17

Problem Solving and Search in Networks

David Lazer and Ethan S. Bernstein

Abstract

This chapter examines the role that networks play in facilitating or inhibiting search for solutions to problems at both the individual and collective levels. At the individual level, search in networks enables individuals to transport themselves to a very different location in the solution space than they could likely reach through isolated experimental or cognitive search. Research on networks suggests that (a) ties to diverse others provide a wider menu of choices and insights for individuals, and (b) strong ties will be relatively more useful for complex information, and weak ties for simple information. At the collective level, these conclusions become less clear. The key question is how the collective operates to coordinate within the group versus beyond it so as to balance experimentation and convergence towards a solution. Collective coordination of search, and collective evaluation of potential solutions, may significantly influence the optimal network structure for collective problem-solving search.

Introduction

Millions of problems go to work each day in search of solutions. The process of search, or "investigation of a question" (as defined in the Oxford English Dictionary), is in part defined by networks. While only a decade ago "problem solving" evoked images of Rodin's *Thinker*, we now think of Obama's Blackberry, IBM's smarter planet campaign, and project managers being able to "Google" all the brains of their organizations (Douglas, 2009). Although these are modern images, our capacity to solve complex problems based in part on the solutions of others is certainly a distinctive feature of human intelligence. When confronted with a problem, the search for a solution may happen in isolation, but it may also involve help from other human or non-human sources

accessible through a network of ties—that is, "networked search." In networked search, the network of sources from which help may be received thus defines problem solvers' access to potential pieces of a solution, whereas strategies to create and search networks define problem solvers' approach to traverse the path to connect those pieces. The construction of problem solving as a networked search, in turn, poses a series of critical questions across multiple levels of analysis, including: Where do people go to find answers? What are the collective, emergent consequences of those behaviors?

Other chapters in this volume highlight the role of asocial search—for example, how individuals search for visual patterns (Wolfe, this volume), or search memory for a relevant fact or word (Davelaar and Raaijmakers, this volume). The search of our networks is analytically distinct but part of the broader picture of search—indeed, as we discuss below, networked search and isolated search can often substitute for each other. Asking one's spouse if they remember where you put the keys may be a substitute for wracking one's memory for where you may have thrown them earlier. Conceptualizing human search as being in part a networked process also offers distinct practical implications: with improved understanding of how collaborative networks operate comes the opportunity, and challenge, to design networks for improved efficiency of networked, social search by individuals, groups, organizations, institutions, and communities.

Here we focus on network search as a core process of problem solving. We begin by providing a typology for understanding existing search research, categorized by types of search behavior. We then turn to our principal task of investigating similarities and paradoxes in network search theory across two levels of analysis: individual and collective. In the process of connecting what has been very disparate literature, our hope is to not only solidify the theory of search in networks but also to distill some important themes and opportunities for future research.

The Role of Networks in Problem Solving

Let us begin by envisioning "solutions" to a problem as a basket of activities, where, generally speaking, the permutations of possible activities are limitless. Searching through the space of possible solutions presents an extraordinary challenge, especially if one assumes (as we do) that synergies among activities are endemic—where, for example, activities A and B may be harmful singly but beneficial together. Given a very large solution space, with high levels of synergy among activities, incremental search (e.g., hill climbing without first determining the highest peak) will be a recipe for being stuck in a local optimum.

A Typology of Search

Search, at its broadest interpretation, may be classified based on the strategy used to navigate through a problem space. Figure 17.1 offers a typology of search useful for focusing our discussion across two common dimensions: whether search is isolated or networked (x-axis) and whether the information source searched is human or non-human (y-axis). Individually, these dimensions are common in relevant literatures; however, the combination is novel and results in four distinct (although combinable) categories of search: cognitive, experimental, social, or nonsocial.
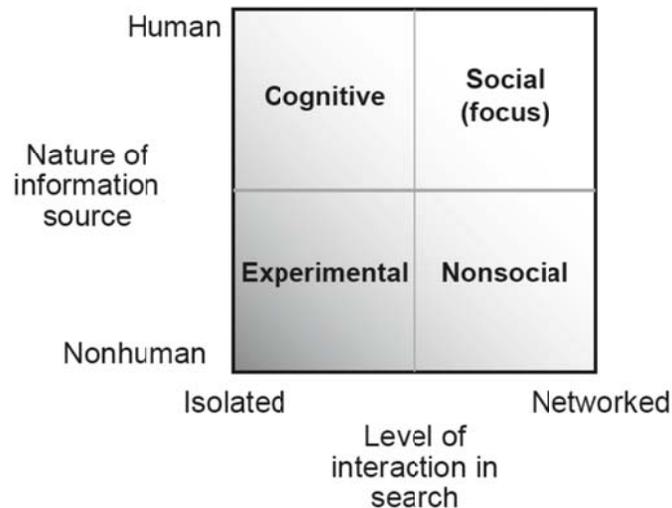
Figure 17.1 Typology of search across two dimensions. X-axis represents the level of interaction in search (isolated or networked); Y-axis indicates nature of the information source (human or nonhuman).

