

# **LEARNING FROM EXPERIENCE IN HIGH-HAZARD ORGANIZATIONS**

John S. Carroll, MIT Sloan School of Management

Jenny W. Rudolph, Boston College Carroll School of Management

Sachi Hatakenaka, MIT Sloan School of Management

Submitted to Research in Organizational Behavior

## Abstract

Learning from experience, the cyclical interplay of thinking and doing, is increasingly important as organizations struggle to cope with rapidly changing environments and more complex and interdependent sets of knowledge. This paper confronts two central issues for organizational learning: (1) how is local learning (by individuals or small groups) integrated into collective learning by organizations? and (2) what are the differences between learning practices that focus on control, elimination of surprises, and single-loop incremental “fixing” of problems with those that focus on deep or radical learning, double-loop challenging of assumptions, and discovery of new opportunities? We articulate these relationships through an analysis of learning practices in high-hazard organizations, specifically, problem investigation teams that examine the most serious and troubling events and trends in nuclear power plants and chemical plants. Our analysis suggests a four-stage model of organizational learning reflecting different approaches to control and learning.

# LEARNING FROM EXPERIENCE IN HIGH-HAZARD ORGANIZATIONS<sup>1</sup>

John S. Carroll, MIT Sloan School of Management

Jenny W. Rudolph, Boston College Carroll School of Management

Sachi Hatakenaka, MIT Sloan School of Management

Submitted to Research in Organizational Behavior

Organizational learning has become a familiar yet controversial concept (Argyris & Schön, 1996; Fiol & Lyles, 1985; Mirvis, 1996). While mindful of the dangers of personification, we treat organizations as learning entities to emphasize particular capabilities and processes. This approach may be especially fruitful and timely as organizations struggle to cope with rapidly changing environments and more complex and interdependent sets of knowledge (Aldrich, 1999; Weick, Sutcliffe, & Obstfeld, 1999). The concept encourages us to confront two central issues: (1) how do organizations integrate local learning (by individuals or small groups) into collective learning (e.g., Crossan, Lane, & White, 1999; Kim, 1993)? and (2) what learning practices help organizations confront the potentially incompatible goals of control/reliability and innovation/discovery (e.g., Argyris & Schön, 1996; Carroll, 1998; March, 1991; Miner & Mezias, 1996; Sitkin, Sutcliffe, & Schroeder, 1994)? We

---

<sup>1</sup> This research was supported by National Science Foundation grants SBR96-1779 and SBR98-11451 and a Marvin Bower Fellowship at the Harvard Business School. Requests for reprints should be sent to the first author, Cambridge, MA 02142. We greatly appreciate the cooperation of the nuclear power and chemical plant sites and the staff of their corrective action programs. Assistance in collecting, coding, analyzing, and interpreting data was provided by Marcello Boldrini, Deborah Carroll, Christopher Resto, Khushbu Srivastava, Annique Un, and Theodore Wiederhold. Dr. William Corcoran provided invaluable insights into the investigation process, the nuclear industry, and the report coding process. We thank Vicki Bier, Bill Corcoran, Mary Crossan, Amy Edmondson, Alfie Marcus, Jim Reason, Nelson Repenning, and Barry Staw for comments on earlier drafts.

articulate these relationships through an analysis of particular learning practices in high-hazard organizations that face dual operational challenges of reliability and safety.

## **BACKGROUND**

### Learning and Knowledge

Learning is typically understood as a description of *individual* human behavior. Humans evolved as adaptive learners with few predetermined behavioral routines beyond what is needed to sustain life. The individual learning process builds upon proven biological principles (cf. Aldrich, 1999; Cambell, 1969): repeat what has been successful (“law of effect,” Thorndike, 1911); make occasional errors that may reveal new opportunities (March, 1991; Weick, 1995); and combine behaviors into more complex routines, while minimizing attention and other scarce cognitive resources (cf., bounded rationality, March & Simon, 1958). Additional characteristics speed up the human learning process: seek novelty (Berlyne, 1960) and imagine possibilities (allowing feedforward<sup>2</sup> predictions, counterfactual thinking, and virtual learning (March, Sproull, & Tamuz, 1991; Morris & Moore, 2000).

In numerous ways, human learning is essentially *social*. We observe others (Bandura & Walters, 1963), get feedback from them, participate in systems of interpersonal interaction (Weick & Roberts, 1993), and use language and other socially constructed conceptions and objects to drive our imagination and facilitate the spread of ideas and practices (e.g., Carlile, in press; Lave & Wenger, 1991).

---

<sup>2</sup> We use feedforward to mean predicting future events from a model or theory of action and feedback to mean comparing outcomes to expectations (e.g., Rasmussen, 1990; Reason, 1997). Crossan et al. (1999) use the same terms differently: feedforward is communicating knowledge from individual exploration to the organization, and feedback is institutionalizing routines to exploit knowledge.

In this context, we define *learning* as a change in linkages between antecedent conditions and imagined or enacted behaviors, and *organizational learning* as an analogous change at an organizational level. This is similar to Argyris & Schön's (1996) definition of theories of action as propositions of the form "if you intend to produce consequence C in situation S, then do [action] A" (p. 13). We preserve the form of these propositions but relax the focus on intentional learning to acknowledge that learning can occur without intention or awareness, and even without observable action (Glynn, Lant, & Milliken, 1994).

Consider, for example, a factory producing some amount of product per unit time. During a visit to another factory, organization members observe that similar machines can be run at higher speeds. Yet, after returning to their factory, production remains at the same rate. Although individual human beings are naturally programmed to learn, organizations are not. Learning may be inhibited by adherence to traditions or bosses who insist, "this is the way we do things around here" (Oliver, 1992; Schein, 1992). Until external pressure, a vision, or an intrinsic motive engages new behaviors, there may be no measurable change in performance.

Nor does learning have to be an improvement (Crossan, et al., 1995; Repenning & Sterman, 2000; cf. superstitious learning, Levitt & March, 1988). The factory may speed up production in response to competition, yet morale may erode, quality may drop, machines may break down, and the factory may ultimately lose its customers. If management decides instead to reorganize the plant into a lean production system rather than simply speeding up the machines and the people, then the factory would need time to try out the new actions and coordinate the various components as it climbs a learning

curve (Argote, 1999). Coordinating more unfamiliar and context dependent actions by more actors requires an iterative process of planning, acting, observing outcomes, analyzing and imagining, and adjusting (Argyris & Schön, 1996; Crossan et al., 1999; Daft & Weick, 1984; Kolb, 1984; Nonaka & Takeuchi, 1995).

Whereas learning is a *process* of change, the *content* of that process, the condition-action linkages, is *knowledge* (broadly construed to include explicit information, tacit know-how, etc.). 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). Different parts of the organization, such as plant operators and corporate executives, “know” different things about how work is done, express their knowledge in different languages (Dougherty, 1992), and keep knowledge in different reservoirs (Argote & Ingram, 2000). Putting this knowledge to use requires combining component-level knowledge and filling gaps by improvisation (Weick, 1998).

