Pre-test	0.055	0.022*		
Second round	7.505	<0.001***	118.31	<0.001***
Third round	10.689	<0.001***	183.04	<0.001***
Factoid set 2	-9.270	<0.001***	-132.86	<0.001***
Factoid set 3	-6.114	<0.001***	-95.87	<0.001***
<i>n</i>	816		51	
Model type	LMM		OLS	
Unit of analysis	Individual		Aggregate	

per minute of all network members), although like above, the differences were not significant for the hierarchy (Model 10). The rewired caveman network performed significantly worse than the caveman network, with a mean performance greater only than the lattice, which had the worst performance (Model 10) except in the first round of play, in which the rewired caveman had the lowest mean performance of all the treatments, including the lattice (Table 5, Models 11-14). Ironically, clustering improved collective results without benefiting individuals themselves—no additional performance benefits accrued to individuals located more highly clustered positions within a network (Model 9).

Table 5 allows a comparative analysis of the variance structure of the four network treatments. Inspection of the random effects variances reveals several interesting points. First, the run-level random variance can be interpreted as the degree to which people in the same organization tended to do well or poorly together: the intra-run correlation of individual performance. This was greatest in the caveman, followed by the lattice, then the hierarchy. Strikingly, the run-level random effect was zero for the rewired caveman, indicating no measurable correlation of performance for people in the same organization.

In the caveman and the lattice, performance varied more by run than by individual, while the opposite was true for the hierarchy and the rewired

Table 5: Pay per minute by network structure

	Treatment							
	CAVE		HIER		RCAVE		LATT	
Model number	11		12		13		14	
Random effects	Variance		Variance		Variance		Variance	
Individuals	9.003		6.950		3.774		4.318	
Run	11.201		3.1480		0.000		9.601	
Fixed effects	Estimate	<i>p</i>	Estimate	<i>p</i>	Estimate	<i>p</i>	Estimate	<i>p</i>
(Intercept)	15.823	<0.001***	22.529	<0.001***	12.385	<0.001***	10.9078	0.002**
Total Facts Rec'd	5.004	0.010*	2.965	0.161	7.012	0.002**	13.818	<0.001***
Pre-test	0.144	<0.001***	-0.060	0.177	0.004	0.918	0.1415	0.009**
Second round	9.868	<0.001***	5.653	0.007**	7.967	<0.001***	1.105	0.716
Third round	8.851	<0.001***	11.723	<0.001***	12.809	<0.001***	3.489	0.357
Factoid set 2	-12.139	<0.001***	-6.556	0.001**	-5.343	<0.001***	-6.371	0.066.
Factoid set 3	-5.978	0.015*	-8.975	<0.001***	-2.083	0.2449	-9.662	0.001**
<i>n</i>	272		176		224		144	
Unique individuals	154		150		135		136	
Runs	17		11		14		9	
Model type	LMM		LMM		LMM		LMM	
Unit of analysis	Individual		Individual		Individual		Individual	

caveman. It is also striking that individual aptitude, as measured by the pre-test, made significant contributions to performance for both the caveman and the lattice, but was uncorrelated with performance in the two more centralized networks.

#### IV.4 The effects of Centralization

Unlike the caveman, the hierarchy is both clustered and centralized, and results for this treatment reveal several subtler points that imply enrichments to our main finding. In particular, the key benefit of clustering – finding less redundant information – did not result in a statistically significant difference from the rewired caveman and the lattice. Moreover, like the other centralized network – the rewired caveman – individual aptitude had no measurable correlation with performance, and the within-run correlation of individuals' performance was low. We have not focused on the language of centralization in the present paper, but we can nonetheless make sense of these results by interpreting centralization in terms of the concepts we have developed.

The critical effect of clustering is to promote mutual awareness of information and solutions. In the hierarchy, sharing mutual knowledge with people in other cliques faces impediments not present in the caveman network. Because of their position in the network, central individuals in the hierarchy may be overloaded with information, disrupted by changes in topic, or simply divided in attention. The peripheral cliques in the hierarchy do not face these difficulties, but their only source of outside

knowledge is the potentially overloaded central clique. In contrast, in the caveman, there are two paths between each pair of cliques, and no clique faces the added information processing challenge of being central.