Isolated search may be *cognitive* when the information source is human (the individual searcher's mind) and the "hardware" is a brain—deriving solutions through mental models of how the world works which remain relatively consistent over time. Alternatively, isolated search may be *experimental*, involving interpreting feedback from attempted activities, adjusting hypotheses, and attempting an activity again, as foraging behavior may be characterized (e.g., McNamara and Fawcett, this volume). When search is networked rather than isolated, external sources of information may be consulted from either other humans (*social*) or nonhuman sources (*non-social*, e.g., files). One might imagine an individual searching for a Starbucks in a city. If they have been to that city, they might recall where a Starbucks is and conduct an isolated, human search based on their memory of that location (cognitive search). They might simply begin walking, assuming that there is a high density of Starbucks and that they will likely find one within a few blocks, circling around the focal location until they succeed in actually finding a Starbucks (experimental search). They might ask someone walking by where there is a Starbucks (social search). Alternatively, they might simply consult their smartphone (nonsocial search).

Using this typology, our first proposition would be to assert that most individuals, when search is isolated, will not make large changes in their solutions to problems they confront, in part because in a complex world, it is difficult or impossible to anticipate the impact of radical change. As in the model of Exploratory Information Foraging presented by Fu (this volume), search may begin without sufficiently precise criteria to judge the relevance of uncovered radical information to form a solution. If the objective of search were to thwart a terrorist attack, it may be hard to determine if increased activity in a suspected cell of a suspicious group is a relevant clue or a false lead. The social component (networked, human), however, offers the capacity for major change, with less risk, through observation of the activities of others. Relevance (or lack thereof) of radical new information to an optimal solution may be inferred from the choices of other human investigators proximate to you; that is, using others as prototypes allows searchers to make larger changes, based on radical new information, with lower risk.

Our focus here is thus on social search, where friendship, trust, belief, and expertise are all dyadic variables that have been examined as drivers of search in networks. Substantial research highlights positional factors that are likely to be related to successful search. For example, Burt (2004) has argued that a position of brokerage (i.e., knowing people who do not know each other) provides advantage by (a) providing ongoing flows of diverse information; (b) facilitating access to non-redundant information in extant cases; and (c) enhancing individual cognitive capacity (see also Burt, 1992).

Levels of Analysis

Social search may be executed by individuals or collectives. Following the vast majority of prior treatments of search, we begin our discussion with the individual level, where the impetus of social search emanates from a single individual. We then turn to the collective level, which becomes analytically relevant when outcomes are not simply the sum of individual efforts, but rather the result of some interaction among individual efforts (e.g., the purposeful coordination of a group or where there are informational spillovers from one individual to the next).

Taking our cues from prior literature, our discussion of the collective will begin by assuming the same theory of networked search operates at both the individual and collective levels, with the primary difference being the expansion of the locus of search impetus from a single individual to a collection of individuals. In effect, collective search simply shifts the boundary between "inside" and "outside": At the individual level, "inside" is an individual brain, and "outside" is everything else; at the collective level, the boundary expands what is "inside" to include multiple individuals (e.g., a group) with a shared set of networks on the "outside." As we progress, we identify what we believe may be at least one key theoretical tension between the individual and collective levels.

Locus of Problem Solving: Individuals

Individual network search, by definition, locates the impetus for problem solving squarely in the hands of one central individual (ego), who may draw on various baskets of activities, situated in diverse networks, to find solutions. Despite the substantial attention devoted to team-based problem solving over the past decade, the locus of most theory on network-related search is still the individual. Here we explore two tensions in the individual network search literature that lie at the heart of current research: creating versus conforming connections (related to number and strength of ties) and tacit versus explicit knowledge (content that flows across the ties).

Individual Creating Versus Conforming

In the quest for theories of performance in networked search, perhaps the most frequently studied tension is that between creation and conformity, exploration and exploitation, innovation and copying (March 1991). Search networks, for example, provide not only information, but also exert control, for example, through conformity pressures.

