

## **Learning From Organizational Experience**

John S. Carroll, MIT Sloan School of Management

Jenny W. Rudolph, Boston College Carroll School of Management

Sachi Hatakenaka, MIT Sloan School of Management

To appear in the Handbook of Organizational Learning and Knowledge

Mark Easterby-Smith and Marjorie A. Lyles (Eds.)

Blackwell Publishers

### **Abstract**

Learning-in-action, the cyclical interplay of thinking and doing, is increasingly important for organizations as environments and required capabilities become more complex and interdependent. Organizational learning involves both a desire to learn and supportive structures and mechanisms. We draw upon three case studies from the nuclear power and chemical industries to illustrate a four-stage model of organizational learning: (1) local stage of decentralized learning by individuals and work groups, (2) control stage of compliance with rules, (3) open stage of acknowledgement of doubt and motivation to learn, and (4) deep learning stage of skillful inquiry and systemic mental models. These four stages differ on whether learning is primarily single-loop or double-loop, i.e., whether the organization can surface and challenge the assumptions and mental models underlying behavior, and whether learning is relatively improvised or structured. The case studies illustrate the details of learning practices and the nature of stage transitions.

## **1 INTRODUCTION**

Learning-in-action, the cyclical interplay of thinking and doing (Argyris & Schon, 1996; Daft & Weick, 1984; Kolb, 1984), is increasingly important for organizations as they struggle to cope with rapidly changing environments and more complex and interdependent sets of knowledge. Organizational knowledge is embodied in physical artifacts (equipment, layout, databases, documents), organizational structures (roles, reward systems, procedures), and people (skills, values, beliefs, practices) (cf., Kim, 1993; Levitt & March, 1988; Schein, 1992). Although organizations may “fill knowledge reservoirs” (Argote & Ingram, 2000) from theoretical principles, by imaginative rumination, or by observing others, enactment or putting this knowledge to use requires combining component-level knowledge and filling gaps by improvisation (Weick, 1998). If the requisite performances are more unfamiliar, tacit, contextual, or contested among stakeholders, then the learning process will be more iterative, unpredictable, and emergent from evolving practice (Carlile, in press; Nonaka & Takeuchi, 1995).

### **1.1 Bridging Across Individuals and Groups**

Carrying out these learning activities as an organization involves complex interdependencies across people and groups (Crosson, Lane, & White, 1999; Kim, 1993). Knowledge is more than lists of facts that can be summed together (e.g., Nonaka & Takeuchi, 1995). Different parts of the organization, such as plant operators and corporate executives, “know” different things about how work is done. Their knowledge is contained in different reservoirs (Argote & Ingram, 2000) and expressed in different languages by groups that live in different thought worlds (Dougherty, 1992). Bridging across these groups requires common experiences and common referents, which are developed in bridging practices (Carlile, in press; Cook & Brown, 1999) including cooperative action, shared representations, collaborative reflection, and exchanges of personnel (Gruenfeld, Martorana, & Fan, 2000).

### **1.2 Skills for Organizational Learning**

Every organization uses multiple routines and tools for learning. As Popper and Lipshitz (1998) suggest, organizational learning involves both a desire to learn and the structures and mechanisms to enact learning effectively. Skills around conflict management and diversity may help groups work with one another, regardless of the substance of their cooperation (Jehn, Northcraft, & Neale, 1999). Specific logical and analytical skills may allow more meaning to emerge from available information. More comprehensive mental models and systems thinking skills may allow attention to a wider variety of information and the interdependencies among them (Senge, 1990).

In this chapter, we draw upon several case studies of organizational learning from the nuclear power and chemical industries in order to illustrate a four-stage model of organizational learning. By breaking the process of developing learning capabilities or absorptive capacity (Cohen & Levinthal, 1990) into stages, we illuminate the underlying dynamics and nature of changes as organizations move from stage to stage.

Nuclear power plants, chemical plants, and other high-hazard or high-reliability production systems (LaPorte & Consolini, 1991) provide interesting cases for examining organizational learning because they have distinct learning strategies (Weick et al., 1999) arising from the need to understand complex interdependencies among systems (Perrow, 1984), and avoid both potential catastrophes associated with trial-and-error learning (Weick, 1987) and complacency that can arise from learning only by successes (Sitkin, 1992). Weick et al. (1999) argue that maintaining high reliability requires mindfulness consisting of attention to hazards and weak signals (Vaughn, 1996), a broad action repertoire (Westrum, 1988), and a willingness to consider alternatives (March, Sproull, & Tamuz, 1991; Schulman, 1993). They theorize that such “inquiry and interpretation grounded in capabilities for action” (p. 91) is encouraged by distinctive organizational processes, including preoccupation with failure and reluctance to simplify interpretations.

## **2 STAGES OF ORGANIZATIONAL LEARNING**

Research specifically on organizational learning and more generally on organizational growth and development suggests a progression in structure, goals, skills, and culture. Whether we use a biological metaphor to talk about individual growth and learning (e.g., Rooke & Torbert, 1998) and organization life-cycles (e.g., Quinn & Cameron, 1983), or an historical analysis of organizational forms over time (Chandler, 1962; Perrow, 1970; Malone & Smith, 1988), we repeatedly find a progression in size, complexity, and interdependence with a more intrusive and unpredictable environment.

In the next section of the paper, we discuss a four-stage framework (see Figure 1) for organizational learning. These stages are presented as a provocative guide to analysis, not as a rigid model of development. In any organization, there will be examples of each stage in operation in different parts of the organization and at different moments in time. However, the latter stages require shared understanding and collaborative effort across the organization, so these capabilities must become relatively widespread and commonly enacted if they are to be sustained. “As Weber noted, ideal types are useful not because they are descriptively accurate – actual instances rarely evince all of the attributes of an ideal type – but because they serve as models that assist in thinking about social phenomena” (Barley & Kunda, 2001, p. 83).

The four stages are organized into a 2 X 2 table in Figure 1, representing two dimensions: (1) single- and double-loop learning (Argyris, Putnam, & Smith, 1985) and (2) improvisation and structure. In single-loop learning, goal-oriented actions are adjusted based on feedback to better achieve the same goal. In double-loop learning, a deeper process surfaces and challenges underlying assumptions and values regarding the

selection of that goal. Improvisation is a process of acting intuitively into an emergent situation rather than following structured procedures or plans (Weick, 1998). Both dimensions should be understood as overall tendencies or relative frequency: no organizations are purely single-loop or double-loop learners or only improvisational or structured.

-----  
Insert Figure 1  
-----

## 2.1 Local Stage

Most organizations begin their lives small, relatively unstructured, and personal or informal, like an entrepreneurial startup (Quinn & Cameron, 1983) or a craft shop (Perrow, 1970). Knowledge is based primarily on the experience and skill of individuals. Organization-specific and task-specific knowledge is local, contextual (Carlile, in press), tacit (Nonaka & Takeuchi, 1995), and sticky or hard to transfer (von Hippel, 1994). Exceptions occur frequently, and the organization relies on technical expertise to cope with surprises and provide flexibility or resilience (Wildavsky, 1988). Decisions are made locally by those steeped in the details, and learning mostly occurs locally as well. Learning is decentralized in individuals or workgroups and primarily single-loop, i.e., behaviors are adjusted after comparison to performance standards or benchmark models, but underlying structures and assumptions are not challenged. The organization is minimal and hardly self-aware.