If centralization creates confusion and bottlenecks, then this counteracts the benefits of clustering which are mutual awareness that leads to more efficient or coordinated search of information space. Therefore, while both the caveman and the hierarchy are highly clustered locally, the centralization of the hierarchy introduces a force akin to anti-clustering at a higher, organization-scale level.

## V. DISCUSSION

One can frequently hear marveling comments about how small our world has become. With the advent and accelerating adoption of increasingly powerful communication technologies and greater international travel, our world is becoming ever more interconnected at every scale. In network terms, small world networks have been long associated with surprisingly extensive diffusion of a given piece of information (Travers and Milgram, 1969; Granovetter, 1973; Watts and Strogatz, 1998), due in part to the existence of short paths between any pair of individuals. It is therefore surprising, if not contradictory to those findings, that in the rewired caveman treatment – a small world network – people depended less on each other than in other network environments. In the small world, correlation of performance of individuals within the same run was zero, suggesting less or less effective collaboration by experimental subjects. The small world runs also had the least amount of sharing and a high redundancy of facts found by search. In sum, the overall character of results from these runs was that people were more on their own than in the other networks.

But short paths between any given pair of individuals are not the only impact of greater communications connectivity. Broadcasts and publications, all manner of digital information systems including, social media, topic-specific RSS feeds, and mobile applications – even, at the global scale, air travel and international trade – all promote mutual awareness and shared knowledge, much as clusters do in smaller scale network terms. Our results can be understood to tell us a little more about the effects of that greater connectivity. We have argued throughout that the key feature of clustering is mutual awareness: within a cluster, everybody is aware what everybody else is doing. In information space, this promotes exploration by allowing a sort of emergent coordination to occur in that people tended to avoid duplicating the work that they knew had already been done. In solution space, it inhibits exploration by allowing more rapid convergence on a consensus about the problems' solutions.

The more connected we are, the more we coordinated we become — either in a self-organized, emergent, “invisible hand” sort of way, or in an intentional delegation sort of way — and the greater the diversity of what

we do and can find out. We can celebrate how improved connectivity is making us ever better at coordinating our exploration of the facts of the world, just as our 16-person networks did in this experiment. Searching for facts is becoming easier and easier as global networks become increasingly dense. Greater interconnectivity can promote other emergent forms of coordination at the global scale as well. Increased geographic division of labor and the creation of niche communities of interest that could not otherwise sustain themselves are examples of increased diversity due to greater interconnectivity – or, one could say, increased aggregate exploration of activity space.

At the same time, the results above should make us extremely cautious about what that clustering means for our interpretation of those now plentiful facts and how broadly we explore the possible answers from these facts that we are increasingly good at finding. And with respect to the increased diversity of actions undertaken by the aggregate of humanity, we might worry that the way we understand those actions is becoming increasingly similar. Although we have more different types of goods and services than ever before, we have little diversity in economic policies, just as we are sometimes also warned of spreading global monoculture and McDonaldization.

For knowledge intensive organizations, the implication is that connecting everybody with increasingly high bandwidth communications technologies may improve coordination, but reduce diversity in the knowledge created within the firm. One possibility is that organizations could adopt different communications structures for different phases of collective problem solving. When information gathering and sharing is important, then clustering is beneficial. For such a phase of information gathering and dissemination, clustering will aid in greater exploration in as much as the information is not yet interpreted. Ordinarily, however, searches of information space are guided by some sort of hypotheses or mental models (whether explicit or tacit). Our notions of the relevant information space to explore derives from our working theories about the world. If a team wishes to find a better protective coating for an electronic product, individual team members will probably not elect to search the space of information about the sugar content of fruits, no matter how well coordinated they are in avoiding duplicating their information gathering work, because no working theory of protective coatings requires such information. In other words, it is inevitable that some degree of interpretation is always occurring in the minds of the information gatherers. If we wish to encourage the widest possible exploration of relevant information space, individuals should be arranged in clusters, extensively share their raw information with their neighbors, but keep their working theories that guided their searches to themselves.