Organizational learning activities engage complex interdependencies across people and groups (Crossan, et al, 1999; Kim, 1993). 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). These bridging practices are supported by both interpersonal skills such as conflict management (Jehn, Northcraft, & Neale, 1999) and networking across boundaries (Yan & Louis, 1999), cognitive skills such as logical analysis and

systems thinking (Senge, 1990), and skills that involve both, such as collaborative inquiry (Isaacs, 1999).

### Organizational Learning and Bureaucratic Control

Theories of organizational learning often contrast incremental and radical learning (Miner & Mezias, 1996). Argyris, Putnam, & Smith (1985) separated single-loop learning that adjusts actions to reduce the gap between desired and actual results, from double-loop learning that challenges the appropriateness of goals and the assumptions and mental models that underlie actions and expectations. March (1991) identified competing learning strategies: exploitation of familiar practices and exploration of new and uncertain possibilities. Organization change theories analogously separate infrequent major reorientations that reshape the fit between the organization and its environment from incremental consolidation that aligns components for smooth operation (Tushman & Romanelli, 1986; see also Orlikowski & Tyre, 1994, regarding introduction of new technologies in punctuated stages of adaptation).

On its face, incremental learning seems consistent with increasing efficiencies in organizations whereas radical learning is paradoxical or threatening to the bureaucratic “iron cage” of rules, controls, and entrenched interests. Sitkin et al. (1994) expanded on March’s (1991) exploitation-exploration distinction in contrasting total quality management (TQM) practices that are implemented within a control orientation with TQM practices in a learning regime. The control orientation seeks to maintain predictable operations, minimize variation, and avoid surprises. The learning orientation seeks to increase variation in order to explore opportunities. Weick and Sutcliffe (2001) argued that investments in plans and routines could restrict an organization’s ability to

sense discrepancies, update knowledge, and act flexibly. Amabile (1996) and O'Reilly and Chatman (1996) agreed that formal bureaucratic controls undermine the intrinsic motivation needed for creative and flexible responses to uncertainty. Senge (1990) contended that humanistic values of caring and individual freedom are essential to building learning organizations.

However, organization theorists since Weber (1978) have insisted that bureaucratic elements of hierarchy and formalization are not inherently linked to control and centralization of power. Perrow (1979) makes a strong case that abuses of power or "specific acts of maladministration" (p. 56) give bureaucracy its negative image of control. Bureaucracies drift toward rigidity and domination because efficient systems and political agreements grow around them. Gouldner (1954) asserted that there are both punitive bureaucracies "based on the imposition of rules, and on obedience for its own sake" and representative bureaucracies "based on rules established by agreement" (p. 24) that encourage worker participation and expertise while benefiting from standardization and coordination of work (although bottom-up controls can become coercive, Barker, 1993). Adler and Borys (1996) make a similar distinction between coercive and enabling bureaucracies. Pugh and Hickson (1976) found empirical evidence that formalization of tasks and centralization of power are uncorrelated across organizations. Adler (1991) offers a detailed case history of a "learning bureaucracy" at the NUMMI joint venture of Toyota and GM that combined highly structured work practices with high worker involvement and continuous improvement.

The conceptual conflation of bureaucracy and domination is supported by commonly shared assumptions about organizations and individuals. For example, in

separate analyses of the TQM movement, Cole (1998) and Repenning and Sterman (2000) showed that managers initially ignored and later misapplied process improvement techniques because they assumed that quality had to cost more, that improvements had to meet a cost-benefit hurdle, that Japanese TQM succeeded because of special characteristics of Japanese workers, that managers and engineers have expert knowledge but not front-line workers, and that problems are due to people rather than systems of production. These assumptions become self-confirming when managerial actions based on mistrust of workers and attention to short-term production goals produce further conflict and rigid control structures (Repenning & Sterman, 2000; cf. Adler, 1991).

Most empirical demonstrations of learning within large, bureaucratic organizations involve repetitive work with measurable outcomes, such as learning curves (e.g., Argote, 1999). Studies of TQM as a learning process, such as those mentioned above, typically focus on structured tasks where inventory can be counted, manufacturing defects are apparent, and statistical techniques are readily applied. In contrast, the management of safety in high-hazard industries such as nuclear power or aviation requires the avoidance of rare events. Weick (1987) calls safety “a dynamic non-event” (p. 118) because uneventful outcomes are produced by “constant change rather than continuous repetition” (Reason, 1997, p. 37). High-hazard organizations learn from experience in a context of potential disaster and therefore have special reasons to manage effectively the relationships between control and learning.

#### High-Hazard Organizations

High-hazard organizations such as nuclear power plants and chemical plants have been an important focus of organizational research since the seminal books by Turner

(1978) and Perrow (1984). High-hazard organizations are distinctive work settings that include potential harm or death to large numbers of individuals in a single event, such as an explosion or crash. Theory about high-hazard organizations developed further with the work of Rasmussen (1990) and Reason (1990) in psychology, LaPorte and Consolini (1991) and Wildavsky (1988) in political science, and Roberts (1990) and Weick (1987) in management.

The special importance of learning in high-hazard organizations was recognized early in both the research and policy literatures (e.g., the investigation of Three Mile Island, Kemeny et al., 1979). As Weick (1987) stated, “organizations in which reliability is a more pressing issue than efficiency often have unique problems in learning and understanding” (p. 112). Such organizations develop 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). Organization theorists argue vigorously over whether high-hazard organizations are distinctive because of the inherent normalcy of accidents (Perrow, 1984; Sagan, 1993; Vaughn, 1996) or because they achieve “high reliability” through special organizational features that allow people to handle hazardous tasks with remarkably few problems (LaPorte & Consolini, 1991; Roberts, 1990; Weick et al., 1999).

Importantly, these organizations undergo intense public scrutiny, particularly when things go wrong. Investigations of the Three Mile Island (Perrow, 1984), Bhopal (Srivastava, 1987), and Challenger (Vaughn, 1996) accidents, for example, provided rich

databases for researchers. Of course, investigators such as Sagan (1993) had to use the Freedom of Information Act to gain access to information on military nuclear weapons mishaps. The post-hoc analysis of accidents has produced a fascinating body of research, limited as it is by a reliance on investigations by others and biases in selection and hindsight (Woods, et al., in press). On-the-ground fieldwork has been more unusual as illustrations and case studies are gathered slowly (e.g., Bourrier, 1999; Perin, 1998; Roberts, 1990).

High-hazard organizations live on the boundary between maintaining control and learning (Sitkin et al., 1994) or exploiting current capabilities while exploring unfamiliar possibilities (Crossan et al, 1999; March, 1991). High-hazard organizations try to anticipate and defend against problems while responding resiliently to surprises (Wildavsky, 1988). On the one hand, such organizations must comply with a large body of regulations and rules to avoid accidents yet, on the other hand, the rules cannot predict every eventuality and humans must still improvise and learn in the midst of action. 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, et al., 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. They further argue that more and more organizations in our fast-paced world are beginning to resemble high-reliability organizations.