For example, from the beginning of the nuclear power industry, design engineers appear to have understood plant construction as a finite project that results in a production machine. Once built and debugged, the plants were expected simply to run, a belief echoed by nuclear utilities and regulators: "Technological enthusiasts heading the AEC [Atomic Energy Commission] believed most accidents were too unlikely to worry about" (Jasper, 1990, p. 52). Given this belief, little attention was paid to "minor" problems in a plant or other plants in the industry, unless those problems affected production. When a combination of minor problems and operators doing what they were trained to do produced the Three Mile Island (TMI) event in 1979, this constituted a "fundamental surprise" (Lanir, 1986) for the nuclear power industry. The information needed to prevent the TMI event had been available from similar prior incidents at other plants, recurrent problems with the same equipment at TMI, and engineers' critiques that operators had been taught to do the wrong thing in particular circumstances, yet nothing had been done to incorporate this information into operating practices (Marcus, Bromiley, & Nichols, 1989). In reflecting on TMI, the utility's president Herman Dieckamp said,

To me that is probably one of the most significant learnings of the whole accident [TMI] the degree to which the inadequacies of that experience feedback loop... significantly contributed to making us and the plant vulnerable to this accident" (Kemeny, et al., 1979, p. 192).

In the local stage, information necessary for learning does not travel easily beyond particular workgroups and contexts.

## 2.2 Control Stage

Growth in terms of size and complexity is a major driver of formalization (e.g., Pugh et al, 1969). To achieve economies of scale, expertise is organized into workgroups and departments that often become classic “silos” of knowledge. To coordinate efficiently among workgroups and other subunits, organizations generate standard operating procedures and other formal routines to make work uniform and predictable, and facilitate communication (Nelson & Winter, 1981; Levitt & March, 1988). Controls are instituted to encourage uniformity, including accounting controls, procedure manuals, training programs, planning processes, and so forth. The “machine” metaphor and technical logic dominates (Carroll, 1998), such that performance is viewed as a summation of component-level, often explicit and measurable, contributions. Learning itself is understood as a set of routines for training, performance feedback, statistical process control (Sitkin et al, 1994), after action review, procedure revision, and so forth. This learning is directed at further control through exploitation of the known rather than exploration of the unknown (March, 1991), single-loop evolutionary enhancements rather than double-loop revolutionary changes (Argyris & Schon, 1996).

For most of its history, the nuclear power industry attempted to improve operations and prevent accidents through creation and enforcement of bureaucratic controls. Elaborate probabilistic analyses were used to anticipate (Wildavsky, 1988) all possible failure paths and to design physical and procedural barriers to these paths. When problems occurred, incident reviews typically identified actions that failed to comply with the rules, such as operators who did not follow procedures or engineers who made erroneous calculations (Carroll, 1995; Reason, 1990). This single-loop learning tended to focus on causes proximal to the problem, with available solutions that could easily be enacted, and were acceptable to powerful stakeholders (Carroll, 1995; Tetlock, 1983). Line managers wanted concrete solutions that “fix” problems (Carroll, 1998) and avoided costly or unpredictable actions: they disparage “trying to solve world hunger.” The suggested corrective actions were usually to strengthen control mechanisms (more training, more supervision, more discipline), create more rules (more detailed procedures, more regulatory requirements), or design hazards and humans out of the system (according to technical design rules, e.g., “inherently safe” nuclear reactor designs). Compliance to industry standards, professional standards, and procedure manuals was backed up by layers of internal and external monitoring and record keeping.

Learning activities were separated from everyday work as part of training or a staff function to analyze problems or utilize industry experience. Staff specialists were accountable for investigation activity and corrective action programs without giving them any real authority. Problems stimulated blame that undermined information flow and learning (Morris & Moore, 2000; O'Reilly, 1978). For example, an inspector criticized one plant after he discovered a set of informal records of problems without a plan to address each problem. As one manager at a well-respected plant stated, “NRC [the US

Nuclear Regulatory Commission] wants crisp problem identification and timely resolution.” The plant’s response was to stop documenting problems for which there were no immediate action plans, thus maintaining the illusion of “control” but decreasing their potential for learning. Learning-in-action would represent a loss of control, both because it would acknowledge the unfinished nature of current routines (Schulman, 1993) and because the learning process itself is decentralized and informal.

### **2.3 Open Stage**

Large, conservative, bureaucratic organizations can be highly successful in stable environments, but in turbulent and unpredictable environments, they do not learn or change fast enough. Bureaucratic controls over behavior fail when routines cannot be written and rewritten for all activities and when learning is restricted to specialized groups such as R&D. For example, Perron & Friedlander (1996) suggest that management systems for Process Safety Management “cannot yet be fully automated” (but notice the “yet”). Even when under pressure from new competitors with new products, rapidly-changing technologies and customer preferences, deregulation, and so forth, large organizations may initially ignore these threats (Freeman, 1999). Eventually, increased pressure and enlightened employees at various levels may open the organization to self-analysis, elaboration of bureaucratic mechanisms, and innovation (Quinn & Cameron, 1993).

The open stage is marked by a climate of psychological safety (Edmondson, 1999) encouraging organization members to ask questions, explore, listen and learn. Assumptions about authority and control give way to recognition of uncertainty and the need for collaborative learning. In the nuclear power industry, regulators and industry groups have long been calling for greater awareness of minor incidents and actions to avoid future trouble (Jackson, 1996; Rochlin, 1993). As Weick, et al. (1999) state, “to move toward high reliability is to enlarge what people monitor, expect, and fear.” Today, a typical nuclear power plant may identify over 2000 problems or incidents per year, 90% of which would have been ignored in the past. Although efforts to accelerate learning may include technological initiatives such as web-based information exchanges and databases of new ideas and best-practice routines (Pan & Scarbrough, 1998; Davenport & Prusak, 1997), the open stage is based on attitudes and cultural values of involvement, sharing, and mutual respect. This goes beyond the typical response of adopting others’ learning practices or buying consultants’ solutions in an effort to treat learning as another activity to fix and control (cf. Sitkin, et al, 1994).

Open stage plants are developing double-loop learning skills. Involvement of groups with different viewpoints provides feedback about varied assumptions and mental models and the impact of these assumptions on plant outcomes. With the stigma of reporting problems minimized, plants at the open stage are able to surface problems early and make the problems actionable (Carroll et al., in press). However, even the best plants still struggle with analyzing below the level of equipment problems, human error, and procedure inadequacies (Carroll, 1995, 1998, Carroll et al, in press). Despite a desire to

improve, investigators and managers seldom look for fundamental or deep, systemic causes because they lack ready-made actions to address such issues and ways of evaluating their success (remnants of the control stage). The skills and discipline of deeper learning often develop later than the willingness to learn.