When information is in hand and it is desirable to generate diverse interpretations of it to generate theories or solutions, then prior literature shows that less clustering is desirable, even within organizational sub-

groups, such that individuals do not prematurely coalesce on a consensus. In our experimental task, this was not a driver of net performance, but we note that it could well be a crucial factor for other problems.

Another organizational response would be to design communications infrastructures that could somehow separate facts from figuring, and adopt differently structured communication networks for each category. In other words, rather than allow the march of technology dictate organizational performance, it is possible to imagine how technology could be harnessed to achieve different performance goals. Even without separation of facts and figuring, these results are likely to be especially relevant for computer-mediated problem solving, because of the ease of manipulating the communication structure of participants. Internal social networking and knowledge management software and external crowdsourcing platforms seem to be fertile ground for testing these implications.

Building on this basic finding, much work could be done to refine the theory and establish boundary conditions. Further experiments should investigate whether these results hold when information and solution spaces are much more rugged than those used here. Empirical work should ask if they hold when structure and social capital are both operating. Moreover, both experimental and observational studies could be enriched by inclusion of embedded or tacit forms of knowledge (cf. Argote and Miron-Spektor, 2011).

Experiments are an especially appealing method to study networks, but our results should strike a note of caution for future experimental designs. Human problem solving has many facets, and if the experimental platform does not measure them separately or omits some entirely, the conclusions are likely to be confused and incomplete. The effects of structure can be very different, depending on the domain in question.

## V.1 The tradeoff between facts and figuring

It is well-established that network structure can influence problem solving performance, and yet a clear understanding of the role of clustering, a basic structural network variable, has remained elusive. By theoretically and experimentally disentangling information search from solution search – two core domains of problem-solving – this study indicates that the effect of clustering is opposite in those two domains. Clustering promotes exploration through information space, but depresses exploration through solution space. Whether increased clustering improves or impairs performance will depend on whether the immediate task or problem-solving stage benefits more from exploration of facts or, instead, the figuring that comes through the exploration of theories that interpret those facts.

Awareness of the differential performance effects of clustering for problem solving in information space from problem solving in solution space presents two challenges. For networks of problem-solving individuals, whether they represent groups, organizational units, whole organizations,



or clusters of organizations, the challenge is one of leadership, such that leaders find ways of pairing the domain of the problem solving task, whether facts or figuring, with an appropriate network structure, whether clustered or not, to improve problem-solving performance. For scholars of both networks and information science, the challenge is one of further research: integrating our basic finding of the distinction between facts and figuring into the examination of how different network structures impact performance may help not only to resolve existing conflicts in disparate yet interconnected literatures, but also open up substantial opportunities for greater, coherent understanding of how we can set the conditions for problem-solving success in networks.

Clustering is a double-edged sword. It has the power to bring members of a network to generate more non-redundant information, but it also has the power to discourage theoretical exploration. Until one knows whether a problem-solving task involves searching for facts or searching for answers, it is impossible to predict the influence of clustering on organizational performance.

