

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Innovation Policy and the Economy, Volume 7

Volume Author/Editor: Josh Lerner and Scott Stern, editors

Volume Publisher: MIT Press

Volume ISBN: 0-262-10121-1

Volume URL: <http://www.nber.org/books/jaff07-1>

Conference Date: April 19, 2006

Publication Date: February 2007

Chapter Title: Innovating under Pressure - Towards a Science of Crisis Management

Chapter Author: Daniel Diermeier, Wallace J. Hopp, Seyed Iravani

Chapter URL: <http://www.nber.org/chapters/c0036>

Chapter pages in book: (p. 125 - 154)

Innovating under Pressure—Towards a Science of Crisis Management

Daniel Diermeier, *Northwestern University*

Wallace J. Hopp, *Northwestern University*

Seyed Iravani, *Northwestern University*

Executive Summary

We propose a rigorous modeling framework for characterizing the structural ability of organizations to respond quickly and effectively to unanticipated events. As such, we seek to provide a theoretical basis for improved crisis management strategies. Our framework conceptualizes organizations as adaptive, responsive networks. Most of the existing models of complex social networks to date, however, have not explicitly modeled human capacity constraints or system congestion. As a result, no viable frameworks exist for investigating the responsiveness of various organizational structures under crisis conditions. Our approach proposes to integrate the social network approach to modeling communication and collaboration with the flow network approach from production systems modeling to represent task processing and flow under crisis conditions. By providing analytic structure to decision making environments currently viewed as not amenable to formal methods, this research, we hope, will help improve the performance of various organizations in both the private and public sectors.

I. Introduction

Crises are an increasingly common feature of modern life. The Munich Olympics, Tylenol tampering, Exxon Valdez, Lockerbie, Asian Currency Crisis, Ford Explorer rollovers, September 11, Northeast blackout, Enron, Worldcom, Indian Ocean tsunami, Hurricane Katrina—hardly a day goes by without some sort of crisis in the news. Of course, this is partly media marketing—tragedy generates better ratings than triumph. But there is no question that the crises are real and involve enormous consequences in human life, financial costs, environmental impacts, and social disruption. As a result, effective crisis management is a challenge that eventually faces most leaders of public and private organizations.

While natural disasters or terrorist attacks are most likely to capture the imagination of the reading and viewing public, crisis management has a much wider scope. Public sector examples range from black-outs or the collapse of critical infra-structure components such as water supply or 911 call systems, to the performance of public services ranging from hospitals, .e.g., in the case of pandemics, to police and fire departments, and to the military.

Unfortunately, while routine day-to-day management has been supported by a steadily systematized body of knowledge, consisting of economic theory, statistical tools, operational models, behavioral science, and results from other disciplines, crisis management remains largely anecdotal and *ad hoc*. As a result, text books and academic courses on the subject are almost exclusively case based and practitioners rely on benchmarking, checklists, and contingency plans. While cases are valuable for building initial insights, they can only go so far in helping to prepare for future events that never fully duplicate the past. Similarly, benchmarks and checklists can help one organization borrow ideas from one another, but cannot create new insights.

This leaves contingency planning, which is by far the most common approach for preparing an organization to manage a crisis. For example, most police, fire, and emergency medical systems now have plans for responding to attacks involving explosive devices and biological agents. The EPA has a plan for responding to an emergency caused by the release of hazardous substances. Emergency rooms have specified protocols for handling specific categories of emergencies, such as chemical agent hazards, and airlines and oil companies have specific protocols for responding to crashes and spills. Moreover, since 9/11, many corporations have hired Chief Security Officers, for whom contingency planning is a major responsibility.

A contingency plan is a reasonable preparation tool for situations in which enough of the future can be foreseen to facilitate a plan. For example, in 2004, FEMA ran a simulation of a fictional Hurricane Pam that predicted the consequences of Hurricane Katrina with disturbing accuracy. Tragically, many of the preparations indicated by the Pam exercise were never made and the response to Katrina was widely criticized. However, even the flawed contingency planning may have done some good; the actual death toll of around 1,400 fell far short of the Pam prediction of 60,000.

But contingency planning is often impossible or impractical. What reasonable fire department would have had a plan for managing the

aftermath of planes crashing into the World Trade Center? Could NASA have drawn up an advance plan for the specific circumstances of Apollo 13? Can we truly fault a Danish dairy company for failing to prepare contingency plans for a dramatic drop in sales triggered by series of cartoons in a newspaper? Should we expect an emergency room to be prepared to be taken hostage by a frustrated father demanding a heart transplant for his son? Admittedly, this last scenario is the plot of a Denzel Washington movie, but is it really any stranger or more shocking than the real crises cited here? The simple fact is that despite an organization's best efforts there are infinitely many scenarios for crises and so it is impossible to prepare in advance for them all.

An instructive example is the preparation of the Tulane University Hospital and Clinic for Hurricane Katrina (Naik 2005). In contrast to many other hospitals, this facility, run by HCA Inc., had extensively studied the specific challenges faced by hospitals in hurricanes and incorporated the insights into crisis preparation and contingency planning. For example, HCA provided satellite phones and back-up generators and stored large quantities of hospital supplies. Despite extensive planning, management had to make bold decisions under extreme time pressure. For example, when the levees collapsed, senior management quickly decided on an evacuation strategy leasing a motley collection of about 20 privately held helicopters for the evacuation of patients and staff, ranging from a privately owned Blackhawk to a Russian made helicopter leased from an owner in Panama City, FL. To make the air-lifts possible even at night a make-shift landing zone was created, illuminated by car headlights. Moreover, management created an ad hoc air traffic control system using amateur ham-radio operators.

Sheffi (2005) (and Fink (2002) before him) used the ubiquitous 2x2 matrix format, which we summarize in figure 5.1 to characterize organizational risks and to classify strategies for mitigating them.

The main insight from this representation is that preparing for emergency situations cannot be done with a "one size fits all" strategy. For events with sufficiently high likelihood of occurrence, it makes sense to build in redundancies or other forms of proactive protection. For example, supply chains may hold safety stock or build in safety lead time as protection against weather related shipment delays or disruptions. Astronauts now carry material with which to repair damaged tiles on the Space Shuttle. Emergency rooms have various specialists available on call to handle types of medical emergencies that exceed the capabilities of the on-site staff.

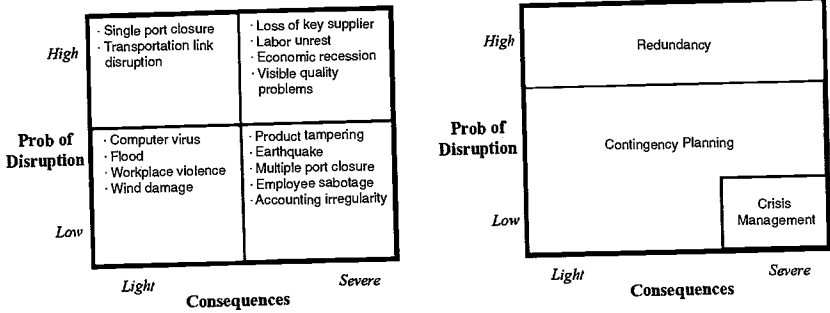


Figure 5.1
Examples of corporate disruptions (left); Classification of response strategies (right).