## **2.4 Deep Learning Stage**

The final stage, as we envision it, would build upon the open stage by adding more capability for double-loop learning that promotes understanding of deep, systemic causes and creating a wider range of action possibilities to address such causes. Organizations at this stage would be capable not only of mutual respect across internal and external boundaries, but also skillful inquiry and facility to gain insights, challenge assumptions, and create comprehensive models (Argyris & Schon, 1996; Senge, 1990). Analyses would be based on facts but connect logically to systemic, organizational, cultural, and political viewpoints and experience with a repertoire of actions that can change these deep structures. Participants transcend component-level understanding and additive models of performance to develop systems thinking skills and more comprehensive mental models.

Deep learning practices are not widespread in the nuclear power and chemical industries. Carroll, Stermann, & Marcus (1998) relate one example of an innovative technique that introduced such practices to Du Pont chemical plants. As part of a company-wide cost-reduction effort, a benchmarking study showed that Du Pont spent more than its competitors on maintenance, yet had worse equipment availability. A culture of reactive fire-fighting had developed, with workers regularly pulled off jobs to do corrective maintenance. Responding to the benchmarking study, a series of cost-cutting initiatives were undertaken that had no lasting impact. Finally, one team questioned the basic assumption that reducing maintenance costs could help reduce overall manufacturing costs; they thought that the effects of maintenance activities were tightly linked to so many aspects of plant performance that no one really understood the overall picture.

Du Pont was able to improve maintenance only after a collaborative conceptual breakthrough. An internal team developed a dynamic model of the system of relationships around maintenance (a "modeling for learning" exercise with the assistance of a researcher/consultant, Senge & Stermann, 1991). However, they were unable to transmit the systemic lessons of the model through ordinary means. Instead, the team created an experiential game in which plant employees play the roles of functional managers and discover new ways to think about plant activities, share their experiences and ideas, and test programs and policies. Having a broad range of employees with a system-wide understanding of the relationships between operations, maintenance, quality, and costs laid the groundwork for a successful pump maintenance pilot program. In the deep learning stage, the organization has the ability to come up with alternative assumptions and models to guide action toward more desirable outcomes, experiment with the new ideas-in-action, and track feedback on their effectiveness.

In the remainder of this chapter, we present three case studies of organizations that made significant transitions from stage to stage. Presenting the cases as transitions clarifies the differences between stages and illustrates the challenges of changing behavior, emotions, and mental models. It also allows us to emphasize that organizations are now, and may always be, in transition. Following the case studies, we draw some lessons about learning-in-action and the stage model.

### **3 A SHIFT FROM LOCAL TO CONTROL STAGE**

In the first case study, a nuclear power plant investigated an incident in which an employee was seriously hurt. This plant was attempting to improve safety and performance in part by using a newly upgraded incident investigation process. The investigation created an opportunity to raise collective awareness about local work practices and helped managers strengthen controls and increase conformity to rules.

#### **3.1 Fall from Roof**

An electrical maintenance supervisor sent three men to replace light bulbs inside the “hot” machine shop, the area used to decontaminate equipment of radiological residue. The men headed off to the work area and discussed among themselves how to reach the light bulbs. They decided that one of them, whom we call Joe, would access the lights by climbing on the roof of a shed within the larger building. Joe and one coworker dressed in anti-contamination suits and propped a ladder against the shed wall. Joe crawled up the ladder and onto the roof. As he was about to reach the lights, one of the roof panels gave way, dumping him 10 feet to the ground below. His injuries included a broken scapula, a broken rib, three fractures to the small bones near the spine, a lacerated lung and arm. His coworkers used a nearby phone to call for help. EMTs arrived shortly and took Joe to the hospital.

#### **3.2 The Plant’s Interpretation**

For an event of this seriousness, a multi-discipline team was assembled to collect information, analyze causes, and make recommendations. The team noted that a number of standard operating procedures regarding safety assessment were not followed. When the electrical supervisor assigned three men to the job, no one was designated to be in charge. The supervisor did not conduct a pre-job brief (explaining the operational and safety issues involved in the job) and no one thought to walk down the job (conduct physical examination and discussion of the safety challenges at the work site) or plan the safest way to do the job. The workers failed to follow rules requiring fall protection (e.g., a harness attached to a fixed support) when working aloft and use of a folding ladder by unfolding it rather than leaning it against a wall.



The team's report noted that these actions and omissions may be part of a local culture of risk-taking. The tone of the task was set, in part, by the most senior electrical worker of the three and the only one who had changed these light bulbs before. He told the others that they would "love this job 'cause it's kind of tight up there." Based on their interviews with Joe and others, the investigators speculated that this challenge struck Joe, who had just transferred to this department, as an "opportunity to succeed." Lastly, the workers ignored warning signs that the job was not routine. Nobody blinked when Joe was advised to stay on the one and a half-inch steel framework of the building because it was the strongest part. Joe failed to reconsider the job when his hand slipped through a skylight and he nearly fell, shortly before slipping again and falling through.

The investigation team's report documented lack of compliance with established safety practices and suggested ways to enhance compliance with existing rules. The report concluded that:

The cause of the accident was a failure of the employee, the employee in charge, and the supervisor to properly follow the Accident Prevention Manual requirements for working in elevated positions. The hazards associated with the job were not properly assessed; a stepladder was improperly used, and fall protection was not used when climbing on a structure.

The report then recommended that the plant should: 1) raise sensitivity to safety on routine jobs by appointing a full-time safety person; require managers to communicate to supervisors and supervisors communicate to employees the plant's expectations regarding industrial safety; and require department managers to provide feedback to the plant manager on each department's safety issues; 2) make more detailed guidelines on working aloft available to employees; 3) consider instituting a company-wide program on "Working in Elevated Positions," and 4) counsel all employees involved in the incident.

### **3.3 Making the Transition Between Organizational Learning Stages**

The incident investigation illustrates the plant's effort to shift its learning orientation from local to control. The report highlights the failure of workers and first line supervisor to comply with existing rules and procedures. The corrective actions' aim to increase awareness and compliance with these rules by appointing a safety advocate, having superiors reinforce the safety message, and improving procedures. Information was generated about local work practices and compliance with rules that could be shared across groups, discussed openly, and used to institutionalize new work procedures. The focus is on changing actions to comply with rules in order to correct a mismatch between desired results (keep people safe) and actual results (Joe is hurt), i.e., single-loop learning (Argyris & Schon, 1996).

Single-loop learning is very compatible with a desire for control and with the norms of the engineering profession that have shaped many industries such as nuclear power (Rochlin & von Meier, 1994). Carroll (1995) argued that there is a "fixing" orientation dominated by analysis of complex situations into additive components, linear cause and

effect thinking, a search among known solutions, a belief in the adequacy of current understanding, and an assumption that “any error is avoidable through engineering design and managerial controls” (p. 187).