The visibility of one's behavior to others creates the opportunity for pressure on the individual to conform to alters' solutions, whether optimal or not. Agent-based simulation models have demonstrated that the more efficient a network is at disseminating information, the better the short-run performance but the worse the long-run performance of the system (Lazer and Freidman, 2007): connectedness encourages fast conformity at the expense of optimality. In Lazer and Friedman's (2007) model of parallel problem solving, a system is made up of a set of agents, each of whom is independently searching for answers. The performance of each agent is independent, in the sense that the performance of agent A has no direct bearing on any other agent, making this a set of individuals rather than a collective. Performance is, however, interdependent in the sense that there is a network connecting agents which allows them to observe the behaviors and performances of other agents (but not otherwise communicate). The essential conclusion from these simulations, as well as similar experimental studies (e.g., Mason et al., 2008), is that for *complex* problems, networks that were inefficient at disseminating information enabled a more thorough search of the problem space by agents, and thus better long-run performance by the system. For performance, the conformity imposed by connectedness was more troublesome than the creativity enabled by access was productive.

Such findings are consistent with an array of research that highlights the dangers of processes of rapid consolidation of individual theories and experimentation (Janis, 1972; Page, 2007), including McNamara and Fawcett's research on premature stopping behavior in this volume. We would expect these dangers to be particularly salient in the case of network closure (Uzzi, 1997),

where one's contacts can see each other. Conversely, we hypothesize that an "opaque" network, which provides more "spatial separation" (increased communication costs) between nodes (Duncan, 1976; Raisch and Birkinshaw, 2008), will therefore encourage experimentation, reduce copying, and lengthen exploration while limiting premature stopping.

We would also note that for a particular problem, there may be multiple useful paths to finding a solution, and a critical and understudied question involves which direction of search an individual takes. Binz-Scharf, Lazer, and Mergel (unpublished), for example, study how individuals in a DNA forensics laboratory search for answers to problems they encounter. In this study, a wide range of sources are utilized, ranging from nonhuman sources (manuals, journals) to institutional support (a help desk for software) to social resources (friends). Further, people often use a distinctive sequence in their search: some, for example, will thoroughly search nonhuman sources first, because no reputational consequences are at stake for getting an answer from a journal, or Google, whereas asking for help from a person could entail loss of face. Given the path dependencies of many answers, study of the behaviors that drive directions for search is an area of promise for future work.

Knowledge Transfer for Simple Versus Complex Problems
The literature on knowledge transfer has distinguished between tacit and explicit knowledge. Explicit knowledge is knowledge that is easy to codify (Nonaka, 1994; Gavetti and Levinthal, 2000; Edmondson et. al., 2003), such as directions to a restaurant (at least in most geographies). Tacit knowledge is knowledge that is difficult to codify, because of its complexity or contingent nature (e.g., an answer that begins with "it depends" likely tends toward the tacit end of the scale).

Because it is easy to codify, explicit knowledge is more likely accessible through nonhuman sources of information, such as a reference manual or materials that could be found, for example, through Google. Alternatively, one could consult with an individual or set of individuals about possible answers, where weak ties will likely suffice in providing an answer. Even if an answer is not provided, a reliable path to the optimal answer—a well-tested routine which has proven to be a fruitful path to find an answer (Nelson and Winter, 1982)—may be provided such that the individual has a stable roadmap to the solution, like a treasure map, that she need only execute to succeed.

For tacit knowledge, nonsocial sources of information become less useful because (by definition) tacit knowledge cannot easily be formally represented. Less trivially, strong ties are particularly important to transfer tacit knowledge (Hansen, 1999). The reason for this is that transferring tacit knowledge smoothly is likely costly, requiring that both actors have background understanding of each other and speak a similar (and similarly situated) language (Bechky, 2003).

These requirements are presumably more likely given strong ties, and thus transfer of tacit knowledge is eased when embedded in a broader set of exchanges between two individuals.

Individual Social Search: Summary

Even our very limited treatment here is sufficient to distill one key issue in individual network search: namely, the complexity of the problem space may dictate the characteristics of an optimal network. We return to this point in more detail below (see section on Discussion and Implications).

<div align="center">

Locus of Problem Solving: Collective

</div>

The collective level is relevant if collective-level consequences result from how individuals are connected together. This might be the case where there is a functional interdependence among the activities of different individuals (e.g., the value produced by activity one by person A depends on whether person B engages in activity two). It would also be the case where there is informational interdependence among actors: person A learns something and transmits it to person B.