But it is not economical or even feasible to carry redundancies for unlikely events. Holding safety stock to protect against a four month disruption in supply resulting from a plant being hit by lightning, or keeping a tropical disease specialist on call to handle unusual emergency room cases, would be prohibitively expensive. So for these less likely, but still predictable, scenarios we rely on contingency planning. For instance, a manufacturer might maintain a list of backup suppliers to contact in case of a disruption. An emergency room might have a network of specialists to consult about rare medical cases.

For truly rare but highly consequential situations, such as product tampering scenarios, attacks on supply chains, unusual medical emergencies, catastrophic mission failures, and many others we fail to imagine until they occur, there is no alternative to crisis management. The key question is how to ensure that such management is a well-orchestrated, highly creative success (Apollo 13) rather than an uncoordinated, ineffectual mess (Katrina).

The main thesis of this essay is that organizations can be predisposed for success in crisis situations by both promoting individual skills *and* structuring collaborative relationships. That is, instead of preparing an organization for specific scenarios, as is done in contingency planning, we feel it is possible to help it prepare for anything by cultivating an ability to respond quickly and adaptively to unfamiliar situations. Moreover, we believe that it is feasible to establish formal metrics of responsiveness and adaptability that can be used to assess an organization's crisis management capabilities. Following this line of reasoning, crisis responsiveness is conceptualized as an organization's ability to find solutions to unanticipated events in high-pressure situations. In

this sense it constitutes a specific form of innovation, but under intense time pressure and with very high stakes. The goal then is to develop a formal framework that can help identify why certain organizations are better at finding solutions than others.

At this point the reader may ask what exactly distinguishes crisis management from other forms of management. Perhaps crisis management is simply management under particularly high stakes, but not a qualitatively different phenomenon and therefore may not deserve a research strategy distinct from the general study of management. Even if crisis management requires rapid innovation, it could still be viewed as a special case of innovation management.

However, even a casual observation of crisis behavior suggests that the existing management literature is ill-suited to study crisis management. Much of the modern management literature focuses on contracts, compensation, and other forms of incentive systems. Yet, it is highly implausible that the excellent performance of New York City's emergency rescue personnel after 9/11 was due to well-chosen incentive schemes or that their performance could have been further enhanced by modifying incentive schemes. The same argument applies to observed differences in performance among military units with the same pay structure. These casual observations suggest that the standard tool-kit of economic analysis (principal-agent theory, theory of contracts, etc.), at least *prima facie*, offers little promise for the study of crisis management. We, therefore, need to look elsewhere for a suitable framework.

Developing the needed models will require a representation of how organizations develop innovative solutions and share knowledge to solve problems. This is a question that has drawn intense interest from both researchers and practitioners in recent years. According to Kogut and Zander (1992), organizations function as "social communities in which individual and social expertise is transformed into economically useful products and services by the application of a set of higher-order organizing principles." In such communities, social connections among individuals form a network through which people share information in order to complete tasks and create new knowledge.

In industry, some large companies, including IBM, HP, and Intel, have begun making use of a special form of knowledge networks, called Communities of Practice (CoP). Promoting learning among members of CoPs has enabled IBM to greatly "decrease the learning curve of new employees, respond more rapidly to customer needs and inquiries, reduce rework and spawn new ideas for products and ser-

vices" (Lesser and Storck 2001). HP has made use of similar methods to implement knowledge management through what they call knowledge communities (Huberman and Hogg 1995). Intel has installed and nurtured knowledge networks to expand their access to university faculty experts and thereby improve the productivity of corporate R&D activities (Chesbrough 2003).

The modeling strategy pursued in this essay utilizes the theory of complex social networks. The idea here is to formally model organizations as knowledge networks as suggested by researchers from various disciplines who have started to explore knowledge and social networks and their effects on organizational performance (e.g., Albrecht and Ropp 1984; Stevenson and Gilly 1991; Wasserman and Faust 1994) and organizational innovations (e.g., Burt 2004; Feldman 1999; Kilduff and Tsai 2003; Kogut 2000; Monge and Contractor 2003). This work has concentrated on the role of node position in the knowledge network (e.g., Burt 2004), strength of ties (e.g., Granovetter 1973), and network evolution (e.g., Fonbrun 1986; Brass 1995). Knowledge management scholars have also examined the relation between informal networks and innovation. Growing interest in using networks to study social systems recently has attracted physical and mathematical scientists to adapt models of physical networks to social and organizational networks. In the next section we give a brief overview on how such models can be used to model knowledge creation and sharing in organizations.

II. Knowledge Networks

A network is a system of nodes with connecting links. The modern theory of complex networks was heavily influenced by sociology and social psychology (e.g., Milgram 1967), which also provided some of the key terminology such as "small world" networks "six degrees of separation" or "cliques." Married with applications from biology and the suitably modified mathematics of random networks, these ideas gave birth to the theory of complex networks (e.g., Watts and Strogatz 1998; Watts 1999; Newman, Strogatz, and Watts 2001; Albert et al. 1999). Once one adopts this viewpoint, networks were found in many different domains (e.g., Barabasi 2002). Individuals exchanging e-mails is one example. Person A sends an e-mail to B; if B replies, A and B are connected. Other clear-cut examples are the Internet, a network of servers, and the World Wide Web, a network of web pages connected by hyperlinks (Albert et al. 1999). Depending on the applications networks can be modeled as directed or undirected.

What makes networks "complex" is a surprising similarity to the literature of critical phenomena (e.g., Stanley 1999). Real networks, whether capturing social, economic, or biological systems, appear to exhibit universal properties that are independent of the specific form of the interactions. Moreover, these similarities can be described by the same mathematical formalism as used in the study of critical phenomena (e.g., Albert et al. 1999).

A recurrent characteristic of real networks is the small-world phenomenon, which is defined by the co-existence of two apparently incompatible conditions, (1) the number of intermediaries between any pair of nodes in the network is quite small and (2) the large local "cliquishness" or redundancy of the network—i.e., the large overlap of the circles of neighbors of two network neighbors. The latter property is typical of ordered lattices, while the former is typical of random graphs.

Recently, Watts and Strogatz (1998) proposed a minimal model for the emergence of the small-world phenomenon in simple networks. In their model, small-world networks emerge as the result of randomly rewiring a fraction p of the links in a d -dimensional lattice. The parameter p enables one to continuously interpolate between the two limiting cases of a regular lattice ($p = 0$) and a random graph ($p = 1$).

A more general question prompted by these results is: how typical are small world regimes? A formal representation of this question corresponds to whether the small-world property emerges for finite values of p when N approaches infinity (e.g., Barthelmy et al. 1999). Numerical results and theoretical arguments show that the emergence of the small-world regime occurs for a value of p that approaches zero as N diverges (e.g., Barthelmy et al. 1999). The implications of this finding are as follows. Consider a system for which there is a finite probability p of random connections. It then follows that independently of the value of p the network will be in the small-world regime for systems with size N close to $1/p$. In other words, most large networks are small-worlds. Importantly, in social networks the agents are likely to be "unaware" of this fact as the vast majority of them have no long-range connections.