In a control-oriented organization, managers are judged by their lack of problems or the speed with which problems are resolved and control reasserted. Challenges to that control are threatening and become political issues (Carroll, 1995; Tetlock, 1983). The investigation process itself is “delegated participation” (Nutt, 1998), a frequently ineffective process in which representatives suggest solutions to managers who may resist implementation (Carroll, in press). One member of the investigation team commented, “When it was becoming apparent what the real problem was, I think the group became (temporarily) unsure where to go—what to do—it looked like a big step.” It appears that the report writers had cause for concern. Another team member reported, “We put together three different drafts and each time someone in upper management disagreed with what we wrote. Finally the plant manager stepped in and accepted our answer.”

The investigation team did not question their assumption that “compliance with safety rules will improve safety.” A focus on compliance distinguishes those who make the rules from those who are being controlled. There is a contest for control between managers and engineers who are labeled as strategists and designers of the plant and operators and maintenance people who are labeled as implementers and doers (Carroll, 1998; Schein, 1996). The rules can become an empty ritual as alienated workers withdraw from the learning process. Without the opportunity to challenge underlying assumptions about why they work the way they do and the chance to reshape work accordingly, employees tend to feel that the corrective actions are simply another layer of control imposed from the outside. The investigators did not ask double-loop learning questions such as, “What frames do supervisors and workers hold that would allow a casual approach to safety develop and endure?”; “How does the status and career advancement system contribute to a culture of risk taking?”; “What frames allowed management to have a design problem (lights in an unsafe place) exist for so long?”; or “How does the work system of separated functions and hierarchical authority inhibit mutual understanding?”

## **4 A SHIFT FROM CONTROL TO OPEN STAGE**

The second case describes an organization-wide change effort in response to a crisis that shut down a large nuclear power station and nearly bankrupted the utility. It is a stark reminder of the importance of people in technically-dominated companies.

### **4.1 The Millstone Turnaround**

In October 1996, the Millstone nuclear power station outside New London, Connecticut, received an unprecedented order from the US Nuclear Regulatory Commission (NRC) to

keep its plants closed until they could demonstrate a “safety conscious work environment.” The problem had come to public attention earlier through a cover story in Time magazine about harassment and intimidation of employees who brought safety concerns to management. An interviewee at Millstone (Carroll & Hatakenaka, 2001) labeled the management culture as “male... militaristic – control and command.” The NRC review (Hannon et al., 1996) concluded that there was an unhealthy work environment, which did not tolerate dissenting views and stifled questioning attitudes among employees, and therefore failed to learn and change. As the report said, “Every problem identified during this review had been previously identified to Northeast Utilities management... yet the same problems were allowed to continue.”

New senior management was brought in to reestablish the trust of regulators, the public, and employees. Investments were made in physical improvements and extensive documentation to meet rising industry standards, but a critical component was culture change. Employees needed to feel safe about reporting concerns, to believe that managers could be trusted to hear their concerns and to take appropriate action. Managers had to believe that employees were worth listening to and worthy of respect. In short, the underlying values had to change from control to openness and trust. It took over two years to shift the culture and learning stage of the plant, but in June 1998 the internal oversight groups and external regulators certified that Millstone could restart its largest unit, and a second unit would restart a year later (the smallest and oldest unit was permanently decommissioned) (see Carroll & Hatakenaka, 2001 for more details).

## **4.2 The Plant’s Interpretation**

In September 1996, the new CEO for Nuclear Power, Bruce Kenyon, set the scene for change by an address to all employees on his first day, in which he introduced his values: high standards, openness and honesty, commitment to do what was right and two-way communications. He immediately revamped the top management team and introduced a new employee concerns program.

His subsequent actions enacted and modeled openness and trust. Throughout the next months, Kenyon met regularly with small work groups and in large all-hands meetings to give information and encourage two-way communication: “It shocked them to get candid answers.” Upon hearing Kenyon say publicly at his first NRC meeting that he found the organizations “essentially dysfunctional,” an interviewee from the NRC remembers thinking, “here’s a fellow who at least recognizes the problem.” Based on recommendations from an employee task force redesigning the employee concerns program, Kenyon agreed to create an Employee Concerns Oversight Panel (ECOP) to have an independent voice and report directly to him. ECOP was staffed with passionate advocates who argued with each other and with management, but over time they evolved a workable role. The panel’s existence “sent a message to the work force that employees could act as oversight of management.”

Kenyon allowed himself to be fallible and to enlist participation. When two contractors were terminated on the grounds of poor performance and the Director of the Employee Concerns Program provided evidence that the terminations had been improper, Kenyon quickly reversed his decision. As one of his senior managers recalls about their working relationship, Kenyon “went along with all my recommendations. He didn’t always agree... [Sometimes he] swallowed hard.” He called upon employees to voice their public support for Millstone to counterbalance media criticism: “when are you going to say what you think?” An ad hoc employee group self-organized, gathered over 1500 signatures on a petition, attended public meetings, wrote to newspapers, and otherwise expressed their commitment to a management that trusted them to become part of the solution.

Individual managers told stories of personal transformations and how they came to understand the nature of the problems. The case of the operations vice president was perhaps the most dramatic. Typical of the old-style management, he was weary of “whiners,” and “didn’t believe anyone would harass someone who brought forth safety concerns.” When the two contractors were terminated and the employee concerns program offered their view that the terminations were improper, “It was one of those moments your perception changes... a watershed for me.” He also remembers vividly his visit with several other Millstone managers to another nuclear power plant that had made a dramatic turnaround, where he learned that safety concerns could make business sense.

Millstone was typical of an industry in which managers are “not high on people skills, for example, few can read nonverbal signals.” They had to appreciate that employees’ perception was their reality. For example, when members of the training and operations departments were disciplined for inaccuracies in training documentation two years earlier, employees immediately assumed that the former training director was being punished because he had been an outspoken critic of management. Management had failed to anticipate reactions or to minimize the impression of retaliation. Managers had to learn new skills, including sensitivity to their own and others’ emotions and perceptions. Through extensive new training programs and coaching by organizational development consultants, they had to “learn the difference between anger, hurt, and a chilling effect” and avoid confusing a fear of reprisal with a lack of confidence that management would take effective action.

Openness and trust emerged organically through multiple mechanisms and venues. We have already mentioned the Employee Concerns Program (ECP) that provided confidential ways to report issues for investigation and the Employee Concerns Oversight Panel (ECOP) that gave a direct connection between employee representatives and the CEO Nuclear. The Executive Review Board was created after the contractor terminations to review all disciplinary actions, comprising senior managers and an ECOP representative as an observer. By opening up the management process, it helped restore employee trust in management, and created an environment for managers to learn and enact new values. The People Team, a coordinating group among human resources, legal department, ECP, ECOP, management, and organizational development consultants, met daily to respond to problems and organize to address issues and monitor

progress. Internal Oversight groups and an independent third-party consulting group required by the NRC provided additional monitoring and advice. These multiple mechanisms and forums allowed broad participation so that managers and employees could share information, develop common language, learn by doing, and build trust by reacting well to challenges.

#### **4.3 Making the Transition Between Organizational Learning Stages**

The NRC requirement that Millstone develop a “safety conscious work environment” and demonstrate this to the satisfaction of an independent third-party consultant was unprecedented in the industry. The NRC offered no guidance. Millstone had to find its own way to move from a control stage characterized by centralized authority and mutual distrust to an open stage characterized by communication, trust, and participation.