There are many constructions of the collective in the social sciences. The literature on groups focuses on small sets of people (typically less than a dozen) with well-defined boundaries, usually structured around some homogeneous, shared purpose. The literature on organizations generally focuses on formal bureaucratic structures, often structured around heterogeneous purposes, on a scale of hundreds or thousands. The literature on communities, broadly construed, can span collectives from thousands to billions. Given the nested nature of these constructs, we focus largely on problem solving in the most fundamental form of collective: groups.

We define a group, following Alderfer (1977), as "an intact social system, complete with boundaries, interdependence for some shared purpose, and differentiated members"—that which Hackman (2012) refers to as "*purposive groups*—that is real groups that exist to accomplish something." This definition incorporates two key identifiers which distinguish a group from other collections of individuals (Hackman, 2012):

1. Members can be distinguished from non-members, by both members and non-members; and
2. Members depend on each other to achieve a collective purpose, accepting specialized roles in the process.

Although other definitions of groups may differ, because of our focus on search as a form of problem solving, we follow Hackman (2012) in excluding "casual gatherings…, reference groups, identity groups, and statistical aggregations of the attributes, estimates, or preferences of people who do not actually interact with one another." We would argue that this definition holds even in an age where groups may never meet face-to-face and may only stay together for a limited time, such

as distributed teams in large organizations (Hackman and Wageman, 2005; O'Leary and Cummings, 2007; Hackman, 2012).

From Individual to Group

The key conceptual question about network search by groups is how one aggregates from the individual to collective. We begin with the proposition that the theory of networked search operating at the individual level remains consistent when analyzing the collective level—only the boundary between "inside" and "outside" shifts outwards to include multiple individuals on the "inside." In other words, with collective problem solving, the search for the best basket of activities to form a solution is distributed among a set of individuals; there is a defined division of labor and rewards for reaching a solution that are distributed across members of the group (although not necessarily equally). The key questions then become:

- How are these network search tasks coordinated inside the group to yield performance?
- Does the relationship between external network structure and search performance change as a result of the actor being a collective rather than an individual?

Coordination of Search Tasks

Many, perhaps even most, important complex problems are not solvable in an efficient manner by an individual because the scale of the effort may be too great, or the scope of the skills required too broad. It is for precisely this reason that, when individuals are the locus of search, they often supplement their own cognitive and experimental search with network search to access solutions and capabilities held by others. Such an individual-centric model for collective search by humans faces, however, a key limitation: failure to account for *coordinated* search through a problem space.

For example, if we take the basic parallel problem solving paradigm (a set of agents, all of whom are working on the same problem, with independent payoffs), but allow agents to communicate about *how* to search through the problem space, as a group would be expected to do, collective search behavior might change dramatically, especially if rewards for finding a solution were shared by the group. It is possible that a group might decide to diversify behaviors so as to make collective search more thorough or to focus search on what are seen as promising areas of the problem space. Thus, within the parallel problem solving framework, the question is: How does the network affect how groups search through a problem space? One might also ask: How does the network affect how groups *decide* to search through a problem space? We are not familiar with research directly on point; partially relevant is the work on the performance implications of transactive memory (e.g., Wegner, 1987; Liang, Moreland, and Argote, 1995; Austin, 2003; Brandon and Hollingshead, 2004) and team familiarity (Huckman, Staats, and Upton, 2009), which

focuses on how people learn to work with specific others, as well as work on self-organization in groups (e.g., Trist et al, 1963; Barker, 1993; Arrow and Burns, 2003; Arrow and Crosson, 2003).

The general set of questions around division of labor and coordination of individual efforts in a collective transcends issues around collective problem solving. Many problems require some division of labor, splitting the problem into sub-problems, each of which is in turn solved by individuals (or smaller groups). Some activities may require efforts by multiple individuals with special and mutually exclusive skills. There is a vast literature in organizational theory on process and coordination (e.g., Mintzberg, 1979; Schein, 1985, 1987; Hackman, 2012). Relatively little of this literature addresses problem solving per se; most focuses on execution of well-defined, if sometimes complex, tasks (e.g., how to create an effective assembly process). A full mapping of how this literature might apply to problem-solving search is beyond the scope of this chapter. Key questions we would highlight include:

- *How does performance of individuals map to the performance of the whole?* Steiner (1976) offers a particularly useful typology: some search tasks are additive, essentially the sum of the contributions of every member of the group; others are disjunctive (only the best performer matters) or conjunctive (only the worst performer matters).