An important characteristic of a graph that is not taken into consideration in the small-world model of Watts and Strogatz is the degree distribution, i.e., the distribution of the number of connections of the nodes in the network. The Erdős-Rényi class of random graphs has a Poisson degree distribution, while lattice-like networks have even more strongly peaked distributions. A perfectly ordered lattice, for example,

has a delta Dirac degree distribution. Similarly, the small-world networks generated by the Watts and Strogatz model also have peaked, single-scale, degree distributions, i.e., one can clearly identify a "typical degree" of the nodes comprising the network.

Against this theoretical background, Barabasi and coworkers found that a number of real-world networks have a scale-free degree distribution with tails that decay as a power law (Albert et al. 1999; Barabasi 1999) as found in the theory of critical phenomena. For example, a social network of movie-actor collaborations, the webpages in the .edu domain, and the power grid of Southern California, all appear to obey distributions that decay in the tail as a power law (Barabasi 1999). Importantly, scale-free networks provide extremely efficient communication and navigability as one can easily reach any other node in the network by sending information through the "hubs," the highly-connected nodes. Moreover, scale-free networks are robust. Their properties survive if nodes or connections are removed randomly. Targeted removal of hubs, however, destroys those beneficial properties.

Recently, researchers have started to address the impact of knowledge networks on organizational performance (Nasrallah and Levitt 2001; Huberman and Hogg 1995). Because problem solving in a crisis setting is often collaborative (e.g., think of the interactive brainstorming of the Apollo 13 engineers as they crafted a return strategy), modeling team interactions (especially if the research focus is on the ability to find innovative solutions) is a key component of the study of crisis management. An important example is Uzzi and Spiro's (2005) empirical study of the Broadway industry. They are interested in identifying network characteristics that encourage innovation, here artistic creativity in the musical industry. Musicals are created by production-specific teams that include a producer, composer, choreographer, etc. Individuals are linked when they collaborate on a given musical production. Uzzi and Spiro then study how these collaboration networks change over time and whether the amount of clustering in the network correlates with commercial or artistic success, measured by box office results, running time, and the like. The main finding is that the likelihood of success indeed correlates with the degree of clustering. Moreover, the influence is non-monotonic. Intermediate levels of clustering are associated with the highest likelihood of commercial success.

Influenced by this line of work, Guimera et al. (2005) study a model of team formation characterized by the propensities that incumbents continue to collaborate on a new project or are matched with other

incumbents or with newcomers. They show that the respective matching probabilities lead to different network topologies and, therefore, different expected performance. The model then is applied to various collaboration networks from different scientific communities.

An independent line of research has focused on the importance of cognitive diversity in teams (March 1991; Page and Hong 2001). Page and Hong, for example, show how groups with diverse search and problem solving strategies outperform experts even if the average competence of the group members is lower compared than in homogeneous teams. The intuition is that the combination of diverse search strategies is more likely to find globally optimal solutions. These results provide a foundation for the common practice in crisis management situations to assemble cross-functional teams. For example, when Mercedes had to manage the recall of its A-class car in the European market in 1997, the crisis team included not only safety engineers, but logistics experts, marketing experts and PR specialists.

From the point of view of crisis management these findings are important because they suggest how organizations can improve their ability to find innovative solutions in crisis situations by changing the characteristics of social interaction whether by increasing diversity, mixing newcomers with incumbents or influencing network clustering. Some insights from this research are already incorporated into management practices (such as the use of cross-functional teams), while others suggest new venues. For example, one of the consequences of the research on team networks suggests that team assignments and training policies should incorporate consequences for the network structure. That is, a joint training course may not only improve the skills of the individual members but it may also create new links (or strengthen existing links) among team members or across teams, with consequences for the performance of the crisis response network as a whole. These insights are beginning to be applied in the area of emergency medicine, in so-called critical care collaboratives. An example is the collaborative of Neonatal Intensive Care Units lead by the University of Vermont.

The application of network theory to knowledge based systems has provided a theoretical framework to study collective problem solving and innovation. By changing the interaction patterns we can expect organizations to improve their ability to find solutions to new problems and aggregate distributed information more efficiently. The existing results suggest various management or policy implications ranging from the increased use of online communication to the introduction of

knowledge brokerage systems and communities of practice. Indeed, the models can be interpreted as micro-foundations for such practical solutions. Finally, as the communication activity can easily be measured (either automatically or through surveys) managers can, at least in principle, monitor the progress of their organization and intervene if necessary.

However, while the knowledge network perspective may be useful in a crisis management context by characterizing information sharing and collective problem solving, they do not help us understand the critical question of how rapidly and reliably the organization can resolve problems or crises.

III. A Simple Model of Collective Problem Solving

Despite these insights' attention, knowledge network analysis is still in its infancy as a management tool. Existing models are useful for characterizing connectivity among agents in an organization, but they are not yet well-suited to predicting impacts on system performance or identifying specific improvement levers. One reason for this is that knowledge-based organizations generally perform two basic functions: (1) knowledge creation, and (2) problem solving. The current state-of-the-art knowledge network modeling provides useful insights into the first function, including how organizations generate, transmit and share knowledge (see Argote 2003 for an overview). But we lack a corresponding understanding of how knowledge based organizations translate this knowledge into tangible outputs, such as timely solutions to problems. Hence, while knowledge networks may be useful in a crisis management context by characterizing information sharing, they do not help us understand the critical question of how rapidly the organization can resolve problems or crises. To put it differently, in a crisis context, organizations not only need to be creative in finding solutions to unanticipated events, they need to do so at high speed and with high accuracy. Such a system would have to satisfy various performance requirements such as:

- *Accuracy.* Agents need to coordinate on the desired collective behavior.
- *Speed.* The desired behavior needs to be reached in a realistic time.
- *Error Tolerance.* Removal of agents or mistakes in processing information should only lead to a moderate decrease in system performance.

- *Scalability.* The resources (e.g., time) to perform the collective task should increase at a slow rate as the number of agents (and the system capabilities) grows more rapidly.

A useful model to study these issues in more detail was suggested in recent papers by Moreira et al. (2004a and 2004b) and then extended by Seaver et al. (2006). The idea is to consider networks of agents who can be in one of two states denoted $\sigma_j = \pm 1$. For example, the states could be two possible solutions to a given task, and the initial state of the system can be interpreted as each agent's initial belief about which solution is correct. The beliefs could be based on receiving a signal that is partially correlated with the correct solution. Agents are connected in a small world network with k_j neighbors with rewiring probability p . Moreover, to capture the effects of misunderstandings and other forms of miscommunication there is a probability η of miscommunication. That is, with probability η agents perceive the state of any connected agent to be -1 when it is $+1$ and vice versa. Collective problem solving is modeled as a density classification task, a widely used measure of coordination and global information processing (Crutchfield and Mitchell 1995). For a system comprised of units whose state is a binary variable, the density classification task is completed successfully if all units converge to the same state *and* the coordinated state is identical to the majority state in the initial configuration. In the case of convergence on the correct state the system therefore successfully aggregates all local information and arrives at the correct solution to the task.