Millstone managers were proud of Millstone’s excellent record in the industry, built on technical leadership of the industry. Financial distress following the completion of the third unit led to an atmosphere of cost-cutting and production emphasis; the heroes were the ones who got the job done without breaking the budget. When employees complained about technical problems or the external regulators criticized them for lack of documentation or growing backlogs of work, managers ignored them or blamed the messengers. Managers believed that Millstone’s design features and managerial controls were sufficient to operate the plant safely and reliably. Complaining employees and meddling regulators were annoying distractions that threatened their sense of control. In short, managers and employees lived in separate thought worlds (Dougherty, 1992) with strong cultural barriers and a perceived contest for control.

New senior management, external intervention, and an infusion of outside employees broke through some of that defensiveness. Management’s basic assumption that “we know everything we need to know” was challenged (cf. Schulman, 1993). And so was employees’ basic assumption that “management can’t be trusted.” Because senior management reacted well to critical events such as the contractor terminations and independent voices were allowed to challenge the status quo, double-loop learning occurred. Multiple venues emerged for managers and employees to talk together and work on the common problem of rebuilding Millstone. Managers began to listen and trust the employees enough to act on what was being said; in turn, employees began to feel safer about speaking out (Edmondson, 1999) and to trust that management would listen and take action. The most powerful way to regain trust is to work together with a common purpose (Kramer, 1999; Whitener, Brodt, Korsgaard, & Werner, 1998).

Managers not only became more open to information coming from employees and external observers, but also became aware of new kinds of information. Control-oriented managers, some of whom get their way by yelling and threatening, are generally unaware of their own emotionality and try to restrict any emotionality in their subordinates. They claim to value facts and rationality, even when they are using fear to exercise control. They did not realize that emotions and perceptions are reality. The more open

environment at Millstone marked an increase in interpersonal skills and emotional intelligence (Goleman, 1995). Emotions and perceptions could be anticipated, considered, discussed, and managed.

## **5 A SHIFT FROM OPENNESS TO DEEP LEARNING**

The third case was part of a larger, plant-wide effort to begin using root cause analysis (RCA) teams as a way to address, simultaneously, a recent history of financial losses, some dangerous incidents, and repeated equipment failures. The idea of using root cause analysis to address adverse incidents in the plant was introduced to plant management as a result of a merger with another petrochemical company, which used the RCA process already. Two headquarters staff at this petrochemical company had been working for a decade to promote more strategic and systemic thinking at operational and executive levels, using RCA as one of several approaches, and their progress was just beginning to accelerate. The Plant Manager had requested that they introduce their RCA practice to his plant with a training intervention. The plant decided to train about 20 plant employees, operators, maintenance staff, engineers, and first line supervisors to conduct RCAs by exploring, in-depth, four significant recent incidents.

Each incident investigation team included some members from inside the plant and some from outside, and at least one experienced root cause facilitator. The overall process included training in investigation, analysis, and reporting methods during the course of a three-week time frame, culminating in reports to plant management. Training was timed to correspond to the needs of the teams as they collected maintenance and operations logs, reviewed physical evidence, interviewed involved parties and knowledgeable experts, analyzed causes, and prepared reports.

### **5.1 Charge Heater Fire**

The charge heater fire investigation examined the explosion and fire in a charge heater that cost \$16 Million for lost production and repairs. Charge heaters are large gas-fueled burners used in the transformation of waste products from oil refining back into usable products through hydrocracking, a dirty and dangerous process requiring very high heat and pressure. The residue of this process is coke (coal dust) which can accumulate on the inside of heater tubes. In addition to unearthing causes of the explosion, plant managers also wanted to discover and ameliorate the conditions which led to this event and might lead to future events.

While the causal analysis presented below may seem extremely straightforward, its simplicity is the result of a rigorous and laborious root cause analysis process that involved four elements: A time line of events; an “Is/Is not” process that differentiates circumstances where the event occurred from similar circumstances where it did not; a detailed causal event diagram; and a process of categorizing the quality of data used to draw inferences in the causal event diagram (is it a verifiable fact, an inference, or a

guess). In doing these analyses, members of the team argued with each other, built on each other's ideas, and alternated between stunned amazement and appreciation at the differences in each other's views of the refinery.

## 5.2 The Plant's Interpretation

Distilling and analyzing the information available, the team concluded that the explosion and fire were due to a tube inside the charge heater that ruptured because the three quarter inch steel skin got too hot and tore. The team found that three factors contributed to the heater explosion: (1) high heat input, (2) low heat removal, and (3) unawareness on the part of operators of the actual tube skin temperature. First, operators ran the burners in the charge heater unevenly to increase heat in order to achieve the desired production level, while avoiding alarms that would signal an unsafe condition. Second, heat was removed more slowly than usual from the tube skin because coke had adhered to the inside of the tubes and was acting as an insulator. There was more coke than usual because it was assumed that a new decoking process worked as well as the previous process and no one had checked for coke build up. Third, the combination of running some tubes hotter (at a higher gas pressure) and the build-up of coke moved the high-heat point up the tube. The thermocouple meant to detect temperature on the tube skin, set at a height specified in the heater design, was now below the hottest part of the tube, so that operators believed the tube temperature was acceptable. The tube ruptured above the thermocouple.

The team noted as a "Key Learning" that plant staff made decisions without questioning assumptions that seemed to underlie them. First, the maintenance department changed decoking processes but did not know and never checked if the new process was effective. Second, operators increased the burner pressure in the charge heater but did not know the consequences of doing so. Third, operators changed the pattern of firing heater tubes (to fire hotter around the perimeter) but again did not know the consequences of doing so. On the basis of these insights, the team's lead recommendation for future action was that the plant identify "side effects" and be more aware of the broader "decision context" when changing production processes.

The team deepened their analysis as they discussed why assumptions about the effectiveness and safety of the new decoking process and the modified charge heater tube firing practices were never questioned at the time that changes were made. They speculated that their colleagues probably were unaware of the assumptions they were making. Our observations of the team's investigation and our *post hoc* interviews with team members highlight the team members' amazement and interest in "how quick we jump to conclusions about things." The team repeatedly mentioned the fact that, prior to learning the new investigation process, they rarely questioned their own conclusion-drawing processes and the assumptions that underlay them. One team member summarized his new approach by saying he now questions his co-workers: "I say, are you sure? Are you sure? Did you look at the initial aspects of what happened?"

As they worked on the investigation, the charge heater team frequently discussed their discovery of unanticipated and previously unknown interactions between apparently unrelated plant processes such as decoking and tube firing. When the team got to the bottom of their cause tree they noticed that each leg was a necessary but not sufficient contributor to the incident. In one of its verbal reports to other investigation teams during the training sessions, the charge heater team noted that, “We are seeing that several things combine over time to create an event.” Independent decisions by maintenance to change decoking, the inspection service to trust that the new decoking was effective, and operators to change burner tube firing practices ended up interacting to produce the heater fire. The team described their learning to other teams by saying, “It appears that in most cases there are elements of human factors (systems) that show up if you dig deep enough.”