- *What is the structure of interdependence among individuals?* Some search tasks require synchronous, coordinated action among different agents; others require asynchronous action. Creating a Wikipedia page with a compendium of facts about some notable individual, for example, requires little coordination; contributors can simply assess what is missing at a given point in time and fill gaps. Investigation of a crime, however, requires coordinated action among the involved investigators.

Both well-functioning groups and networks facilitate coordination, and thus a simple hypothesis would be that higher levels of interdependence among agents require denser networks between those agents. There are, however, many mechanisms to facilitate coordination beyond networks in human systems. Standardization, for example, is one major mechanism for coordinating behavior without communicating. The need to communicate at a given moment to allow for synchronized action is eliminated by our ability to track time accurately and the convergence of particular conventions around time keeping. The question regarding the role of networks thus becomes one of the (sometimes large) residual: given the other mechanisms for coordination, what network structures, both inside and outside of the collective, support group performance in solving a problem?

Relationship between Network Structure and Search Performance for Collectives

Just as was the case for individuals, a blanket assertion that more and stronger ties would be better is clearly not the right answer for the collective.

"Inside"

With regard to the network inside the group, consider the classic Asch experiment (1956), in which the choices of subjects were visible to the other members of the group and they thus conformed to the (false) group norm. One difference between individual and collective networked search lies on the "inside" of the collective: collectives do not always have access to all of the knowledge within the collective the way that individuals ordinarily do. Put differently, collectives "forget" a lot more than individuals do. This finding comes from the particularly substantial thread of related research focused on the issue of information aggregation within groups, especially on "hidden profile" tasks (Stasser & Titus, 1985, 2003). In a hidden profile task, information is distributed among group members—some of which is redundant, some of which is (privately) held by single individuals—and the group is searching for a "right" answer that requires individuals to combine their privately held information. The robust, and paradoxical, finding of this research is that despite incentives to maximize group performance, individuals will tend to focus their discussion on commonly held information, and not discuss (or reveal) information that is privately held, even though that information is necessary for group success. This has led to substantial research on the conditions that will lead individuals to reveal the information that they alone have (Sunstein, 2006).

The hidden profile paradigm is based on group discussion; however, an older vein of research, which came out of the Small Group Network Laboratory at MIT in the 1950s (Bavelas, 1950; Leavitt, 1951; also see Guetzkow and Simon, 1955), examined information aggregation in the context of distributed information in networks. In this research, information would be distributed in a group, where members each had a signal about the state of the world. Successful answering of the problem by the group required pooling all of these signals together, and disseminating the right answer to the entire group. Individuals were connected to a subset of the entire group and could pass a signal on to one of their contacts. The key question was: what network structure facilitated group success? The robust answer was that centralized networks worked best for simple problems, whereas decentralized networks functioned best for complex problems that required more individual effort.

Neither of the above research paradigms, however, incorporated the idea of individual experimentation (i.e., individuals might proactively seek information that the group does not have) through coordinated search via external networks. Therefore, we turn our attention to the other side of the collective boundary.

"Outside"

Just as with individuals, groups that are well-connected to external networks run a risk of suffering reduced performance: Bernstein (2012) demonstrates, through a field experiment in a manufacturing context, that a modest increase in group-level privacy (reduced observability through stronger group boundaries) can improve, sustainably and significantly, assembly-line performance by as much as 10-15% on a simple assembly task, by supporting productive deviance, localized experimentation, distraction avoidance, and continuous improvement.

In the context of problem solving, experimentation necessarily requires nonconformity; visibility through the network (whether internal or external to the collective) may therefore stymie new behaviors. Given that perhaps the dominant small group unit is, even today, the family (i.e., a group made up of people with shared genetics), we speculate as to whether certain individual behaviors that might be viewed as dysfunctional (e.g., stubbornness) might actually be beneficial at the group level because they maintain diversity within the group, yielding greater group success and thus improved propagation of genes.

While the importance of maintaining access to diverse perspectives was also relevant at the individual level, the question of how network structure influences performance at the collective level is complicated by the fact that the collective has both an internal and external component to its network. We believe this complication calls for more research to build a more nuanced theory of networked search by collectives, a recommendation we explore in the next section.