High-hazard organizations are strongly compliance oriented – they are heavily regulated and concerned with reliable operations and prevention of accidents rather than exploration of new opportunities (March, 1991; Sitkin et al., 1994). For example, since its founding, the nuclear power industry attempted to improve operations and prevent accidents through creation and enforcement of bureaucratic controls (a similar story could be told for many industries). Although all organizations generate standard operating procedures and other formal routines to make work predictable and facilitate coordination (Nelson & Winter, 1981; Levitt & March, 1988; Pugh et al., 1969), “the managers of hazardous systems must try to restrict human actions to pathways that are not only efficient and productive, but also safe” (Reason, 1997, p. 49). Elaborate probabilistic analyses (e.g., US Nuclear Regulatory Commission, 1975) are used to anticipate (Wildavsky, 1988) possible failure paths and to design physical and procedural barriers such as redundant safety systems, detailed procedures, training, and supervision.

Layers of internal and external monitoring and record keeping help enforce compliance with procedures, codes, and standards. Deviations are dealt with by evolutionary enhancements, including more controls: “Safe operating procedures... are continually being amended to prohibit actions that have been implicated in some recent accident or incident” (Reason, 1997, p. 49). Performance is understood as the absence of deviation or error, a *prevention* focus that is associated with anxiety, loss aversion, avoidance of errors of commission, and a strong moral obligation to comply with rules (Higgins, 1998). Learning is understood as a set of routines for training, performance feedback, statistical process control (Sitkin et al, 1994), after action review, procedure revision, and other forms of incremental improvement (Miner & Mezas, 1996).

Corrective actions usually 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).

It is very challenging for organizations intent upon compliance to develop a full range of learning capabilities because assumptions underlying compliance and control can be in conflict with efforts to learn (Carroll, 1995, 1998; Sitkin et al., 1994).

Engineers and managers are trained to plan, analyze complex situations into understandable pieces, avoid uncertainty, and view people as a disruptive influence on technology or strategy (Rochlin & von Meier, 1994; Schein, 1996). Problems stimulate blame that undermines information flow and learning (Morris & Moore, 2000; O’Reilly, 1978). For example, an inspector from the US Nuclear Regulatory Commission (NRC) 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 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 ignoring the unfinished nature of current routines (Schulman, 1993) and maintaining the *illusion* of control (Carroll, 1995; Langer, 1975), while possibly decreasing the potential for learning.

Yet high-hazard organizations can develop capabilities for challenging common assumptions and exploring new possibilities, as exemplified by an innovative approach at Du Pont (Carroll, Sterman, & Marcus, 1998). 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 reactive culture 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 & Sterman, 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 how programs and policies can improve performance. 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. With enhanced learning capabilities, the organization had 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.

Research on high-hazard organizations has primarily focused on case studies of major accidents such as Three Mile Island (Perrow, 1984) or Challenger (Vaughn, 1996),

or on organizations that operate reliably under extreme conditions, such as the fast-paced actions on aircraft carrier flight decks (Roberts, 1990). Perrow's (1984) seminal sociological analysis attributes accidents to organizational and technological structure (complexity, coupling), and the institutions that surround organizations such as regulators, extra-organizational networks, and the interests of elites. The high-reliability theorists (e.g., Weick et al., 1999) focus on organizational resilience achieved through behavioral processes such as heedful interrelating, respect for expertise over position when danger arises, improvisation in the midst of action, and diversity of viewpoints to avoid simplification in the face of uncertainty. But neither body of work includes empirical studies of how high-hazard organizations create learning capabilities or develop the characteristics needed for high reliability operations. Our own research has focused on organizations that are attempting to change how they learn, and therefore how they perform.

### Our Research Program

Our approach begins with the assumption that organizational learning takes place through specific learning practices. Carroll (1998) and Popper and Lipshitz (1998) list some common practices or mechanisms such as after action review and benchmarking. Some learning practices are carried out by individuals, some by small groups, and others by large units or the organization as a whole. Crossan et al. (1999), for example, theorize that discovery and creativity are primarily the responsibility of individuals and small groups, whereas institutionalization of improved routines involves organizational actions. They conceptualize organizational learning as an integration of processes happening at different levels.

Rather than catalogue or describe learning practices, we focus our work on a particular learning activity carried out by nearly all high-hazard organizations: problem investigations within corrective action programs (Carroll, 1995, 1998; Carroll, et al., 2001; Schaaf, Lucas, & Hale, 1991). More frequent than the massive investigations triggered by rare accidents such as Three Mile Island, these smaller-scale self-analyses and problem solving efforts follow detection of small defects, near misses, and other lesser failures (Sitkin, 1992) or precursors (Reason, 1990). This problem investigation process is a form of off-line reflective practice (Argyris, 1996; Schön, 1987; Rudolph, Taylor, & Foldy, 2000): sensemaking, analysis, and imagining of alternatives take place outside of the regular work process, often carried out by individuals not immediately involved in the problem itself. Problem investigations are part of a corrective action program that starts with reporting of problems and continues with investigation of facts and opinions, attribution of causes, generation of insights and recommendations, implementation of interventions to improve performance, and verification that these interventions were carried out and produced the expected results.

Our research program has focused on studies of problem investigation teams at several nuclear power and chemical plants. Although individuals investigate most problems once they are reported, teams handle the most serious, persistent, causally ambiguous, and organizationally complex problems. Problem investigations illustrate how team learning, organizational learning, and individual learning are distinct yet interconnected (Crossan et al, 1999; Kim 1993). Teams are asked to imagine and interpret on behalf of the organization (Huber, 1991). The team report is a partial record of *team learning*, which feeds a change implementation process and explicit databases

intended to capture *organizational learning*. The team learning then must be embedded or institutionalized in knowledge reservoirs such as procedure manuals and databases (Argote & Ingram, 2000), physical changes, and altered routines (Levitt & March, 1988). Team learning influences *individual learning* as the team members develop their personal knowledge and skills and bridge their communities of practice (Brown & Cook, 1999). Individual learning indirectly influences organization learning as team members share information with managers and coworkers in the departments to which the team members return (Gruenfeld, et al. 2000).

In the next section of this paper, we provide four empirical illustrations of organizations seeking to learn how to improve their performance, using three different types of studies. The first type is represented by two case studies of specific problem investigation teams. One case study is based primarily on written questionnaires from team members and managers collected more than a year after the investigation, combined with our own analysis of the team's written report. The other case study is based on direct observation of the team during its investigation process and follow-up interviews with team members and managers several months later, along with analysis of the written report. The second type of study is a quantitative analysis of questionnaire responses from team members and managers involved in 27 investigations along with coding of the written reports. The final case study examines an organizational transformation.

Although this diverges from our studies of problem investigations, it represents the same themes of control and learning at an organizational level. The first author visited this plant regularly as part of a team advising the board of directors, combining first-hand observation, document review, and interviews with a wide range of employees and other

key informants over a four-year period of time. Analysis of the four empirical examples helps us develop a *four-stage model* of organizational learning.