Based on the insights from this team and from the other teams, the plant decided to implement a “Management of Change Process” to address the unanticipated side effects and interactions that caused problems. According to follow-up interviews with team members six months after their investigation, the actual results are mixed. One team member felt the plant Management of Change process had teeth:

“The biggest issue that came out [of the root cause analysis training] was management of change. MOC. Now people pay more attention to adhering to the MOC process. It may be that the RCA training helped focus attention on MOC. MOC is serious. It is real. If you don’t do it, your job is on the line. If not do it, have to explain why not.”

However, another team member felt, “There are no legs on the management of change effort. It is just a lot of talk.”

### **5.3 Making the Transition Between Organizational Learning Stages**

The charge heater investigation provides examples of an organization increasing both openness and deep learning. The independent decisions that changed decoking and heater tube firing practices illustrate aspects of the local stage of organizational learning. In our observations of the training session, it was evident that at least some participants were anxious about being open with colleagues in their own department or in other departments, or with management. Would operators talk to engineers? Would an operator working on this investigation be perceived as having sold out? Would managers listen to reports that were critical of their own behavior? The investigation could have blamed the operators for “getting around” the tube temperature alarms, ignored the role of management decisions about production goals, and instituted more monitoring and rules. A control approach to learning could have reinforced barriers to the open flow of information and discouraged participation, and failed to get at the underlying, systemic causes of the event.

However, plant management was not approaching its problems from the viewpoint of control. Instead, there was a desire to create more openness, and to demonstrate the value of openness and deep learning for achieving better performance. During the course of the



training and investigation, teams experienced more openness and collaboration than they expected. There was a willingness to confront reality and to surface underlying assumptions about “how we do work around here.” Support from a new plant manager helped encourage full participation. That support was itself an outcome of the training team who were working publicly with the investigation teams but privately meeting with management to reduce their defensiveness and enlist their visible engagement. And, it was evoked and reinforced by specific features of root cause analysis that require close attention to factual details, data quality, and cause-effect relationships.

The training team was very deliberate in bringing a deep learning approach to the plant. Their goal was to educate management by challenging their mental models with rich and compelling data and interpretations. The underlying concept is that managers establish the conditions for performance, i.e., they manage the system. Managers do not control behavior, but rather provide the resources (people, time, money, equipment, plans, opportunities, legitimacy, procedures, etc.) by which the system will operate.

The team investigation began to create deep learning when they started addressing operations at the plant from a systemic perspective and challenging assumptions. Paradoxically, the process of “drilling down” precisely and narrowly into causes of this incident allowed the team to develop new awareness of interdependencies across the system. They recognized interactions among components of the system and began to understand a central tenet of the quality movement (e.g. Goldratt & Cox, 1992) that working to optimize individual components does not automatically add up to an optimized system. Most importantly, they developed and practiced double-loop learning capabilities to recognize assumptions and mental models as separate from reality, understand that assumptions and mental models affect behaviors and outcomes, imagine alternative mental models to guide action toward more desirable outcomes, and take action with the new mental models (Friedman & Lipshitz, 1992; Argyris, et al., 1985).

The process of root cause analysis encouraged awareness of mental models and ability to work *on* mental models rather than *through* them. The team became aware of their own mental models and the distinction between model and reality, which is a necessary step to double-loop learning (Friedman & Lipshitz, 1992). In our interviews with team members, they universally highlighted the benefit of having a diverse team because of the surprising differences among people’s ways of looking at the same problem. The rigor of the root cause analysis process encouraged them to “hold their assumptions lightly” as the analysis held these views up to comparison and disconfirmation.

The cause-effect diagrams worked as a boundary object to help reveal tacit assumptions about plant processes that were key links in the causal chains leading to the heater explosion. The team’s cause-event tree included the three assumptions: “[desired] charge rate [desired production rate] dictates heater firing”; “sandjetting works as well as steam air decoking [to remove coke from inside heater tubes]”; and “there are no ‘hot spots’ [overheated areas] on the tubes.” Developing a gut sense that assumptions matter in shaping action and outcomes is important to overcome fears about “trying on” new mental models (Rudolph, Taylor, & Foldy, 2000). Their recommendation that

“identifying side effects and documenting decision context become a central part of decision making at [the plant]” implies a new insight: “Decisions made in one context may have side effects in other contexts and these are important to consider.”

## **6 Insights About Organizational Learning**

Our case studies and other data from the nuclear power and petrochemical industries (Carroll, 1998; Carroll et al., in press) suggest that transformations to double-loop learning and systemic thinking is difficult. It is rare to find companies that are skilled in deep learning. The most advanced of the companies we have studied are only beginning a transformation into deep learning, motivated by a few subversive visionaries and the intense pressures of competition and regulation. Anecdotal evidence suggests that companies adopt programs such as TQM and learning organization more to copy success stories and achieve legitimacy rather than through commitment and understanding. No wonder they rapidly move on to the next management fad.

### **6.1 Stages of Organizational Learning**

One possibility for unpacking the difficult transformation to deep learning is to recognize that there are multiple elements that must be brought together to support this transformation. Awareness and measurability are important, but so are cultural values and motives, supportive structures and resources, specific skills and tools, and concepts and mental models. The large, technologically-driven organizations that we have studied do not appear to transform all these elements at once or even to change them gradually. Instead, there is a natural order to their focus. Control through measurement, monitoring, incentives, and other traditional bureaucratic mechanisms seems to come naturally to managers and engineers (Carroll, 1998; Schein, 1996). However, the control stage can become a competency trap that inhibits learning (Levitt & March, 1988).

Most of the organizations we have studied seem to move out of the control stage by recognizing the limitations of top-down control and promoting more participation and open exchange of information throughout the organization and between the organization and the outside world. The “questioning attitude” and “safety culture” advocated in these industries are directed at acknowledging doubt (Schulman, 1993), increasing awareness or mindfulness (Weick et al., 1999), respecting the contributions of others, and placing a positive value on teamwork and learning. Such trust can only be developed by observations of the experience of courageous pioneers who take early risks to tell the truth. When open behavior is validated by others, trust is built and openness spreads in a virtuous cycle.

The open stage is also characterized by an awareness of people as different from machines. An ability to acknowledge emotions, conflicts, and different perceptions that underlie work relationships and political contests allows for discussion of the human side of organization. Of course, managers are uncomfortable and initially incompetent in this

domain, but openness to its importance builds a greater emphasis on managers, employees, and consultants with people skills. Over time, people learn by doing and through feedback from colleagues and coaches.

Openness to learning becomes linked to a discipline for learning in the transition to what we call the deep learning stage. The complexity and pace of change of modern organizations requires more than a desire to learn. Special circumstances for learning and concepts and techniques that make learning more efficient are needed to break through long-held assumptions and cognitive habits. Deep learning is not simply the use of particular techniques such as root cause analysis. There are many versions of “root cause analysis,” most of which are used with minimal training to find and fix problems (Carroll, 1995) rather than to challenge deep assumptions with rigorous and systemic thinking, just as TQM can be used for control rather than learning (Sitkin et al., 1994). It is not particular tools such as root cause analysis that lead to learning, but rethinking actions and assumptions in the context of new *concepts* that underlie the tools, such as data quality, rigorous cause-effect connections, systems thinking, mutual respect across groups, insight into personal and political relationships, and double-loop learning. The tools and the learning activities are only an opportunity to have new conversations, enact new behaviors, develop new skills, and build new relationships.