Moreira et al. (2004a) show that a simple heuristic ("do what the majority of your neighbors do") leads to rapid and robust convergence to the correct state provided the interaction structure is characterized by moderate noise and constitutes a small-world network. Importantly, both conditions are necessary for effective problem solving. Moreover, more complicated decision-rules that work well in the case of $\eta = 0$ and $p = 0$, for example, the Gacs-Kurdyumov-Levin (GKL) rule (Crutchfield and Mitchell 1995) fail to function when communication is noisy or interaction occurs in an asynchronous fashion. Intuitively, this implies that random connections to other members of the network not only are important for problem solving, but that if they exist, even extremely simple decision heuristics can be successful. Moreover, these systems satisfy all the criteria critical for crisis situations. They are accurate, fast, error tolerant, and scalable.

Consider the following example adapted from Moreira et al. (2004b). In this variant of the model, a modular or "island" network is considered. In a modular network N individual agents are divided into communities of equal size S . Then, each agent is connected to k other agents. With probability p the connection is directed to a random unit in the network and with probability $(1 - p)$ the connection is established within the unit community. The fraction of extra-community connections p controls the network topology: for $p = 0$ one has completely disconnected communities, while for $p = 1$ one has a random graph. To see the importance of p consider figure 5.2.

As can be seen from the figure, the effect of moving to a small-world regime is striking. Problem solving efficiency (measured by the percentage that the system converges to the correct solution in $2N$ time steps) dramatically increases as random connections reach a critical threshold. Intuitively, an organization will be in this state if on average at least 20 percent of each agent's collaborative relationships are outside their immediate unit or work group.

An important application in the context of public sector applications is the case of multiple agencies sharing responsibilities, a common organizational feature especially in Federalist political systems. Examples range from identifying and controlling epidemics to emergency response or intelligence gathering agency networks. This research not only suggests the critical importance of collaborative task forces (such as joint counter-terrorism task forces that include members of the FBI, CIA, and other members of the intelligence community) but the importance of facilitating serendipitous connections across existing units, as suggested by the small-world model. These can be created by infor-




Figure 5.2
Majority rule performance in modular networks.

mal means such as joint training or preparation workshops or more formalized integrations of response and communication networks. The point is that for successful crisis management we not only need to create efficient structures for expected scenarios, but to improve the innovativeness of response networks. It follows that the main benefit from a joint training exercise of state, local, and Federal law enforcement, for example, may lie less in the developed common protocols than in the resulting mutual trust that may make it more likely that a member of the local policy force calls his counterpart at the state level.

This insight can be easily captured in a rule of thumb. Suppose we measured the network of trusted relationships, e.g., in a simple survey, as follows: A has a trusted relationship with B if A is willing to ask B for help in a case where A does not know the answer to a given problem. In our computer model (Moreira et al. 2004b) we establish a necessary condition for responsive organizations: each individual should have at least 20 percent of his or her trusted relationships outside his immediate organizational sub-unit (whether it's a group, field office, or department). Take the FBI as an example. For a typical field agent 1 in 5 of all his or her trusted relationships should be outside of her field office. Similarly, 1 out of 5 of all trusted relationships of, e.g., an intelligence analyst, should be members of other agencies. The research on knowledge networks suggests that organizations that have fewer average cross-unit relationships are unlikely to effectively respond to unexpected challenges.

Seaver et al. (2006) consider more general interaction structures where agents may have various forms of decision biases. Specifically, they consider organizations that may include "conservative" or "partisan" agents which hold a bias toward a particular state. Specifically, a conservative agent requires a "qualified" majority of her neighbors to convince her to change her state. But if a conservative agent changes her state, it will again take a qualified majority to change that new state to yet another state. "Partisan" agents, on the other hand prefer a particular state, e.g., -1 . In that case it will take a qualified majority to cause the agent to change her state to $+1$, but only a simple majority to change back to -1 . Of course, different individuals can have different bias strengths as measured by the size of the qualified majority necessary to lead to a change of state.

Adding partisans to the model (even if they are distributed evenly between partisans for $+1$ or -1) dramatically decreases the performance of the system. In the case of conservative agents the model yields a

surprising conclusion. For moderate levels of bias (for a majority of 5/7 or less of the neighbors) the system not only shows remarkable levels of efficiency, but system performance actually *increases* as the fraction of conservatives increases, provided the inter-agent noise level is sufficiently high. However, there is an important trade-off with the speed with which a solution is reached. That is, the time to reach consensus grows in the fraction of conservatives. If the fraction is larger than 30 percent, consensus cannot be reached in $2N$ time steps.

To summarize, the simple model of information aggregation yields two important insights. First, once knowledge networks have enough connectivity, collective problem solving can be highly accurate, fast, robust, and scalable. Second, these properties largely survive when we consider agents with a cognitive bias ("conservatives"). However, there now is a trade-off between accuracy and speed. The system may even reach higher accuracy levels, but at significantly reduced speed.

The fact that partisanship and conservatism behave so differently is of particular importance in the current debate concerning intelligence failures as in the case of Iraq's lack of Weapons of Mass Destruction or the alleged failure of the intelligence agencies to "connect the dots" in the advance of 9/11. The results suggest that even modest levels of partisanship (i.e., bias towards a particular solution) can lead to dramatic drops in system performance.

IV. Responsiveness and Adaptation

The insights from simple knowledge works suggest the importance of modeling trade-offs between different performance dimensions in more detail. One particularly useful representation is network flow models (e.g., Gordon and Newell 1967; Jackson 1957) which were originally created to understand the impacts of capacity and congestion in production systems. They have provided many useful insights into behavior of systems ranging from job serial production lines to complex supply chains (see e.g., Altiock 1997; Askin and Goldberg 2002; Hopp and Spearman 2000). Considerable research has been devoted specifically to the problem of promoting flexibility and responsiveness in production and service systems (see e.g., Degroote 1994; Wadhwa and Rao 2003 for overviews). The insights from this research have played a role in the evolution of practices (labeled variously as time based competition, quick response manufacturing, flexible production, just in time, agile manufacturing, and lean production) for facilitating fast, flexible response in production and service systems.

The challenge is to combine these two approaches in the context of a crisis situation. However, at this point there simply is no framework for assessing and enhancing responsiveness in knowledge management contexts. To develop such a framework to support the vast number of public sector crisis response organizations, and the even larger number of business organizations that must face occasional crises, basic research is needed to investigate the factors that influence responsiveness of a knowledge based organization.

In the remainder of this paper we outline a path of research aimed at generating a more general formalism for crisis management. We focus on two issues that are essential in managing emergency situations: (1) responsiveness, which represents the ability of a system to perform tasks or resolve problems quickly, and (2) adaptability, which measures the system's ability to dynamically adjust to changing environmental conditions. By combining insights from both social networks and production flow networks, we eventually hope to generate a new class of models, which we call adaptive response networks, that will enable us to diagnose and improve the preparedness of organizations. Ultimately, the goal is to create a rigorous framework to support the design and management of highly responsive and effective crisis management organizations. At this stage such a general framework does not yet exist. However, we can provide some insights in the case of an important sub-class of crisis response infra-structure: emergency call centers.

In addition to becoming a large service industry, employing roughly 3–4 million Americans and many more internationally, call in-bound centers provide critical emergency services such as 911, police, ambulance, fire dispatching, etc. Workforce management (i.e., workforce training and scheduling) in call centers is a very difficult task, due to the high variability in call arrivals and response times, especially in crisis situations.