### **LEARNING IN ACTION: EMPIRICAL ILLUSTRATIONS**

In the first case study, a nuclear power plant investigated an incident in which an employee was seriously hurt. The plant was attempting to improve safety and performance in part by using a newly upgraded problem investigation process. In the nuclear power industry, regulators and industry groups have long been calling for greater awareness of minor problems and actions to avoid future trouble (Jackson, 1996; Rochlin, 1993). As Weick and Sutcliffe (2001) state, “to move toward high reliability is to enlarge what people monitor, expect, and fear” (p. 91). A typical nuclear power plant identifies over 2000 problems or incidents per year, 90% of which would have been ignored a decade ago. The investigation created an opportunity to raise collective awareness about local work practices and helped managers strengthen controls and increase conformity to industrial safety rules.

#### 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 several broken bones and a lacerated lung and arm. His coworkers used a nearby phone to call for help. Emergency medical technicians arrived shortly and took Joe to the hospital.

### *The Team Investigation and Managers' Reactions*

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 proper 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 heeded the implications when Joe was advised to stay on the one and a half-inch steel framework of the roof 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.

The team's analyses and recommendations had to be negotiated with and implemented by managers. Preliminary reports are typically discussed with management before being formally issued, in order to increase clarity, comprehensiveness, and buy-in. A team member suggested that management had been defensive about initial drafts and the report had "pulled its punches." 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.” When asked if the team’s recommendations had been implemented successfully, a team member responded, “Management took renewed emphasis on safety. Procedures (pre-job briefs) were changed and working aloft programs were implemented.” But a different team member replied plaintively, “If top level managers aren’t willing to listen to the people doing the work, and respond to their findings, it all becomes a waste!”

*Discussion: The Compelling Nature of Control*

This problem investigation illustrates the plant’s effort to strengthen its controlling orientation. The report highlighted the failure of workers and first line supervisor to comply with existing rules and procedures. The corrective actions were aimed at increasing awareness and compliance with these rules by appointing a safety advocate, reinforcing 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 was 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 et al., 1985; see Fig. 1).

\*\*\* Insert Fig. 1 \*\*\*

In striving for control, identifying a presumed “root cause”<sup>3</sup> or source of the mismatch reduces uncomfortable ambiguity (cf. Frisch & Baron, 1988). Causes are

---

<sup>3</sup> Industry uses many definitions of root cause such as “the primary or direct cause(s) that, if corrected, will prevent recurrence of performance problems, undesirable trends, or specific incident(s).” In practice, it means a manipulable cause, i.e., a place to intervene to improve performance (Carroll, 1998).

found that are familiar (Rasmussen & Batstone, 1991), proximal to the problem (White, 1988), with available solutions that can easily be enacted, and are acceptable to powerful stakeholders (Carroll, 1995; Tetlock, 1983). Observers commonly make the fundamental attribution error of finding fault with salient individuals in a complex situation (Nisbett & Ross, 1980), such as the operators or mechanics who had their hands on the equipment when the problem arose (Carroll, 1995; Reason, 1990). Decision makers see problems “out there” rather than questioning how their own internal assumptions and beliefs contribute to the structures in which they work (Repenning & Sterman, 2000; Torbert, 1972).

In a controlling-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, 1999), a frequently ineffective process in which representatives suggest solutions to managers who may resist implementation (Carroll et al., in press) for various reasons. Managers avoid expensive actions, or actions without a clear contribution to specific problems in their domain. As one member of the investigation team commented, “When it was becoming apparent what the real problem was [i.e., in management’s domain], I think the group became (temporarily) unsure where to go—what to do—it looked like a big step.”

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 workers 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 (Adler, 1991). The investigators did not ask double-loop learning questions such as, “What frames do supervisors and workers hold that let a casual approach to safety develop and endure?”; “How does the status and career advancement system contribute to a culture of risk taking?”; “What assumptions and practices allowed a design problem (lights in an unsafe place) to exist for so long?”; or “How does the work system of separated functions and hierarchical authority inhibit mutual understanding?”

#### Quantitative Analysis of Problem Investigation Teams

The Fall From Roof investigation was one of 27 investigations we studied at 3 nuclear power plants (Carroll, Hatakenaka, & Rudolph, 2001). Although two plants were selected for study because of their prior contact with the first author, the third was added because industry experts generally considered it a leader at conducting problem investigations. These investigations addressed the most serious problems and most troubling trends during the previous two years, and were therefore assigned to temporary teams of three or more members from multiple departments who were released from their regular work for the duration of the investigation.

After preliminary interviews at each plant to understand their investigation process and to identify team investigations, lengthy questionnaires were distributed to

team members and a shorter version went to managers who identified themselves as sponsors or customers of the reports. The team member questionnaires contained open- and closed-ended items about the investigation, the team process, the report, and any resultant changes in the plant. There were also measures of cognitive complexity, cognitive style, age, education, work experience, and so forth. The manager questionnaires focused on the difficulty of the investigation task, the quality of the report, its usefulness for creating improvements, and their reasoning about why improvements had or had not occurred. Two members of our research team coded the team reports for various characteristics loosely grouped into causes (e.g., complete causal chain, role of culture), barriers/defenses (i.e., gaps in procedures, supervision, training), learning (e.g., failures to learn from previous problems), corrective actions (e.g., logically connected to causes), and narrative features (e.g., memorable story, active voice). In this section of the paper, we concentrate on a subset of variables directly related to team learning (the team report), organizational learning (changes to the plant), and characteristics of the team (team member demographics and work history).

### *Team Learning*

We assumed initially that team learning would be captured in the official team product, the written report. After all, the report is the artifact that drives the corrective action process and is encoded into work management and problem trending databases. However, we came to understand that the team report is negotiated with managers and therefore represents a subset of that learning. We will first discuss the contents of the team reports and the characteristics of the teams and the support they received, and then consider other aspects of what the teams learned.

Our coding of the reports showed a disappointing level of depth and completeness, insight and clarity. The focus of the reports was most frequently on changes to action that involved improving compliance with existing procedures. Consistent with the example of Fall From Roof, teams rarely looked for fundamental or deep, systemic causes. They sometimes did a good job of understanding problems in terms of failed barriers (e.g., design, procedures, supervision) that had actionable responses. Corrective actions could be misaligned with the supposed causes, either leaving causes without corrective actions, or introducing corrective actions without specifying how they would address the causes. Few reports were well written -- they were sometimes confusing, often redundant, lacking a good story line, and usually in passive voice. There was little evidence of inquiry or proposed actions that could challenge assumptions and change deep structures (Argyris & Schön, 1996). For example, while some reports identified assumptions and beliefs about work practices that contributed to the problem, the proposed corrective actions almost never addressed these assumptions and beliefs.