## **6.2 Learning-in-Action**

The cases reinforce the importance of learning through action. Although some kinds of knowledge are represented explicitly (numbers, words) and easy to store and transfer, many kinds of knowledge are difficult to represent or separate from their context. These kinds of knowledge have to be reconstructed by users, improvised, tried out and modified to suit the occasion. Enactment is problematic, and learning is in the doing, individually and collectively. Learning is in the connection between action and reflection, in making knowledge actionable (implementation) and action knowledgeable (sensemaking) (Argyris, et al, 1985; Crosson, et al., 1999).

From this viewpoint, incident investigations involve both access to information and socially constructive processes such as imagination and negotiation. The collective analysis of factual details (with a disciplined logic that identified gaps) helped to drive a systemic understanding. The cause-effect diagrams were boundary objects (Carlile, in press) negotiated by the team in a process of knowing (Cook & Brown, 1999) that helped surface previously unarticulated mental models of the work environment, compare them, and arrive at new, shared views. Some of the learning was articulated in the written report, another boundary object negotiated between the team and managers that initiates corrective actions and feeds databases, but much remained unwritten (although discussed as part of the reporting out process).

Experience with learning cycles increased tolerance for short-term difficulties and occasioned resource shifts away from production toward learning. A systemic view suggests that changes take time to unfold, that things get worse before they get better

(since resources are shifted away from immediate needs), and that leverage points must be identified for selective investment in changes that are not simply ceremonial but actually transform practice. Problems are not simply someone's fault, but rather a feature of the system; altering that system takes deep understanding of which way to go and mobilization of broad support. This is more than "controlling" people. A good system may be difficult to understand; its principles may be hard to verbalize yet possible to learn through action or instruction. For example, a rigid grip of a rowing oar may increase the feeling of control but decrease absorption of the shock of uneven waters, thereby decreasing actual control. Managers may use "heavyhanded" incentives and authority to increase their feeling of control and drive noncompliance out of sight, simultaneously increasing the discrepancy between rules and actual behavior.

In summary, we have argued for the importance and difficulty of learning from experience. Nuclear power plants and chemical plants are challenged by the hazards in their work processes to learn from problems and to overcome barriers to learning. The history of these industries and the case studies we have examined suggest that there is a common progression from local learning to a control orientation associated with single-loop learning, which is then held in place by managerial and professional culture. Yet problems continue to occur and many organizations seek to be more proactive by becoming a learning organization, which incorporates mutually-reinforcing elements of attitudes and thinking patterns. Our results suggest that, to some degree at least, attitudes favorable to learning precede double-loop learning skills. The concepts and skills of deep learning seem to be difficult to master and to require significant commitment, discipline, and learning-in-action. Future research will undoubtedly put more flesh on the bones of this framework, and contribute alternative ways to think about organizational learning.

## REFERENCES

- Argote, L. & Ingram, P. (2000). 'Knowledge transfer: A basis for competitive advantage in firms'. *Organizational Behavior and Human Decision Processes*, 82, 150-69.
- Argyris, C., Putnam, R., & Smith, D. M. (1985). *Action Science: Concepts, Methods and Skills for Research and Intervention*. San Francisco: Jossey-Bass.
- Argyris, C. & Schon, D. (1996). *Organizational Learning II: Theory, Method, and Practice*. Reading MA: Addison-Wesley.
- Barley, S. R. & Kunda, G. (2001). 'Bringing work back in'. *Organization Science*, 12, 76-95.
- Carlile, P. R. (in press). 'A pragmatic view of knowledge and boundaries: Boundary objects in new product development'. *Organization Science*.
- Carroll, J. S. (1995). 'Incident reviews in high-hazard industries: Sensemaking and learning under ambiguity and accountability'. *Industrial and Environmental Crisis Quarterly*, 9, 175-197.
- Carroll, J. S. (1998). 'Organizational learning activities in high-hazard industries: The logics underlying self-analysis'. *Journal of Management Studies*, 35, 699-717.
- Carroll, J. S. & Hatakenaka, S. (2001). 'Driving organizational change in the midst of crisis'. *MIT Sloan Management Review*, 42, 70-9.
- Carroll, J. S., Rudolph, J. W., Hatakenaka, S., & Wiederhold, T. L. 'Learning in the Context of Incident Investigation Team Diagnoses and Organizational Decisions at Four Nuclear Power Plants'. In G. Klein & E. Salas (eds.) *Naturalistic Decision Making*. Lawrence Erlbaum, in press.
- Carroll, J. S., Sterman, J., & Marcus, A. A. Losing the Maintenance Game: How Mental Models Drive Organizational Decisions. In R. N. Stern & J. J. Halpern (Eds) *Debating Rationality: Nonrational Aspects of Organizational Decision Making*. Ithaca, NY: Cornell University ILR Press, 1998.
- Chandler, A. D. (1962). *Strategy and Structure*. New York: Doubleday.
- Cohen, W. M. & Levinthal, D. A. (1990). 'Absorptive capacity: A new perspective on learning and innovation'. *Administrative Science Quarterly*, 35, 128-152.
- Cook, S. D. N. & Brown, J. S. (1999). 'Bridging epistemologies: The generative dance between organizational knowledge and organizational knowing'. *Organization Science*, 10, 381-400.
- Daft, R. L. & Weick, K. E. (1984). 'Toward a model of organizations as interpretation systems'. *Academy of Management Review*, 9, 284-295.
- Davenport, T. & Prusak, L. (1997). *Working Knowledge*. Boston, Harvard Business School Press.
- Dougherty, D. (1992). 'Interpretive barriers to successful product innovation in large firms'. *Organization Science*, 3, 179-202.
- Edmondson, A. (1999). 'Psychological safety and learning behavior in work teams'. *Administrative Science Quarterly*, 44, 350-383.
- Freeman, S. F. (1999). 'Identity maintenance and adaptation: A multilevel analysis of response to loss'. *Research in Organizational Behavior*, 21, 247-294.
- Freidman, V. J. & R. Lipshitz (1992). 'Teaching people to shift cognitive gears: Overcoming resistance on the road to Model II'. *Journal of Applied Behavioral Science*, 28(1), 118-136.