Discussion and Implications

Two focal points have been identified for thinking about search in networks and problem solving: the individual and the collective. At the individual level, network research highlights the value of having diverse information sources, as well as strong ties for complex information and weak ties for simple information. At the collective level, we argue that problem complexity requires an external network structure that slows the consolidation processes in the system (see Figure 17.2) given the internal tendencies of groups.
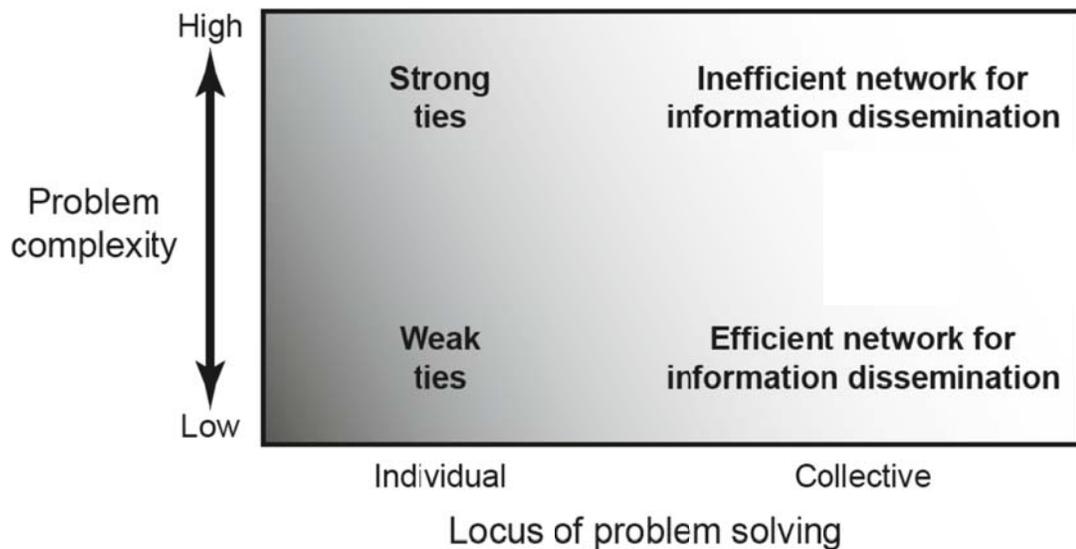
Figure 17.2 Preconditions for successful search in external networks.

There is a potential tension between what is optimal at these two levels. As Lazer and Friedman (2006) describe, when groups confront complex problems, there is a potential "tragedy of the network": individuals have an interest in not engaging in costly experimentation, but in positioning themselves to receive information from others quickly, whereas the group has an interest in individual experimentation and putting the brakes on information spread.

Research into that tension is made more interesting, and more urgent, in light of modern information and communication technologies, which offer the potential for vastly more efficient communication outside of collectives. Bernstein's (2012) finding of substantial value to group-level privacy, even in contexts like simple assembly manufacturing where complete transparency and observability is the standard, suggests the possibility of rather perverse consequences to rewiring the network of communication to make it more efficient. Indeed, given that most of human existence has been under conditions of much more limited global communication, this raises the question of whether there is a fundamental misalignment between individual psychology (which is to work on problems together) and collective outcomes in an age where the technical limits to working together and learning from one another have been radically reduced.

The more nuanced question regarding modern information and communication technologies is how the mechanisms for global coordination interplay with those of local coordination. Nonhuman foraging, with some exceptions, typically allows for just local coordination, because the means to communicate non-locally is lacking. Thus, a key analytic question must focus on the global logic of purely local search behaviors. If one is looking at swarming, for example, key questions become: how do agents distribute themselves over a landscape, based on local spacing decisions?; and how does the swarm allocate enough agents to resource retrieval when a resource is found? Non-local communication is, however, eminently possible for humans. There are

mechanisms to apprise a whole population of certain facts, and shortcuts that rapidly diffuse relevant information.

While we only raise the questions here, a useful avenue for further research would be on the interplay of these different levels of mechanisms to coordinate search, particularly in this contrast between the operation of theory at the individual and collective levels. One potential flaw in a hive/swarm metaphor for human behavior is that information is constantly flowing into groups, in part because most individuals belong to many distinct groups. There is a natural osmosis of information among groups, where groups actively manage ties to the external environment as necessary to help in problem solving. These observations point to the need for study of search in networks with many methodologies, as individual and collective research tend to be executed in different contexts with different methods. Classic experimental work offers great promise, for example, for studying how people choose to balance experimentation and exploration with manipulated network structures. However, it is equally clear that problem solving in organizations, families, and societies, comes with contexts that interplay and structure the directions that people can search. This highlighting the need for field research that uses replicable methods systematically across settings to evaluate what patterns of search are robust, and which patterns interplay with particular contextual factors.