Figure 2 presents statistically significant relationships from a series of stepwise multiple regression analyses controlling for plant. We assumed that characteristics of teams<sup>4</sup>, shown in the left-most column of variables, would influence our coding of report characteristics, given in the next column. Controlling for plant, reports with corrective actions addressed to fundamental causes and “out-of-the-box thinking” (a composite of four codes that reflected an ability to see a problem in terms of past similar problems and potential future problems, cross-context comparisons to other departments and plants, and

---

<sup>4</sup> Responses from team members were averaged by team to create team-level variables, except for departmental affiliations, where we counted unique departments represented.

attention to informal routines) were associated with team members who had more experience in the nuclear power industry, fewer years in the military, and more varied plant experience (total number of departments team members had worked for in prior jobs). Depth of causal analysis (a composite of five codes about the completeness and integration of causal explanation, attention to culture, management, and oversight) was not related to any of the predictor variables (and is not shown in Fig. 2). Reports with more generic lessons learned came from teams with more average years of plant experience, those with stronger narrative qualities came from teams with more training in teamwork, and reports with identification of more failed barriers came from teams with more training in teamwork and more varied plant experience.

\*\*\* Insert Figure 2 \*\*\*

Although there were only three plants in our study, one of them was generally considered to be a national leader in problem investigations. The seven reports we collected from that plant did not exhibit more depth of causal analysis or identification of more generic lessons learned. However, reports did have statistically significant higher scores on corrective actions matching causes, use of an interdisciplinary approach, and “out-of-the-box thinking.” Reports from this plant had better grammatical structure and made better use of graphics. This plant also had significantly larger teams, more departments represented on each team, more managers or supervisors on the teams, and team members with the least years of military service. Plant respondents also had significantly lower scores on Need for Closure, i.e., relative to those from other plants, their cognitive style was more comfortable with ambiguity and novel problems rather than structure and clear answers (Webster & Kruglanski, 1994). Interviews at the plant

consistently discussed how the plant manager strongly supported self-analysis and learning, including requesting involvement of other plants in investigations. This suggests that any advantage of this plant lay not in their analytical and causal models, but rather in their open and cooperative attitude toward learning, their linking of problems across contexts, and their ability to turn analysis into action. These capabilities helped institutionalize a stronger learning orientation.

### *From Team Learning to Organizational Learning*

In nuclear power plants, investigation teams do not have authority to implement changes. The teams report to managers who “sign off” on the teams’ recommendations, and then the teams disband. Managers oversee any actual changes to routines and physical equipment. Although we lack direct observations of such changes, we asked both team members and managers to report on changes and improvements. Fig. 2 also shows how ratings of changes at the plant by team members and managers were associated with team characteristics, report characteristics, and ratings of report quality by team members and managers.

In general, both team members and managers reported that changes had been made that addressed the problems, based on the team recommendations. However, correlations between team and manager ratings were modest and not significant ( $r = .33$ ). Many team members reported not knowing what had happened as a result of their report, perhaps because there was infrequent feedback to ex-team members after the teams disbanded. Some team members reported that management had been defensive and therefore reports had been less than candid. Teams that reported more change resulting from their reports also indicated that their teams had better relations with management,

more frequently had managers or supervisors on the teams, and had more training for teamwork.

From their side of the negotiation, managers complained about long lists of causes and corrective actions that undermined the impact of the report and seemed to yield little value for the investment. Managers rated change as greater when the team members had more investigation experience. In general, the results reinforce the importance of both skills in teamwork and investigation and the external function of teams -- their ability to negotiate resources and goals with management, to obtain information from other groups, and later to “sell” the report (Ancona & Caldwell, 1992; Dougherty & Hardy, 1996; Dutton & Ashford, 1993) from a weak power position analogous to minority influence (Wood et al. 1994).

The challenges of communicating between team members and managers were clarified by examining their differing reactions to the team reports. On average, team members and managers thought the teams had “produced excellent reports” (means of 1.80 and 2.12 with 1=strongly agree and 6=strongly disagree), however, these judgments correlated only slightly with each other ( $r=.18$ ,  $n=27$ , n.s.). As shown in Fig. 2, teams’ self-ratings of the excellence of their reports were associated with reports that identified failed barriers and generic lessons that could be applied across contexts, self-ratings of better access to information, and more departments represented on the team. Managers’ ratings of the excellence of the team reports were associated with better access to information, more team investigation experience, and reports that suggested effective corrective actions. Given such differences, particularly the teams’ focus on causes and managers’ focus on corrective actions, and the open-ended comments about difficulties

communicating between teams and managers, we suspect that much of team learning remained tacit and unwritten (and possibly was lost) due to time pressure, management embarrassment, or inability to articulate assumptions.

One outstanding counterexample was a report from the well-regarded plant that offered an explicit description of the negotiation process between the team and its management customers. In its initial draft report, the team made over twenty recommendations, which management evaluated and reduced to six, only four of which were implemented. The final report provided a uniquely candid discussion of managers' cost-benefit analyses of all the original recommendations. Reports from this plant were generally more candid than the other plants, perhaps because the plant manager was more supportive of learning. In contrast, at one chemical company we visited, investigation reports were considered a legal requirement to be met in a minimal way that would not expose the plant to additional scrutiny; the important learning was conveyed through oral reports and discussion. Unwritten knowledge was referenced in reports and databases by the names of knowledgeable individuals (cf. transactive memory, Moreland & Mayaskovsky, 2000).

#### *Discussion: Control and Learning*

Given these differences between team members and managers, it is not surprising that plants have difficulty learning from investigations of operating experience. There are different "knowledge reservoirs" (Argote & Ingram, 2000) with different content and different interpretations arising from different mental models or logics (Carroll, 1998). Plants therefore have difficulty putting knowledge into action or implementing anything that managers cannot support. In control-oriented plants, only a limited number of

knowledge reservoirs have sufficient legitimacy to influence management; in more open plants, more knowledge reservoirs are created and used. Interestingly, at a chemical company that we will discuss later, problem investigation activities had the explicit goal of *educating managers*, not solving problems! In this company, teams presented facts and tightly reasoned causal connections, but did not make recommendations. It was managers' collective job to understand the problem and its context, discuss improvement opportunities, commission solution development activities, and implement changes.

Organizations that encourage participation from a wider range of members are more likely to surface underlying assumptions and mental models and to address more fundamental changes in structure and culture (Bartunek, 1984). However, the relationship between diversity and organizational performance is complex and contested (Williams & O'Reilly, 1998). Interestingly, teams from our "industry leader" plant rated diversity of team inputs as significantly more important than did teams from the other plants. Informational diversity (Jehn, et al., 1999) or "conceptual slack" (Schulman, 1993) can arise from mixing occupational and educational backgrounds (Dougherty, 1992; Rochlin & von Meier, 1994) and varied cognitive styles (Jackson, 1992; White, 1984) that together combine abstract, systemic issues with concrete, operational details and technical complexity with human ambiguity. Weick et al. (1999) similarly assert that high reliability requires a reluctance to simplify, which is supported by having a rich and diverse set of perspectives and frequent boundary-spanning activities. This creates skills and opportunities for engaging in a process of knowing (Cook & Brown, 1999) that can help surface previously unarticulated mental models of the work environment, compare them, and arrive at new, shared views.