- Goldratt, E. M. & Cox, J. (1992). *The Goal: A Process of Ongoing Improvement*. Great Barrington, MA: North River Press.
- Goleman, D. (1995). *Emotional Intelligence*. New York: Bantam.
- Gruenfeld, D. H., Martorana, P. V., & Fan E. T. (2000). 'What do groups learn from their worldiest members? Direct and indirect influence in dynamic teams'. *Organizational Behavior and Human Decision Processes*, 82, 45-59.
- Jackson, S. A. (1996). *Challenges for the Nuclear Power Industry and its Regulators: The NRC Perspective*. Speech presented at the Regulatory Information Conference, Washington, D. C., April 9.
- Jasper, J. M. (1990). *Nuclear Politics: Energy and the State in the United States, Sweden, and France*. Princeton, NJ: Princeton U.
- Kemeny, J. G., Babbitt, B., Haggerty, P. E., Lewis, C., Marks, P. A., Marrett, C. B., McBride, L., McPherson, H. C., Peterson, R. W., Pigford, T. H., Taylor, T. B., and Trunk, A. D. (1979). *Report of the President's Commission on the Accident at Three Mile Island*. New York: Pergamon.
- Jehn, K. A., Northcraft, G. B., & Neale M. A. (1999). 'Why differences make a difference: A field study of diversity, conflict, and performance in workgroups'. *Administrative Science Quarterly*, 44, 741-63.
- Kim, D. H. (1993). 'The link between individual and organizational learning'. *Sloan Management Review*, 35, 37-50.
- Kolb, D. A. (1984). *Experiential Learning as the Source of Learning and Development*. Englewood Cliffs, NJ: Prentice-Hall.
- Kramer, R. M. (1999). 'Trust and Distrust in Organizations: Emerging Perspectives, Enduring Questions'. *Annual Review of Psychology*, 50, 569-98.
- LaPorte, T. R. & Consolini, P. M. (1991). 'Working in practice but not in theory: Theoretical challenges of "High-Reliability Organizations"'. *Journal of Public Administration Research*, 1(1), 19-47.
- Lanir, Z. (1986). *Fundamental Surprise*. Eugene, OR: Decision Research.
- Levitt, B. & March, J. G. (1988). 'Organizational learning'. *Annual Review of Sociology*, 14, 319-340.
- Malone, T. W. & Smith, S. A. (1988). 'Modeling the performance of organizational structures'. *Operations Research*, 36, 421-36.
- March, J. G. (1991). 'Exploration and exploitation in organizational learning'. *Organization Science*, 2, 71-87.
- March, J. G., Sproull, L. Tamuz, M. (1991). 'Learning from samples of one or fewer'. *Organization Science*, 2, 1-13.
- Marcus, A. A., Bromiley, P., and Nichols, M. (1989). *Organizational Learning in High Risk Technologies: Evidence From the Nuclear Power Industry*. Minneapolis: U. Minnesota Strategic Management Research Center, Discussion Paper #138.
- Morris, M. W. & Moore, P. C. (2000). 'The lessons we (don't) learn: Counterfactual thinking and organizational accountability after a close call'. *Administrative Science Quarterly*, 45, 737-65.
- Nelson, R. & Winter, S. (1981). *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard U.

- Nonaka, I., and Takeuchi, H. (1995). *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. New York: Oxford University Press.
- Pan, S. L. & Scarbrough, H. (1999). 'Knowledge management in practice: An exploratory case study'. *Technology Analysis and Strategic Management*, 11, 359-74.
- Perrow, C. (1970). *Organizational Analysis: A Sociological View*. Belmont, CA: Wadsworth.
- Perrow, C. (1984). *Normal Accidents*. New York: Basic Books.
- Perron, M. J. & Friedlander, R. H. (1996). 'The effects of downsizing on safety in the CPI/HPI'. *Process Safety Progress*, 15(1), 18-25.
- Pugh, D. S., Hickson, D. J., Hinings, C. R., & Turner, C. (1969). 'The context of organization structures'. *Administrative Science Quarterly*, 14, 19-114.
- Quinn, R. E. & Cameron, K. (1983). 'Organizational life cycles and shifting criteria of effectiveness: Some preliminary evidence'. *Management Science*, 29, 33-51.
- Rasmussen, J. & Batstone, R. (1991). *Toward Improved Safety Control and Risk Management*. Washington, D.C.: World Bank.
- Reason, J. (1990). *Human Error*. New York: Cambridge U.
- Rochlin, G. I. (1993). 'Defining "high reliability" organizations in practice: A taxonomic prologue'. K. H. Roberts (Ed), *New Challenges To Understanding Organizations*. New York: Macmillan.
- Rochlin, G. I. & von Meier, A. (1994). 'Nuclear power operations: A cross-cultural perspective'. *Annual Review of Energy and the Environment*, 19, 153-187.
- Rudolph, J. W., Taylor, S. S., & Foldy, E. G. (2000). 'Collaborative off-line reflection: A way to develop skill in action science and action inquiry'. In P. Reason & H. Bradbury (Eds), *Handbook of Action Research*. Thousand Oaks: Sage.
- Schein, E. H. (1992). *Organizational Culture and Leadership*, 2nd ed. San Francisco: Jossey-Bass.
- Schein, E. H. (1996). 'The three cultures of management: Implications for organizational learning'. *Sloan Management Review*, 38, 9-20.
- Schulman, P. R. (1993). 'The negotiated order of organizational reliability'. *Administration and Society*, 25, 353-72.
- Sitkin, S. B. (1992). 'Learning Through Failure: The Strategy of Small Losses'. *Research in Organizational Behavior*, 14, 231-266.
- Sitkin, S. B., Sutcliffe, K. M., & Schroeder, R. G. (1994). 'Distinguishing control from learning in total quality management: A contingency perspective'. *Academy of Management Review*, 18(3), 537-64.
- Tetlock, P. E. (1983). 'Cognitive style and political ideology'. *Journal of Personality and Social Psychology*, 45, 118-126.
- Torbert, W. (1994). Cultivating post-formal adult development: Higher stages and contrasting interventions. In M. Miller & S. Cook-Greuter (Eds), *Transcendence and Mature Thought in Adulthood: The Further Reaches of Adult Development*. Lanham, MD: Rowman & Littlefield.
- Vaughn, D. (1996). *The Challenger Launch Decision*. Chicago: U. Chicago.
- von Hippel, E. "Sticky information" and the locus of problem solving: Implications for innovation'. *Management Science*, 1994, 40(4), pp. 429-439.

- Weick, K. E. (1987). 'Organizational culture as a source of high reliability'. *California Management Review*, Winter, 112-127.
- Weick, K. E. (1995). *Sensemaking in Organizations*. Thousand Oaks, CA: Sage.
- Weick, K. E. (1998). 'Improvisation as a mindset for organizational analysis'. *Organization Science*, 9, 543-55.
- Weick, K. E., Sutcliffe, K. M. & Obstfeld, D. (1999). 'Organizing for high reliability: Processes of collective mindfulness'. *Research in Organizational Behavior*, 21, 81-123.
- Westrum, R. (1988). *Organizational and inter-organizational thought*. Paper presented at the World Bank Conference on Safety Control and Risk Management.
- Whitener, E. M., Brodt, S. E., Korsgaard M. A., & Werner, J. M. (1998). 'Managers as Initiators of Trust: An Exchange Relationship Framework for Understanding Managerial Trustworthy Behavior'. *Academy of Management Review*, 23, 513-30.
- Wildavsky, A. (1988). *Searching for Safety*. New Brunswick, NJ: Transaction Press.



Figure 1  
The Four Stages of Organizational Learning