While an important topic in their own right, understanding the performance of emergency call centers also provides some more general insights into the designing, training, and supporting of crisis teams. As discussed above, one important insight from the knowledge network literature is that how crisis teams are formed will affect the quality of the solution. But we also know that in a crisis, there isn't always time to form the best teams. Even if a search could be carried out quickly enough, the most desirable individuals may be unavailable due to an overload. For example, in the aftermath of Hurricane Katrina, the U.S. Fish and Wildlife Service was available to help with search-and-rescue operations in New Orleans, but was unable to reach someone

from FEMA to get instructions. So, instead of waiting for coordinated assignments, the Fish and Wildlife team went ahead as best they could and rescued some 4,000 people.

Hence, our framework must characterize the tradeoff between speed and precision in assigning teams. We must also recognize that crisis situations are precisely when organizations are likely to be overloaded and hence, at least in the short term, team formation may be extremely difficult. Hence, the structure should be robust enough to function when highly disrupted, but adaptable enough to return quickly to more efficient operation as time permits.

One way to increase responsiveness of a system is to make it more flexible. In this section we present some of our preliminary work aimed at designing flexible system structures in the context of call centers (Iravani et al. 2005, Iravani et al. 2004). This can be illustrated in the specific case of a call center, and we show how the well-known Average Path length (APL) metric of Small World Networks (SWN) can be used to capture the flexibility of cross-training structures of call center agents.

Managers of call centers have found that careful attention to cross-training of their workforce can help avoid lost calls and reduce long waiting times. Cross-training allows labor capacity to be dynamically reallocated in response to shifts in call volume and mix. Even when there are no apparent trends in demand, cross-training reduces the frequency with which agents starve for work due to intrinsic variability in inter-arrival and service times. Effective use of cross-trained agents can reduce caller wait times and/or staffing requirements.

However, full cross-training of every agent for every call type is very costly and sometimes impossible (e.g., with call centers that serve clients in several languages). Hence, the problem becomes one of finding a partial cross-training program that results in a flexible structure that can mitigate the disruptive effects of variability. We have found that the concept of small world networks discussed above can also be applied to capture the flexibility offered by different cross-training structures in the form of an index that can be used to choose an effective structure from feasible alternatives.

To illustrate our method consider the four cross-training structures illustrated in the following figures.

In figure 5.3, there is no cross-training, and, therefore, each agent can only answer a single call type. Suppose that there is high variability in call inter-arrival times and call service times. This may result in situations where the queue of call type A becomes empty, while the queue of

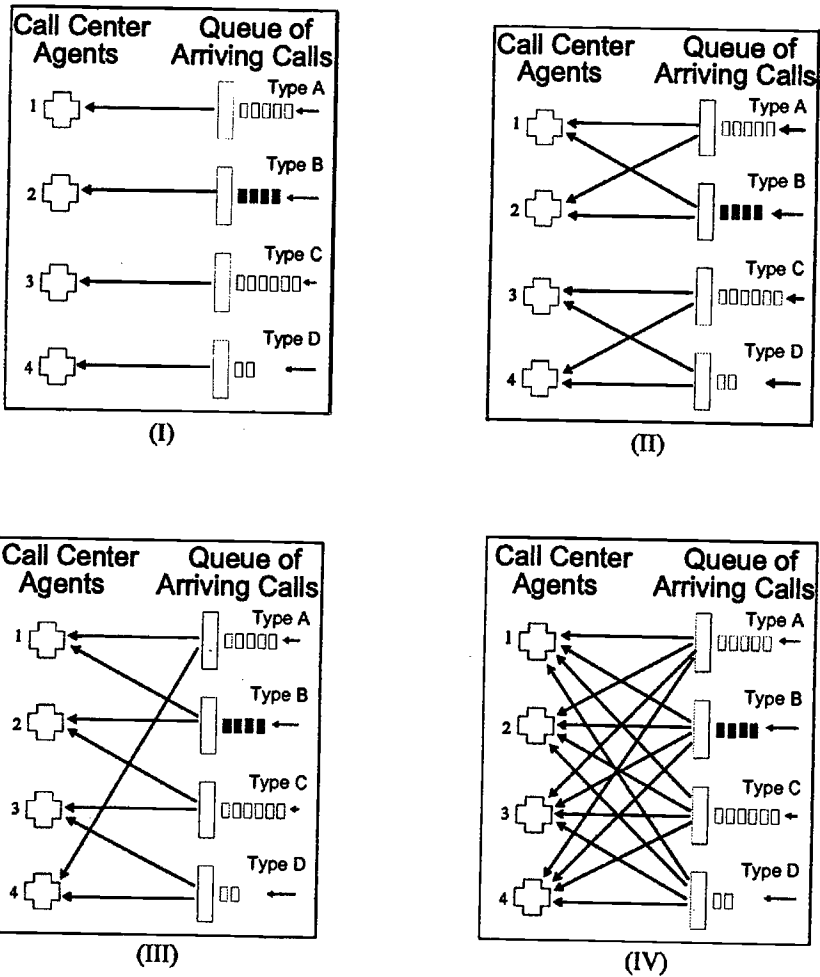



Figure 5.3
Examples of cross-training structures in a call center.

call type B is very long, and therefore customers in that queue experience a long waiting time. Since Agent 1 is not cross-trained for call type B, her available unused capacity cannot be used to help Agent 2. Hence, in the presence of variability, some workers will occasionally be starved for work while others are overwhelmed, which may cause long queue lengths.


Additional cross-training (i.e., partial server pooling) can alleviate such congestion. In systems (II) and (III) Agent 1 is cross-trained to serve calls of type B. By using this skill appropriately (for example, at

times when queue A is empty and queue B has 2 or more calls), then customers of type B will experience less waiting time in the line. Both structures (II) and (III) have the same total number of skills. Furthermore, in both (II) and (III) all agents are cross-trained for two skills and every call type can be answered by two agents. However, it has been shown (see Jordan and Graves 1995) that in the presence of variability, structure (III) is more effective than structure (II). One reason is that in (II) Agents 1 and 2 cannot help Agents 3 and 4, while in (III) Agents 1 and 2 can (directly or indirectly) help Agents 3 and 4.


Figure 5.4 shows another example in which all agents are cross-trained for all call types, and therefore each agent can help every other agent in responding to any call type. This structure is known as complete server pooling or full cross-training, and is well known to reduce congestion (under an appropriate service policy). Note that, while (IV) is the best performing cross-training structure, it is also the one with the highest training and/or wage costs.




(I)



(II)



(III)



(IV)

Figure 5.4

WS network representations for cross-training structures.

The above examples show how and why the addition of a skill can improve system performance, provided the coordination of workers is effective. The main role of the additional skill is to give an agent the capability (flexibility) to help another agent service a particular call type when needed. This example also suggests that the more ways agents can help each other, the more effective the cross-training structure is in improving the system performance (something we have found to usually be the case in our study of other systems with cross-trained labor).

To characterize the flexibility of a given cross-training structure more generally, we turn to the network literature discussed above. Specifically, we use the Average shortest Path Length (APL) and the cluster coefficient as measures of the network topology. We have found the APL to have an obvious intuitive relationship that corresponds to the efficiency and responsiveness of cross-training structures in call centers.