Problem investigations are precisely the kind of opportunity that could bring diverse perspectives together and facilitate learning and change. Yet this was only partially successful. Managers seeking efficiency delegated participation (Nutt, 1999) to the team and waited to respond to a draft report. They sought logical and practical solutions to problems, and disparaged complex analyses as attempts to “solve world hunger.” Deeper inquiry that could reveal new insights and systemic understandings would have required time to grapple with uncertainty and complexity and therefore that managers permit themselves temporarily to “not know” and “not act” (Schulman, 1993; Weick et al., 1999). Directly involving managers in the investigation activity could be an important strategy, but it is resisted by overworked managers who are encouraged to produce, not learn -- exploit, not explore. As we mentioned earlier, the industry-leading nuclear power plant made it easier for managers or supervisors to be team members.

The team report was only a weak boundary object (Carlile, in press; Star, 1989). Boundary objects are concrete means for team members and managers to specify and learn about their differences and dependencies through joint practice that collectively transforms knowledge. However, sharpening and bridging differences between managers and teams requires rigorous attention to a shared approach. As we noted, managers were often not full participants on the teams, reports were rather casual in connecting causes and recommended actions, and negotiations over the final report tended to be about authority rather than reasoning. It takes mindful attention to build shared understanding around diffuse issues such as “culture” and “accountability” that have very different meanings and implications to professional groups (Carroll, 1998; Carroll *et al.*, in press). Because of their emphasis on controlling deviations from existing procedures and rules,

few teams and managers were willing or able to work hard to clarify meaning and build shared mental models. They therefore missed opportunities to deepen their understanding and produce organizational learning and change.

### Charge Heater Fire

At a petrochemical plant, the new plant manager initiated a plant-wide effort to use problem investigation teams as a way to address, simultaneously, a recent history of financial losses, some dangerous incidents, and repeated equipment failures. The plant ownership had recently changed to a joint venture among chemical companies, one of which used problem investigations as a regular practice at some (but not all) of its plants. Two headquarters staff from this company had been working for a decade from positions of little formal authority to promote more strategic and systemic thinking at operational and executive levels in their company, using problem investigations (which they called “root cause analyses”) as one of several approaches. Only in the past two years had they received significant management support and institutionalization of new practices at several plants.

They contracted with the plant manager to support introduction of their root cause analysis practice to his plant, beginning with a large-scale learning intervention.<sup>5</sup> Managers decided to expose about 20 plant employees, operators, maintenance staff, engineers, and first line supervisors to root cause analyses by dividing them into teams to explore four significant recent problems.

Each problem investigation team included some members from inside and some from outside the plant and at least one experienced root cause facilitator (most from other

---

<sup>5</sup> They avoid calling this “training,” which implies top-down control, and misses the larger desired impacts on managers and the culture in general.

plants). The overall process included presentations and practice sessions on investigation techniques, analytical tools, and reporting methods during the course of a three-week time frame, culminating in reports to plant management. The facilitation team timed each day's preparation 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.

One of the investigation teams examined a 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 the causes of the fire, plant managers also wanted to discover and ameliorate the conditions that 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, sometimes detailed to the minute;
- an “Is/Is not” process that differentiates circumstances where the event occurred from similar circumstances where it did not, e.g., what differed between this charge heater and a similar one that did not catch fire, or between this moment in time and earlier times when the same charge heater ran uneventfully (Kepner & Tregoe, 1981);

- a detailed causal event diagram working backward in small logical steps from the fire to the causes and conditions that were collectively necessary to produce the incident, and working back further and further to uncover deeper causes; and
- a process of categorizing the quality of data used to support each link in the causal event diagram (as 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.

### *The Team Investigation*

Distilling and analyzing the information available, the team concluded that the fire was due to a tube rupture inside the charge heater that occurred when the three quarter inch steel skin of the tube got too hot and tore. The team found that three factors contributed to the heater fire: (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 and thereby achieve the desired higher 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 maximum 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 to increase production 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 first 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 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 major causal factor was a necessary but not sufficient contributor to the incident. In one of its verbal reports to other investigation teams during discussion session, 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 shown results:

The biggest issue that came out [of the root cause analysis intervention] 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 you do not do it, you 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.”

*Discussion: The Challenges of Organizational Learning*

The charge heater investigation provides examples of an organization striving to motivate openness and learning and also to enhance inquiry and analysis skills. In our observations of the intervention, it was evident that at least some participants were anxious about being open with colleagues in their own department, 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 punitive response or a controlling 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. We observed the plant manager and the facilitation team repeatedly articulate the value of rethinking and deeper learning for achieving better performance. During the course of the investigation, employees discovered they could tell managers about how resource allocation decisions were creating unanticipated safety and quality issues, and that managers would listen, respond to questions, and focus on making things better. Operators told engineers about technical problems and the engineers started to listen. There was a willingness to confront reality and to surface underlying assumptions

about “how we do work around here.” Support from the new plant manager helped encourage full participation from the teams and from managers. That support was itself an outcome of the facilitation team who worked publicly with the investigation teams but met *privately* 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. Yet, employees’ final verdict on management openness would await management action following the investigations.

Team members developed and practiced new learning capabilities as they started addressing operations at the plant from a systemic perspective and challenging assumptions. The rigor of the root cause analysis process encouraged team members to “hold their assumptions lightly” as the analysis held these views up to comparison and disconfirmation. They enacted double-loop learning by recognizing assumptions and mental models as distinct from reality, understanding that assumptions and mental models affect behaviors and outcomes, imagining alternative mental models to guide action toward more desirable outcomes, and taking action with the new mental models (Friedman & Lipshitz, 1992; Argyris, et al., 1985). 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 (Deming, 1986) that working to optimize individual components does not automatically add up to an optimized system.

The cause-effect diagrams worked as boundary objects to help reveal tacit assumptions about plant processes that were key links in the causal chains leading to the heater fire. As team members created these diagrams, they were continually confronted with the fact that each saw the same situation differently, forcing them to recognize that their image of reality was distinct from others' images and from the reality itself. 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. They discovered how mental models affect behavior when they recognized that three assumptions in their cause-effect tree had contributed to the charge heater fire: "[desired] charge rate [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, et al., 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."

An explicit goal of the investigation and the overall intervention at the plant was to educate management by challenging their mental models with rich and compelling data and interpretations. Because written reports were considered to pose a potential legal liability, a detailed oral reporting process encouraged management involvement and discussion. The use of cause and effect diagrams and tight factual logic appealed to norms of rationality shared by managers, engineers, and participants in general. Thus,

new practices were linked to legitimate aspects of the culture (Schein, 1992). The facilitation team reduced resistance by coaching managers to spend time with the teams and to advocate learning even when their own beliefs were being challenged.

### The Millstone Turnaround

Moving beyond problem investigations, the final case examines an organization in crisis and recovery. When managers pushed a classic controlling orientation to the extreme, employees and external regulators lost trust and the organization's learning capabilities eroded. External pressure mandated culture change, but enacting change required new learning capabilities. The changes emerged as a collective improvisation. We examine the organizational conditions that form the context for problem investigations, reflective self-analyses, and learning. We again see themes of control and learning, this time with a different focus.