Conclusion

In concluding, we would like to step beyond future research possibilities to offer one hypothesis we find particularly intriguing. A common theme throughout this chapter has been the importance of complexity in influencing how increasing connectedness will impact performance. The literature suggests, as presented in Figure 17.2 at the individual level, that more complexity requires more connectedness. This seems consistent with common perceptions of the world today: the world's problems have become more complex, but our ability to deal with complex problems has improved with modern technologies which permit substantially better connectedness for network search. The anomaly, however, is in the increased prevalence of teams. As network search has become more efficient, current theory would suggest that groups would become less important, not more—that the network has become so powerful that individuals could harness it for problem solving without the need for well-performing teams, which are neither easy nor costless to build. Why, then, in the age of Google, Facebook, e-mail, texting, instant messaging, and costless phone calls, is there apparently increased reliance on teams as a key unit of production (Arrow and McGrath, 1995; Hackman, 2002; Edmondson and Nembhard, 2009; Hackman, 2012)?

One possibility is that it is precisely *because* network search has become so much more powerful that teams have become more prevalent. By giving the group some power to jointly

balance "inside" and "outside" activities, collective search solves many of the counterproductive aspects of increasingly powerful networks in networked search. The construction of boundaries may buffer individuals from outside control, allowing a more deliberative, exploratory space within the group. If that were the case, in a world where ever escalating connectivity enables exploitation of what is collectively known, teams would be increasingly important as instruments of exploration.

References

Alderfer, C. P. 1977. Group and intergroup relations. In: Improving Life at Work, ed. J. R. Hackman and J. L. Suttle, pp. 227–296. Santa Monica: Goodyear.

Arrow, H., and K. L. Burns. 2003. Self-organizing culture: How norms emerge in small groups. In: The Psychological Foundations of Culture, ed. M. Schaller and C. Crandall, pp. 171–199. Mahwah, NJ: Lawrence Erlbaum.

Arrow, H., and S. Crosson. 2003. Musical chairs: Membership dynamics in self-organized group formation. *Small Group Res.* **5**:523–556.

Arrow, H., and J. E. McGrath. 1995. Membership dynamics in groups at work: A theoretical framework. In: Research in Organizational Behavior, ed. B. M. Staw and L. L. Cummings, pp. 373–411. Greenwich, CT: JAI.

Asch, S. E. 1956. Studies of independence and conformity: A minority of one against a unanimous majority. *Psychol. Monogr.* **70(9)**:1–70.

Austin, J. R. 2003. Transactive memory in organizational groups: The effects of content, consensus, specialization, and accuracy on group performance. *J. Appl. Psychol.* **88(5)**:866–878.

Barker, J. R. 1993. Tightening the iron cage. Concertive control in self-managing teams. *Adm. Sci. Q.* **38**:408–437.

Bavelas, A. 1950. Communication patterns in task oriented groups. *J. Acoust. Soc. Am.* **57**:271–282.

Bechky, B. A. 2003. Sharing meaning across occupational communities: The transformation of understanding on a production floor. *Organ. Sci.* **14(3)**:312–330.

Bernstein, E. S. (2012). The transparency paradox: a role for privacy in organizational learning and operational control. *Adm. Sci. Q.* **57**:181-216.

Binz-Scharf, M. C., Lazer, D., & Mergel, I. (unpublished). Searching for answers: networks of practice among public administrators. American Review of Public Administration.

Brandon, D. P., and A. B. Hollingshead. 2004. Transactive memory systems in organizations: Matching tasks, expertise, and people. *Organ. Sci.* **15(6)**:633–644.

Burt, R. S. 1992. Structural Holes: The Social Structure of Competition. Cambridge, MA: Harvard Univ. Press.

Burt, R. S. (2004). Structural holes and good ideas. *Am. J. Soc.* **110(2)**:349–399.

Douglas, P. 2009. What If We Could Google Our Own Brains? http://www.techradar.com/news/world-of-tech/what-if-we-could-google-our-own-brains-533445. (accessed 26 Oct. 2011).

Duncan, R. B. 1976. The ambidextrous organization: Designing dual structures for innovation. In: The Management of Organization Design, ed. R. H. Kilmann et al., pp. 167–188. New York: North Holland.

Edmondson, A. C., and I. M. Nembhard. 2009. Product development and learning in project teams: The challenges are the benefits. *J. Prod. Innovat. Manag.* **26(2)**:123–138.

Edmondson, A. C., G. Pisano, R. M. J. Bohmer, and A. Winslow. 2003. Learning how and learning what: Effects of tacit and codified knowledge on performance improvement following technology adoption. *Decision Sciences* **34(2)**:197–223.