#### REFERENCES CITED

- Abrahamson, E., & Rosenkopf, L. (1997). Social network effects on the extent of innovation diffusion: A computer simulation. *Organization science*, 8(3), 289-309.
- Alavi, M., & Leidner, D. E. (2001). Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS quarterly*, 107-136.
- Argote, L. 1999; *Organizational Learning: Creating, Retaining and Transferring Knowledge*. Boston: Kluwer Academic Publishers.
- Argote, L., & Ingram, P. (2000). Knowledge transfer: A basis for competitive advantage in firms. *Organizational behavior and human decision processes*, 82(1), 150-169.
- Argote, Linda and Ella Miron-Spektor (2011) "Organizational Learning: From Experience to Knowledge" *Organization Science* 22:5 1123-1137
- Argote, L., McEvily, B., & Reagans, R. (2003). Managing knowledge in organizations: An integrative framework and review of emerging themes. *Management science*, 49(4), 571-582.
- Bailey, D. E., Leonardi, P. M., & Chong, J. (2010). Minding the gaps: Understanding technology interdependence and coordination in knowledge work. *Organization Science*, 21(3), 713- 730.
- Baldwin, T. T., Bedell, M. D., & Johnson, J. L. (1997). The social fabric of a team-based MBA program: Network effects on student satisfaction and performance. *Academy of Management Journal*, 1369-1397.
- Bates, Douglas, Martin Maechler and Ben Bolker (2012). lme4: Linear mixed-effects models using Eigen and Eigen. R package version 0.999999-0. <http://CRAN.R-project.org/package=lme4>

- Bavelas, A. (1950). Communication patterns in task-oriented groups. *The Journal of the Acoustical Society of America*, 22(6), 725-730.
- Burt, Ronald S. 2004. Structural Holes and Good Ideas. *American Journal of Sociology* 110:2, pp. 349-399
- Burt, R. S. (2009). *Structural holes: The social structure of competition*. Harvard University Press.
- Butts, C. T., Petrescu-Prahova, M., & Cross, B. R. (2007). Responder communication networks in the World Trade Center disaster: Implications for modeling of communication within emergency settings. *Mathematical Sociology*, 31(2), 121-147.
- Cabrera, A., & Cabrera, E. F. (2002). Knowledge-sharing dilemmas. *Organization studies*, 23(5), 687-710.
- Carlisle, P. R. (2004). Transferring, translating, and transforming: An integrative framework for managing knowledge across boundaries. *Organization science*, 15(5), 555-568.
- Centola, D. (2010). The spread of behavior in an online social network experiment. *Science*, 329(5996), 1194-1197.
- Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3), 702-734.
- Cummings, J. N., & Cross, R. (2003). Structural properties of work groups and their consequences for performance. *Social Networks*, 25(3), 197-210.
- Edmondson, A. (2011). Strategies for learning from failure. *Harvard Business Review*, 89(4).
- Edmondson, A. C., Winslow, A. B., Bohmer, R. M., & Pisano, G. P. (2003). Learning how and learning what: Effects of tacit and codified knowledge on performance improvement following technology adoption. *Decision Sciences*, 34(2), 197-224.
- Enemark, D., McCubbins, M., Paturi, R., & Weller, N. (2011, June). Does more connectivity help groups to solve social problems. In *Proceedings of the ACM Conference of Electronic Commerce* (pp. 21-27).
- Enemark, D., McCubbins, M. D., & Weller, N. (2012). Knowledge and networks: An experimental test of how network knowledge affects coordination. *Social Networks*. <http://dx.doi.org/10.1016/j.socnet.2012.10.001>
- Galbraith, J. R. (1974). Organization design: An information processing view. *Interfaces*, 4(3), 28-36.
- Granovetter, M. S. (1973). The strength of weak ties. *American journal of sociology*, 1360-1380.
- Guetzkow, H., & Simon, H. A. (1955). The impact of certain communication nets upon organization and performance in task-oriented groups. *Management science*, 1(3-4), 233-250.
- Haas, M. R., & Park, S. (2010). To share or not to share? Professional norms, reference groups, and information withholding among life scientists. *Organization Science*, 21(4), 873- 891.
- Hadfield, Jarrod D (2010). MCMC Methods for Multi-Response Generalized Linear Mixed Models: The MCMCglmm R Package. *Journal of Statistical Software*, 33(2), 1-22. URL <http://www.jstatsoft.org/v33/i02/>.