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 three 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." A 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. 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.”

Millstone Station was required to establish programs and institutions to define, develop, and maintain a safety conscious work environment, yet the NRC itself lacked any regulations or experience on the topic. Millstone’s response was closely watched by the industry as the utilities and regulator continued to learn together (sometimes through painful experimentation) how to expand their focus from engineered safety features and compliance with procedures to the management of a high-hazard work system.

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 psychologically safe (Edmondson, 1999) 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 compliance to openness and trust, from efficiency to performance, and from denial of problems to preoccupation with failure (Weick et al., 1999). It took over two years to shift the culture and learning orientation 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).

## *The Culture Change*

In September 1996, the new CEO for Nuclear Power, Bruce Kenyon, set the scene for change with 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 quickly revamped the top management team and introduced a stronger employee concerns program.

His subsequent actions were a role model for 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 [employees] to get candid answers.” Upon hearing Kenyon say publicly at his first NRC meeting that he found the organization “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 was not prepared, however, to deal with the safety conscious work environment issue and the sensitivity of Millstone employees. In his previous experience, management had dealt with employee issues in an atmosphere of trust, without needing special confidential programs or the layers of oversight required by the NRC Order. He assumed that by articulating and living his values, and strengthening the

Employee Concerns Program, he could rebuild trust. He had “never encountered a culture as broken.” At Millstone, management actions were regarded with suspicion and distrust. For example, when Kenyon relieved the VP of Oversight (responsible for audits and quality control), some employees immediately interpreted it as another attack on independent Oversight.

But Kenyon allowed himself to be fallible (Schulman, 1993) and he enlisted participation from others so that the change initiative could emerge as a collective product. When two contractors were terminated for alleged poor performance and the Director of the Employee Concerns Program immediately protested and 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. In response, 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 experienced personal transformations in shifting their assumptions and mental models. The operations vice president had been a typical old-style manager: he described himself as weary of “whiners,” and he “didn’t believe anyone would harass someone who brought forth safety concerns.” When the two contractors were terminated and the Director of the Employee Concerns Program offered his 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 attention to safety concerns could make business sense.

The nuclear power industry, including Millstone, is staffed with managers who are “not high on people skills, for example, few can read nonverbal signals.” Managers had strong assumptions that employees shared their perception of reality; they had to learn 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. Since management had no intention of being retaliatory, they 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 (cf. emotional intelligence, Goleman, 1995). Through extensive new training programs and coaching by organizational development consultants, they began to “learn the difference between anger, hurt, and a chilling effect”.<sup>6</sup> Managers needed to respond differently to employees who were afraid of reprisals for speaking up and those who simply lacked confidence that management would take effective action. Controlling-oriented managers, some of whom get their way by yelling and threatening, generally deny their own emotionality and claim to value facts and rationality, even when they are using fear to exercise control. Openness to the

---

<sup>6</sup> A chilling effect is a perception that employee concerns will be met with retaliation, thereby suppressing them.

thoughts and feelings of others creates feelings of dependency and vulnerability (Hirschhorn, 1993). Dozens of managers who could not change were moved out of management or left the company, which sent a clear message that managers were accountable for relationships as well as performance.

Openness and trust emerged organically through multiple mechanisms and venues that encouraged participation (employee voice, Adler & Borys, 1996), diversity of viewpoints, and attention to solving and averting problems. We have already mentioned ECP and ECOP that encouraged reporting and gave employees more voice. Management created an Executive Review Board after the contractor terminations to review all disciplinary actions. Offering an opportunity for discussion among senior managers and an ECOP observer, it opened up the management process, 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, 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.

*Discussion: Improvisation and Emergent Change*

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 regime of centralized authority and mutual distrust to a culture of open communication, trust, and participation.

Millstone managers had been proud of Millstone's technical leadership of the industry and excellent performance record. Managers believed that Millstone's design features and managerial controls were sufficient to operate the plant safely and reliably. Management's basic assumption was, "we know everything we need to know." 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. As a result, employees developed a basic assumption that "management can't be trusted." In short, managers and employees (and regulators and publics) 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. Although employees were suspicious of management and biased toward self-confirming perceptions, they began to change their assumptions because senior management reacted well to critical events such as the contractor terminations. Independent voices were allowed to challenge underlying assumptions as multiple venues emerged for managers and employees to talk together and work on the common problem of rebuilding Millstone (and retaining jobs!). 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. Trust emerged through collective action

and perception of shared purpose (Kramer, 1999; Whitener, Brodt, Korsgaard, & Werner, 1998).

### A STAGE MODEL OF ORGANIZATIONAL LEARNING

Our studies of the nuclear power and chemical industries highlight some frequent challenges of control and learning that resonate with the literature on high-hazard and high-reliability organizations, notably: (1) integrating local learning with organizational learning (Crossan et al., 1999), which depends upon trustful relating (Edmondson, 1999) between groups and across hierarchical levels and a balance of structure and improvisation (Weick & Westley, 1996), and (2) surfacing assumptions and mental models (Argyris & Schon, 1996) in order to act mindfully (Weick et al., 1999) upon operational details in the context of a tightly-coupled system (Perrow, 1984; Repenning & Sterman, 200).

We introduce a stage model as a way of capturing our observations that these organizations seemed to exhibit a common order or progression in the initiatives they have taken to address these challenges and build learning capabilities. Stage models are widely used in organization theory, sometimes based on biological metaphors of life-cycle development or increasing maturity (e.g., Rooke & Torbert, 1998; Quinn & Cameron, 1983), and sometimes on historical analyses (e.g., Chandler, 1962; Malone & Smith, 1988; Perrow, 1970). Bohn (1994) provided a stage model of increasing process control knowledge. Reason (1997, pp. 61-65) offered a stage model of organizational control moving from predominantly feedforward embedding of anticipated problems into procedures to predominantly feedback learning from experience.

However, we are not proposing that organizations necessarily go through these stages as ecological adaptations or that these stages are an invariant unfolding of organizational maturity. Although environmental factors such as regulation, public opinion, and managerial ideology can influence organizations (e.g., Guillen, 1994; Hannon & Freeman, 1977; Powell & DiMaggio, 1991), as shown in the Millstone example, there is no particular direction to these forces. Further, organizations may ignore signals (Cole, 1998; Freeman, 1999) or define situations as either threats to be defended against or opportunities to lead (e.g., Dutton, 1993). The changes we observed in our studies were a product of creativity and improvisation, in the face of uncertainties and risks, with both beneficial and problematic outcomes.

Thus, the four stages in Figure 3 are presented as a provocative guide to analysis, not as a rigid model of development. “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). 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, since the latter stages require shared understanding and collaborative effort, sustaining these capabilities requires their relatively widespread use. Although we propose that these stages and capabilities tend to emerge in a particular order, being “at” a stage means that there is relatively more behavior consistent with that stage and *earlier* stages. Indeed, organizations that enact multiple learning orientations and processes at many organizational levels (individual, team, department, and so forth) may generate a useful creative tension (Crosson & Hurst, 2001; Crosson et al., 1999;

Weick et al., 1999). Our studies show several examples of such partial transitions and continuing inconsistencies between local knowledge, improvisation, formalized rules, hierarchical power, and systematic inquiry.