We have developed a methodology based on the APL metric for converting the design of a cross-training structure into a useful small-world representation. This is based on the Work-sharing (WS) Network, which we define by letting nodes represent agents and arcs represent an overlap in the skill sets of two agents. Furthermore, the length of an arc connecting two nodes, i and j , is the reciprocal of the number of call types that can be served both by Agents i and j .

For example, consider the following network representation of the cross-training structures displayed above.

In figure 5.2, there is an undirected arc connecting nodes $i = 1$ and $j = 2$ because Agents 1 and 2 in figure 5.2 can both serve (i.e., help each other in serving) at least one common call type (A or B). In the WS network in figure 5.3 there is also a unidirectional arc connecting node $i = 1$ to node $j = 2$. The reason is that, as figure 5.3 shows, Agents 1 and 2 can both help each other in serving as least one (i.e., in this case only one) common call type B.

Note that figure 5.1 has no link between any nodes, because the agents are not cross-trained and therefore have no skills in common. Although systems (I) and (II) have illustrative value, our methodology is intended for structures such as (III) and (IV) that are connected, a typical small-world network (Watts and Strogatz 1998). Figure 5.4 has every pair connected, and this represents the fully cross-trained case in which every agent can help each other.

Given variability in the demand and/or service processes, it is clear that the more call types agents can help each other with (i.e., greater

number of shared call types), the more effective the cross-training structure would be. To capture this, we set the length of the arc between i and j , $arc(i,j)$, to be the reciprocal of the number of call types that both Agents i and j can serve. Therefore, if Agents i and j help each other in more call types, the length of the $arc(i,j)$, becomes smaller. Consequently, the network structures with smaller lengths between their nodes (and thus smaller APL) will represent cross-training structures in which agents can help each other in more call types.

To appreciate the arc lengths in a network, again consider structures (II) and (III) and their corresponding WS network. In figure 5.2, Agents 1 and 2 can help each other to serve two call types, A and B. Therefore, the length of the arc connecting nodes 1 and 2 in WS network in figure 5.2 is the reciprocal of 2, which is 0.5. On the other hand, in figure 5.3, agents 1 and 2 can help each other in serving only one call type, namely type B. Hence, the length of the arc connecting nodes 1 and 2 in WS network in figure 5.3 is 1 (i.e., the reciprocal of 1). Similarly, in figure 5.4, there are 4 call types (call types A, B, C, and D) which Agents 1 and 2 can both serve. Therefore, in figure 5.4 the length of the arc connecting nodes 1 and 2 is 0.25, the reciprocal of 4.

For a graph with N nodes, computation of the APL metric requires the calculation of the minimum distance between every possible pair (i,j) , denoted as $L_{ij}^{(\min)}$. Since the path between node i and i has no meaning in our WS network, and since the shortest path from node i to node j is the same as the shortest path from node j to node i (i.e., $L_{ij}^{(\min)} = L_{ji}^{(\min)}$), we only need to calculate $N(N-1)/2$ shortest paths. The APL of a Work Sharing network with N nodes is therefore the average length of these $N(N-1)/2$ paths, which can be calculated as:

$$APL = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^N L_{i,j}^{(\min)}$$

Note that our SW network is defined such that a smaller APL number generally corresponds to a cross-training structure in which agents have greater versatility in helping each other. Thus, the smaller the APL, the more flexible and more effective the cross-training structure is in reducing the customer average waiting times. We call our methodology based on the WS network and the APL metric the *WS-APL method*.

As an illustrative example, consider a call center that receives 12 different types of calls, which are labeled A,B,C,...,L. Calls of type i arrive randomly with a rate specified by i^{th} element of the demand rate vector $D = (0.5,0.5,0.5,0.75,0.75,0.5, 1.167,0.667,0.667,0.333,0.333,0.333)$. Suppose

that the call center has seven agents. We assume that call handling/service times are stochastic with an average of 0.9 units of time (which corresponds to a system utilization of 90 percent under an aggregate arrival rate of seven calls per unit time). Figure 5.5 shows two agent cross-training structures that are easily capable of handling demand vector D , because under both structures, all call types receive enough capacity.

The key question, therefore, is which cross-training structure is more flexible (i.e., will yield a smaller average customer waiting time over a range of operating conditions that include high variability in call arrival and call response time, as well as high call volumes during the peak hour)?

The WS networks of the same structures are presented in figure 5.6.

Although Structure 1 has more total number of kills than does Structure 2, the WS network of Structure 2 has a smaller APL than that of Structure 1 ($APL_2 = 1.21 < APL_1 = 1.63$), which implies that the cross-training of Structure 2 should be more effective than Structure 1.




Figure 5.5
Alternative structures for demand vector.




Figure 5.6
WS network representation for cross-training structures.

To see which structure actually does yield the lowest average customer waiting time, we developed a discrete-event computer simulation program and estimated the average customer waiting times under each cross-training structure. The underlying model is a queuing network with parallel, infinite-buffer queues. Call inter-arrival times and call service times are modeled using Gamma distributions (which covers a wide range of variability scenarios with coefficients of variation of less than, equal to, or greater than one). Over a range of variability levels in the call arrival and service processes, as well as the utilization (load) of the system and peak hours, simulation reveals that Structure 2 does indeed outperform Structure 1 in minimizing the mean waiting times. In fact, for different variability and utilization scenarios, Structure 2 resulted in 15.6 percent to 24.6 percent (with an average of 20.2 percent) smaller waiting times than Structure 1. This is consistent with the prediction of our WS-APL method that suggests Structure 1 is the more flexible alternative.

We performed similar experiments on close to 1,000 more cases that include: (1) systems ranging from 6 to 10 call types, (2) systems with different levels of uncertainty in call arrival process and call response time (i.e., we considered coefficient of variation $CV = 1$ and $CV = 2$ for call inter-arrival and service times), (3) systems with random shocks where a shock significantly increases call arrival rate during a particular interval (i.e., peak hours). For details of our numerical study see Iravani et al. (2005).

We found that in 90 percent of the cases in our numerical study, the WS-APL index was able to detect the more flexible structure among

any pair of alternative structures. In the remaining 10 percent where the APL prediction was wrong, the performance (i.e., the average customer waiting time) of the structure chosen by the WS-APL index was, on average, only 2 percent worse than that of the better structure. This supports our assertion that in general, the smaller the world of the WS network (i.e., the smaller the APL), the better the performance of the corresponding cross-training structure.

As our preliminary results show, system structure, if properly designed, can significantly increase system responsiveness by mitigating the negative effects changes in the environment (including disruptions due to routine variability and/or unpredictable shocks). This system responsiveness can be characterized quantitatively through the use of network analysis.

An important next goal is to extend this approach for measuring structural flexibility to develop general metrics that characterize the ability of an organization to respond quickly and accurately to emergency situations. To do this, one will need to consider more complex workflows than those in an emergency call center. The task network involved in responding to a crisis will generally involve precedence constraints, collaborative work, and learning over time. But the same basic underlying structure of a network of capacity constrained agents responding to uncertain and dynamically varying workloads, is still valid. Hence, we expect that a structural analysis in the same vein as that described above will provide part of a general characterization of adaptive response networks. The following section will outline some ideas of how this can be accomplished.