- Hansen, M. T. (1999). The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative science quarterly*, 44(1), 82-111.
- Hansen, M. T. (2002). Knowledge networks: Explaining effective knowledge sharing in multiunit companies. *Organization science*, 13(3), 232-248.
- Hossain, L., & Wu, A. (2009). Communications network centrality correlates to organisational coordination. *International Journal of Project Management*, 27(8), 795-811.
- Huckman, R. S., Staats, B. R., & Upton, D. M. (2009). Team familiarity, role experience, and performance: Evidence from Indian software services. *Management Science*, 55(1), 85-100.
- Judd, S., Kearns, M., & Vorobeychik, Y. (2010). Behavioral dynamics and influence in networked coloring and consensus. *Proceedings of the National Academy of Sciences*, 107(34), 14978 -14982.
- Kauffman, S., & Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of theoretical Biology*, 128(1), 11-45.
- KC, Diwas, Staats, B. R., & Gino, F. (2013). Learning from My Success and from Others' Failure: Evidence from Minimally Invasive Cardiac Surgery. *Management Science*.
- Kearns, M., Judd, S., Tan, J., & Wortman, J. (2009). Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences*, 106(5), 1347-1352.
- Kearns, M, S. Suri and N. Montfort (2006) "An Experimental Study of the Coloring Problem on Human Subject Networks" *Science* 313: 824-827
- Kogut, B., & Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization science*, 3(3), 383-397.
- Kuhn, F., & Wattenhofer, R. (2006, July). On the complexity of distributed graph coloring. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing* (pp. 7-15). ACM.
- Lakoff, G (1987). *Women, Fire and Dangerous Things: What Categories Reveal about the Mind*. University of Chicago, Chicago.
- Lazer, David L. and Ethan S. Bernstein. Forthcoming, 2012. "Problem Solving and Search in Networks." In Todd, P. M. & Robbins, T. (Eds.), *Cognitive Search: Evolution, Algorithms, and the Brain*, Chapter 17, 261-273. Cambridge, MA: Strüngmann Forum Reports, MIT Press.
- Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4), 667-694.
- Leavitt, H. J. (1951). Some effects of certain communication patterns on group performance. *The Journal of Abnormal and Social Psychology*, 46(1), 38.
- Levitt, B. & March, J. G. (1988). Organizational learning. *Annual Review of Sociology*, 14, 319-340.
- Levine, S. S., & Prietula, M. J. (2012). How Knowledge Transfer Impacts Performance: A Multilevel Model of Benefits and Liabilities. *Organization*

- Science*, 23(6), 1748-1766.
- Ma, M., & Agarwal, R. (2007). Through a glass darkly: Information technology design, identity verification, and knowledge contribution in online communities. *Information Systems Research*, 18(1), 42-67.
- Madsen, P. M., & Desai, V. (2010). Failing to learn? The effects of failure and success on organizational learning in the global orbital launch vehicle industry. *Academy of Management Journal*, 53(3), 451-476.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, 2(1), 71-87.
- Marsh, S. J., & Stock, G. N. (2006). Creating dynamic capability: The role of intertemporal integration, knowledge retention, and interpretation. *Journal of Product Innovation Management*, 23(5), 422-436.
- Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, 137(3), 422.
- Mason, W., & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, 109(3), 764-769.
- Mulder, Mauk (1960) "Communication Structure, Decision Structure and Group Performance" *Sociometry* 23(1), 1-14.
- McCubbins, M. D., Paturi, R., & Weller, N. (2009). Connected coordination network structure and group coordination. *American Politics Research*, 37(5), 899-920.
- Newman, M. E. J. (2006) Modularity and community structure in networks *Proceedings of the National Academy of Sciences* 2006 103 (23) 8577-8582;
- O'Reilly III, C. A., Caldwell, D. F., & Barnett, W. P. (1989). Work group demography, social integration, and turnover. *Administrative science quarterly*, 21-37.
- O'Reilly, C. A. (1980). Individuals and information overload in organizations: is more necessarily better?. *Academy of Management Journal*, 23(4), 684-696.
- Quigley, N. R., Tesluk, P. E., Locke, E. A., & Bartol, K. M. (2007). A multilevel investigation of the motivational mechanisms underlying knowledge sharing and performance. *Organization Science*, 18(1), 71-88.
- R Core Team (2012). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.R-project.org/>.
- Raisch, S., & Birkinshaw, J. (2008). Organizational ambidexterity: Antecedents, outcomes, and moderators. *Journal of Management*, 34(3), 375-409.
- Raisch, S., Birkinshaw, J., Probst, G., & Tushman, M. L. (2009). Organizational ambidexterity: Balancing exploitation and exploration for sustained performance. *Organization Science*, 20(4), 685-695.
- Reagans, R., & McEvily, B. (2003). Network structure and knowledge transfer: The effects of cohesion and range. *Administrative Science Quarterly*, 48(2), 240-267.