Gavetti, G., and D. Levinthal. 2000. Looking forward and looking backward: Cognitive and experiential search. *Adm. Sci. Q.* **45(1)**:113–137.

Guetzkow, H., and H. A. Simon. 1955. The impact of certain communication nets upon organization and performance in task-oriented groups. *Management Science* **1(3–4)**:233–250.

Hackman, J. R. 2012. From causes to conditions in group research. *Journal of Organizational Behavior* **33**:428-444.

Hackman, J. R. 2002. Leading Teams: Setting the Stage for Great Performances. Boston: Harvard Business School Press.

Hackman, J. R., and R. Wageman. 2005. When and how team leaders matter. In: Research in Organizational Behavior, ed. B. M. Staw and R. D. Kramer. New York: Elsevier.

Hansen, M. T. 1999. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Adm. Sci. Q.* **44(1)**:82–111.

Huckman, R. S., B. R. Staats, and D. M. Upton. 2009. Team familiarity, role experience, and performance: Evidence from Indian software services. *Management Science* **55(1)**:85–100.

Janis, I. L. 1972. Victims of Groupthink: A Psychological Study of Foreign-Policy Decisions and Fiascoes. Boston: Houghton Mifflin.

Lazer, D., and A. Friedman. 2006. The Tragedy of the Network. In: Intl. Sunbelt Social Network Conf. Vancouver: INSNA.

Lazer, D., & Friedman, A. 2007. The network structure of exploration and exploitation. *Adm. Sci. Q.* **52**:667–694.

Leavitt, H. J. 1951. Some effects of certain communication patterns on group performance. *J. Abnorm. Soc. Psychol.* **48**:38–50.

Levinthal, D., & March, J. G. (1981). A model of adaptive organizational search. *Journal of Economic Behavior and Organization* **2**:307–333.

Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal* **14**:95-112.

Liang, D. W., Moreland, R., & Argote, L. (1995). Group versus individual training and group performance: the mediating role of transactive memory. *Personality and Social Psychology Bulletin* **21(4)**:384-393.

March, J. G. 1991. Exploration and exploitation in organizational learning. *Organ. Sci.* **2(1)**:71–87.

Mason, W. A., A. Jones, and R. L. Goldstone. 2008. Propagation of innovations in networked groups. *J. Exp. Psychol. Gen.* **137(3)**:422–433.

Mintzberg, H. 1979. The Structuring of Organizations: A Synthesis of the Research. Englewood Cliffs, NJ: Prentice-Hall.

Nelson, R. R., and S. G. Winter. 1982. An Evolutionary Theory of Economic Change. Cambridge, MA: Harvard Univ. Press.

Nonaka, I. 1994. A dynamic theory of organizational knowledge creation. *Organ. Sci.* **5(1)**:14–37.

O'Leary, M. B., and J. N. Cummings. 2007. The spatial, temporal, and configurational characteristics of geographic dispersion in teams. *MIS Quarterly* **31**:433–452.

Page, S. E. 2007. The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies. Princeton: Princeton Univ. Press.

Raisch, S., and J. Birkinshaw. 2008. Organizational ambidexterity: Antecedents, outcomes, and moderators. *J. Manag.* **34(3)**:375–409.

Schein, E. H. 1985. Organizational Culture and Leadership. San Francisco: Jossey-Bass.

Schein, E. H. 1987. Process Consultation. Reading, MA: Addison-Wesley.

Stasser, G., and W. Titus. 1985. Pooling of unshared information in group decision making: Biased information sampling during discussion. *J. Pers. Soc. Psychol.* **57**:67–78.

Stasser, G., and Titus, W. 2003. Hidden profiles: A brief history. *Psychol. Inq.* **14(3–4)**:304–313.

Steiner, I. 1976. Task-performing groups. In: Contemporary Topics in Social Psychology, ed. J. Thibaut et al., pp. 393–421. Englewood Cliffs, NJ: General Learning Press.

Sunstein, C. R. 2006. Infotopia: How Many Minds Produce Knowledge. New York: Oxford Univ. Press.

Trist, E., G. Higgin, H. Murray, and A. Pollock. 1963. Organisational Choice: The Loss, Rediscovery and Transformation of a Work Tradition. London: Tavistock.

Uzzi, B. 1997. Social structure and competition in interfirm networks: The paradox of embeddedness. *Adm. Sci. Q.* **42(1)**

Wegner, D. M. 1987. Transactive memory: A contemporary analysis of the group mind. In: Theories of Group Behavior, ed. B. Mullen and G. Goethals, pp. 185–208. New York: Springer Verlag.