V. Towards a General Approach

Our general goal is to understand the impact of organizational structure on the responsiveness and effectiveness of crisis management by developing a theory of adaptive response networks. To do this, we need to model an organization as a network in which nodes represent agents and arcs represent potential collaboration links. Each agent has specific skills (the nature of which may be only partially observable) and is capacitated (so that system congestion will cause delays in completing tasks). Arcs represent formal and/or informal links which can be exploited to search for information and form problem solving teams.

A "crisis" can be modeled as the random arrival of a problem, which can be resolved by completion of a set of interrelated tasks. The qual-

ity and speed of the crisis resolution will depend on task sequencing, assignments of teams to tasks, and allocation of agent time to tasks. In this modeling context, the crisis management problem devolves to one of assigning agents to tasks over time in a fashion that yields a high quality solution in a time responsive manner. One may begin by assuming that the problem structure is given and that agent assignments are centrally controlled and will seek insights into the network structure and control policy that yields the best results. A next step would be to then relax the central control assumption and assume that agents make local decisions regarding their time allocations and will seek further insights into the impact of training and communication policies on performance. Finally, one should consider the case where the nature of the problem is revealed over time and seek to identify effective "act and adapt" strategies for simultaneously learning about and responding to a crisis situation.

From a structural perspective, the parallel server structure of the call center represents one of the simplest possible environments. Tasks (customers) are all single step operations. Skills are represented with simple "on" or "off" switches. And disruptions are restricted to fluctuations in workloads. The literature on social networks in organizations provides a basis for representing more complex problem solving environments. For example, figure 5.7a represents a pure hierarchical organization (adapted from Watts 2003, figure 9.1). Both NASA and many emergency rooms have hierarchies like this. If this structure were used rigidly, for two agents to communicate, they must connect through formal channels, which may require many steps. In contrast, figure 5.7b represents a hierarchical organization that has adapted by evolving




Figure 5.7

(a) Pure hierarchical organization; (b) hierarchical organization with informal adaptive links.

informal links that serve as communication shortcuts. Such informal communication links can be promoted by policies ranging from office layout (Davenport 2005) to use of formal "knowledge brokers" (Cross and Prusak 2002).

As discussed above, facilitating communication via shortcut links has been widely studied in the complex networks literature (see e.g., Watts and Strogatz 1998). However, to our knowledge, none of this literature has attempted to explicitly model task processing times and therefore no previous model can predict the impact of network evolution on response times. The analysis then should be extended to other collaborative structures such as matrix organizations, team based structures, multi-agency collaborative environments, etc.

A natural extension of a model where agent teams are assigned (or self-assigned) to tasks, which they will carry out as expeditiously as their workloads will permit, would then also consider the possibility of adjusting task assignments or team structures as the crisis evolves. Moreover, in most crisis situations, however, conditions change dramatically over time. New information becomes known. Agents who were either unavailable or overloaded become accessible. Completion of tasks reveals the need for other tasks. Hence, in practice, the ability to adapt to changing conditions is a vital crisis management skill. This suggests extending the task resolution model of crisis management to incorporate act and adapt policies under which agents strike a balance between acting on their currently assigned tasks and searching for new information and/or collaboration opportunities. Acting immediately on the basis of local information speeds response, but spending time to search for better solutions may ultimately produce a faster good response.

The recent Hurricane Katrina experience offers some instances where too much search and coordination led to resources being idled while people were in immediate peril. As a tragic example of the consequences of too little of a search, we may consider the case of a patient who came into an emergency room exhibiting symptoms of confusion and odd behavior. The initial agent responsible for the situation was the receptionist, who had to decide whether to have the patient wait or be seen immediately. In this case, the receptionist called a nurse to examine the patient, and it became her decision to seek additional expertise immediately or to have the patient wait. Had the nurse called a physician to see the patient, the physician would have had to decide whether to call in outside expertise. At each step, the individual examining the patient also had to decide how much time to spend with the patient

and what steps to take. Sadly, in this case it turned out that the patient was suffering from heat stroke aggravated by an impermeable weight loss suit and the patient died in the waiting room before being seen by a physician. Failure of the system to respond quickly enough had fatal consequences.

To extend a modeling framework to include adaptability, as well as responsiveness, one should therefore:

1. Incorporate agent utilization as a variable in the team formation process. A natural starting point is the study of flexible work systems in which queue length information is used to dynamically allocate agents (see e.g., Van Oyen, Senturk-Gel, and Hopp 2001). The policies from this research offer ways to balance the need for working on the most pressing tasks in the short term with the need to level agent workloads to provide good performance in the longer term.
2. Incorporate information search into the adaptive response network framework. That is extending the problem of search in complex networks (see e.g., Watts 2003, Chapter 9 for an overview) to include models of search in the presence of constraints on agent capacity and/or system congestion.
3. Incorporate learning into the adaptive response network framework. To do this, one must model changes in the available information (e.g., the skill level needs of a particular task or group of tasks). We must also consider the possibility of new tasks for addressing the crisis becoming available (e.g., new developments occur which require attention of agents). The resulting problem could be modeled, for example, as a dynamic stochastic control problem or involve the use of plausible heuristics.

VI. Conclusion and Policy Implications

In this essay we suggest that knowledge network and network flow models offer a highly flexible, yet integrated approach to modeling crisis responsiveness and preparedness. Clearly, we are in the very early stages of exploring these modeling approaches. Yet, even our very preliminary results suggest that they may provide new insights that can guide practitioners in designing improved crisis management practices. Specifically, our approach suggests at least three novel perspectives on crisis management.

1. Rather than developing response strategies to more potential “contingencies,” or “scenarios” crisis management performance may be enhanced more readily by focusing on the responsiveness of organizations: the ability to find solutions to unanticipated problems and tasks.
2. Organizations vary in their responsiveness. But the factors that can account for these differences are not yet well understood. Once we conceptualize organizations as networks of individuals that need to quickly solve an incoming flow of critical problems, we can begin to model the performance drivers of such knowledge networks systematically.
3. The insights from such modeling activities combined with careful data analysis and field study will likely yield important policy conclusions. Even our initial results suggest that as governmental actors are facing an increasing number of complex and unanticipated events, investment in inter-agency connections is likely to yield large benefit as long as these connections lead to trusted relationships between members from different agencies or backgrounds. Indeed the benefits from improving the interaction between agents may far outweigh the benefits from investing in expensive individual training programs.

Acknowledgements

The authors wish to thank Scott Stern for valuable suggestions. All remaining errors are our own. Financial support for this project was provided by the Ewing Marion Kauffman Foundation.

References

- Albert, R., H. Jeong, and A.-L. Barabasi. 1999. *Nature* 401: 130–131.
- Albrecht, T. L., and V. A. Ropp. 1984. “Communicating about Innovation in Networks of Three U.S. Organizations.” *Journal of Communication* 34: 78–91.
- Altioik, T. 1997. *Performance Analysis of Manufacturing Systems*. New York: Springer.
- Argote, L., B. McEvily, and R. Reagans. 2003. “Managing Knowledge in Organizations: An Integrative Framework and Review of Current Themes.” *Management Science*: 571–582.
- Askin, R. G., and J. B. Goldberg. 2002. *Design and Analysis of Lean Production Systems*. New York: John Wiley & Sons.
- Barabasi, A. 2002. *Linked: The New Science of Networks*. New York: Perseus.
- Barabasi, A., and R. Albert. 1999. “Emergence of Scaling in Random Networks.” *Science* 286(15): 509–512.