- Reagans, R., & Zuckerman, E. W. (2001). Networks, diversity, and productivity: The social capital of corporate R&D teams. *Organization science*, 12(4), 502-517.
- Rivkin, J. W. (2000). Imitation of complex strategies. *Management Science*, 46(6), 824-844.
- Rivkin, J. W., & Siggelkow, N. (2003). Balancing search and stability: Interdependencies among elements of organizational design. *Management Science*, 49(3), 290-311.
- Robert, L. P., Dennis, A. R., & Ahuja, M. K. (2008). Social capital and knowledge integration in digitally enabled teams. *Information Systems Research*, 19(3), 314-334.
- Schneider, S. C. (1987). Information overload: Causes and consequences. *Human Systems Management*, 7(2), 143-153.
- Shalizi, C. R., & Thomas, A. C. (2011). Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods & Research*, 40(2), 211-239.
- Shaw, Marvin E. (1954). "Some effects of problem complexity upon problem solution efficiency in different communication nets." *Journal of Experimental Psychology* 48(3) 211-217
- Siggelkow, N., & Rivkin, J. W. (2005). Speed and search: Designing organizations for turbulence and complexity. *Organization Science*, 16(2), 101-122.
- Snijders, T. A. B. (2001), The Statistical Evaluation of Social Network Dynamics. *Sociological Methodology*, 31: 361–395. doi: 10.1111/0081-1750.00099
- Sparrowe, R. T., Liden, R. C., Wayne, S. J., & Kraimer, M. L. (2001). Social Networks and the Performance of Individuals and Groups. *Academy of management journal*, 44(2), 316-325.
- Speier, C., Valacich, J. S., & Vessey, I. (1999). The influence of task interruption on individual decision making: An information overload perspective. *Decision Sciences*, 30(2), 337 -360.
- Suleiman, J., & Watson, R. T. (2008). Social loafing in technology-supported teams. *Computer Supported Cooperative Work (CSCW)*, 17(4), 291-309.
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic management journal*, 18(7), 509-533.
- Travers, J., & Milgram, S. (1969). An experimental study of the small world problem. *Sociometry*, 425-443.
- Tsai, W. (2002). Social structure of "coopetition" within a multiunit organization: Coordination, competition, and intraorganizational knowledge sharing. *Organization science*, 13(2), 179-190.
- Tushman, M. L. (1979). Work characteristics and subunit communication structure: A contingency analysis. *Administrative Science Quarterly*, 82-98.
- Tushman, M., & O'Reilly, C. (1996). "Ambidextrous Organizations: managing evolutionary and revolutionary change". *California management review*, 38, 1-23.

- Von Krogh, G., Nonaka, I., & Rechsteiner, L. (2012). Leadership in organizational knowledge creation: a review and framework. *Journal of Management Studies*, 49(1), 240-277.
- Von Nordenflycht, A. (2010). What is a professional service firm? Toward a theory and taxonomy of knowledge-intensive firms. *Academy of Management Review*, 35(1), 155-174.
- Wasko, M. M., & Faraj, S. (2005). Why should I share? Examining social capital and knowledge contribution in electronic networks of practice. *MIS Quarterly*, 35-57.
- Watts, D. (1999). "Networks, dynamics, and the small-world phenomenon." *American Journal of Sociology*, 105: 493-527.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684), 440-442.
- Zollo, M., & Winter, S. G. (2002). Deliberate learning and the evolution of dynamic capabilities. *Organization science*, 13(3), 339-351.