- Barthelemy, M., and L.A.N. Amaral. 1999. "Small-World Networks: Evidence for a Cross-over Picture." *Phys. Rev. Lett.* 82: 3180–3183.
- Brass, D. J. 1995. "A Social Network Perspective on Human Resources Management." *Research in Personnel and Human Resources Management* 13: 39–79.
- Burt, R. S. 2004. "Structural Holes and Good Ideas." *American Journal of Sociology* 110(2): 349–399.
- Chesbrough, W. H. 2003. "A Better Way to Innovate." *Harvard Business Review* 81: 12–13.
- Cross, R., and L. Prusak. 2002. "The People Who Make Organizations Go — Or Stop." *Harvard Business Review* 80(6): 104.
- Crutchfield, J. P., and M. Mitchell. 1995. "Evolving Globally Synchronized Cellular Automata." *Proceedings of the National Academy of Sciences* 92: 10742.
- Davenport, T. 2005. *Thinking for a Living: How to Get Better Performance and Results from Knowledge Workers*. Harvard Business School Press.
- De Groote, X. 1994. "The Flexibility of Production Processes: A General Framework." *Management Science* 40(7): 933–945.
- Diermeier, D. 2005. "From Minimizing Liability to Maximizing Opportunity — Crisis Management with Applications to the Natural Resource Industries." *Rocky Mountain Mineral Law Review*. Westminster, CO.
- Feldman, P. M. 1999. "The New Economics of Innovation, Spillovers and Agglomeration: A Review of Empirical Studies." *Economics of Innovation and New Technology* 8: 5–25.
- Fink, S. 2002. *Crisis Management: Planning for the Inevitable*. Lincoln, NE: iUniverse.
- Fombrun, C. J. 1986. "Structural Dynamics with and between Organizations." *Administrative Science Quarterly* 31: 403–421.
- Gordon, W., and G. Newell. 1967. "Closed Queueing Systems with Exponential Servers." *Operations Research* 15: 254–265.
- Granovetter, M. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78(6): 1360–1380.
- Guimera, R., B. Uzzi, J. Spiro, and L. Amaral. 2005. "Team Assembly Mechanisms Determine Collaboration Network Structure and Team Performance." *Science* 308: 697–702.
- Hopp, W. J., and M. L. Spearman. 2000. *Factory Physics: Foundations of Manufacturing Management*, Second Edition. Burr Ridge, IL: McGraw-Hill.
- Huberman, B. A., and T. Hogg. 1995. "Communities of Practice: Performance and Evolution." *Computational and Mathematical Organization Theory* 1: 3–92.
- Iravani, S. M. R., B. Kolfal, and M. P. Van Oyen. 2004. "Call Center Labor Cross-Training: It's a Small World After All." Forthcoming in *Management Science*.
- Iravani, S. M. R., K. Sims, and M. P. Van Oyen. 2005. "Structural Flexibility: A New Perspective on the Design of Manufacturing and Service Operations." *Management Science* 51: 151–166.
- Jackson, J. R. 1957. "Networks of Waiting Lines." *Operations Research* 5: 518–521.

- Jordan, W. C., and S. C. Graves. 1995. "Principles on the Benefits of Manufacturing Process Flexibility." *Management Science* 41(4): 577–594.
- Kilduff, M., and W. Tsai. 2003. *Social Networks and Organizations*. Thousand Oaks, CA: Sage Publications.
- Kogut, B. 2000. "The Network as Knowledge: Generative Rules and the Emergence of Structure." *Strategic Management Journal* 21(3): 405–425.
- Kogut, B., and U. Zander. 1992. "Knowledge of the Firm, Combinative Capabilities and the Replication of Technology." *Organization Science* 3: 383–397.
- Lesser, E. L., and J. Storck. 2001. "Communities of Practice and Organizational Performance." *IBM Systems Journal Knowledge Management* 40.
- March, J. G. 1991. "Exploration and Exploitation in Organizational Learning." *Organization Science*: 71–87.
- Milgram, S. 1967. *Psychology Today* 2: 60–67.
- Monge, P. R., and N. Contractor. 2003. *Theories of Communication Networks*. Cambridge: Oxford University Press.
- Moreira, A., D. Diermeier, and L. Amaral. 2004a. "Efficient System-Wide Coordination in Modular Networks." *Proceedings of the 2004 Workshop on Social Agents: Ecology, Exchange, and Evolution*. Argonne National Laboratory.
- Moreira, A., A. Mathur, D. Diermeier, and L. Amaral. 2004b. "Efficient System-Wide Coordination in Noisy Environments Using Heuristic Methods." *Proceedings of the National Academy of Sciences*, 101, 12085–12090.
- Naik, Gautam. 2005. "Planning and Private Resources at the Hospital Giant HCA Made Rescue Operation Possible." *The Wall Street Journal* (September 7): B1.
- Nasrallah, W. F., and R. E. Levitt. 2001. "An Interaction Value Perspective on Firms of Differing Size." *Computational and Mathematical Organization Theory* (7): 113–144.
- Newman, Mark, Steve Strogatz, and Duncan Watts. 2001. "Random Graphs with Arbitrary Degree Distributions and Their Applications." *Physical Review*: E 64.
- Page, Scott, and Lu Hong. 2001. "Problem Solving by Heterogeneous Agents." *Journal of Economic Theory* 97: 123–163.
- Seaver, S., R. D. Malmgren, M. M. Moreira, M. Sales-Pardo, D. Diermeier, and L. A. N. Amaral. 2006. "Micro-Bias and Macro-Performance." Working Paper. Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL.
- Sheffi, Y. 2005. "The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage." Cambridge, MA: MIT Press.
- Stanley, H. E. 1999. *Rev. Mod. Phys.* 71, S358–S364. [Special Issue for the Centennial of the American Physical Society.]
- Stevenson, W. B., and M. C. Gilly. 1991. "Information Processing and Problem Solving: The Migration of Problems through Formal Positions and Networks of Ties." *Academy of Management Journal* 34: 918–928.
- Uzzi, B., and J. Spiro. 2005. "Collaboration and Creativity: The Small World Problem." *American Journal of Sociology* 111(2): 447–504.

- Van Oyen, M. P., E. G. Senturk-Gel, and W. J. Hopp. 2001. "Performance Opportunity of Workforce Agility in Collaborative and Noncollaborative Work Systems." *IIE Transactions* 33: 761-777.
- Wadhwa, S., and K. S. Rao. 2003. "Flexibility and Agility for Enterprise Synchronization: Knowledge and Innovation Management towards Flexagility." *Studies in Informatics and Control* 12(3): 111-128.
- Wasserman, S., and K. Faust. 1994. *Social Network Analysis: Methods and Applications*. Cambridge, UK: Cambridge University Press.
- Watts, D. 2003. *Six Degrees: The Science of the Connected Age*. New York: Norton
- Watts, D. 1999. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton, NJ: University Press.
- Watts, D., and S. Strogatz. 1999. "Collective Dynamics of 'Small-World' Networks." *Nature* 393(4): 440-442.
- Watts, D. J., and S. H. Strogatz. 1998. *Nature* 393: 440-442.