\*\*\* Insert Figure 3 \*\*\*

### Local Learning 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). Despite the focus of research on large corporations, the vast majority of firms employ under 100 workers (Aldrich, 1999). Early nuclear power plants started as small demonstrations of possible designs. Although they were more proceduralized than fossil fuel plants, nuclear power plants were similar in drawing knowledge from the experiences and skills of craft workers and sharing the attitude of “run it ‘til it breaks.” During the next two decades, reactor designs rapidly increased in size to achieve greater economies of scale, increased challenges to their inherent safety because of the sheer amount of radioactive fuel in one place, and increased their complexity as more safety systems and people were added to compensate.

In what we label the Local Stage, there is considerable organization-specific and task-specific knowledge that is 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). The organization is minimally self-aware: decisions are made locally by those steeped in the details, and learning mostly occurs locally as well. The challenge at the local stage is to move from habitual or idiosyncratic

actions with no comparisons to standards to single-loop learning, i.e., behaviors are adjusted after comparison to performance standards or benchmark models, but underlying structures and assumptions are not challenged (see Figure 1). Although academics often privilege double-loop learning, single-loop learning is an accomplishment, undoubtedly more common, and often very powerful (Argyris, 1996, p. 80; Miner & Mezias, 1996).

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. Although the experience of failure can prompt learning, what is detected and defined as a failure, and who is blamed for the failure, depends on how information is shared and interpreted (Carroll, 1998; Repenning & Sterman, 2000). 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 local information into plant-

wide or industry-wide 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)

The Fall From Roof and Charge Heater Fire cases offer examples of local work practices among maintenance workers, operators, and engineers that were learned over time because they appeared effective and efficient to the work groups involved, but were also detrimental to the plant. In the case of Fall From Roof, industrial safety was being compromised in a way that primarily affected the workers themselves, but secondarily reinforced a generally casual attitude toward rules that could signal a broader risk. A major lesson from the Charge Heater Fire investigation was that local groups were making changes in routines that appeared more effective from their viewpoint, but they never checked with other groups about the impact of those changes on the plant as a whole system. Their local practices only came to light when implicated in a serious event, and even then only when investigators were able to gather rich information and analyze it with a highly-disciplined set of conceptual tools.

#### Constrained Learning Stage

The constrained stage refers to how learning is enacted within a strong and coherent set of assumptions about division of labor, formalization of tasks, and linear cause and effect. Learning is “constrained” by assumptions, mental models, habitual

routines, and entrenched interests seeking to preserve their status and perceived competence (cf. constrained action, Crosson & Hurst, 2001; Pfeffer, 1982).

In the nuclear power industry, for example, the decade after Three Mile Island was characterized by dramatic increases in regulations, formal procedures, internal and external oversight, reporting requirements, and staffing.<sup>7</sup> The industry developed sophisticated probabilistic techniques for anticipating problems (US Nuclear Regulatory Commission, 1975; Wildavsky, 1988) and designed “defense in depth” such that a serious accident could occur only if a series of barriers were breached (Reason, 1997). Industry organizations (e.g., Institute of Nuclear Power Operators, see Rees, 1993) were founded to promote training and exchanges of information and best practices. The industry recognized that unanticipated interactions among complex processes were possible (Perrow, 1984), but believed that proper design and procedures could reduce the likelihood of accidents to acceptable levels.

In our studies of nuclear power industry problem investigation teams, most were working to establish and maintain control over local work processes and to prevent future problems. As we noted, their analyses rarely went very deep into underlying processes and assumptions, but rather executed single-loop learning to efficiently fix problems. Similarly, when best practices are transferred without rethinking assumptions about why they work, the best practices are being used for fixing or exploiting what is known rather than challenging mental models and exploring new possibilities.

Control through measurement, monitoring, incentives, and other traditional bureaucratic mechanisms seems to come naturally to managers and engineers (Carroll, 1998; Schein, 1996). Supported by cognitive biases such as the fundamental attribution

error (Nisbett & Ross, 1980) and the illusion of control (Langer, 1975), this orientation may be universal, or at least deeply embedded in Western industrialized culture. In the stable and protected environment of a regulated industry with a relatively fixed technology base and public pressure for reliable and safe performance, a control orientation and skilled use of single-loop learning may be appropriate (cf. Miner & Mezias, 1996; Reason, 1997).

However, the constrained learning stage can become a competency trap (Levitt & March, 1988) buttressed by self-confirming attributions (Repenning & Sterman, 2000). Working harder to achieve consistency and efficiency may preclude “working smarter” for fundamental process improvement because such initiatives divert resources away from current production (Repenning & Sterman, 2000). The Millstone case illustrates a controlling orientation pushed so far that information contrary to strongly held assumptions was resisted by authority, and learning was inhibited. Management reacted to internal and external criticism by accusing the critics of error or ulterior motives, which led to more insularity and more criticism. When organizations are forced by repeated failure or stakeholder pressure to initiate changes, they may seek to copy success stories (e.g., start a TQM program, Cole, 2001) and achieve legitimacy rather than to rethink assumptions and reshape the organization. As the nuclear power industry deregulates and downsizes, the challenge is to avoid increased constraints on learning or mindless adoption of management fads. As we pointed out in the Charge Heater Fire case, it was the concepts underlying the tools, not simply the tools themselves, that enabled movement beyond the constrained stage.

---

<sup>7</sup> As suggested by Barry Staw, the Enron debacle may have a similar impact on many corporations.

## Open Learning Stage

Large, conservative, bureaucratic organizations can be highly successful in stable environments, but as we observed in the preceding section, in turbulent and unpredictable environments, they may be stuck in a competency trap that inhibits learning. Facing pressures from new competitors with new products, rapidly changing technologies and customer preferences, deregulation, and so forth, large organizations may initially ignore these threats (Cole, 1998; Freeman, 1999). Even when repeated internal problems or external threats garner attention, there may be little learning if powerful stakeholders feel threatened or existing norms prevent change (Vaughn, 1996). Bureaucracies have features “which are designed to repress or forget confusing or contradictory qualities” (Weick & Westley, 1996, p. 445). Even in trying to respond to problems, constrained learning practices can be overwhelmed by the sheer number of exceptions, the need to rewrite and retrain routines, and the work necessary to assimilate learning from various sources (Rudolph & Reppenning, 2002). For example, Perron & Friedlander (1996) suggest that management systems for Process Safety Management in the chemical industry “cannot yet be fully automated” (but notice the “yet”).

Some of the organizations we studied were trying to move out of the constrained stage, having recognized some limitations to top-down control. They were promoting more participation and open exchange of information throughout the organization and between the organization and the outside world. For example, interviewees characterized the industry-leading nuclear power plant in our questionnaire study as having an open learning environment supported by the plant manager (in contrast to sister plants that were more constrained). The nuclear power industry in general had been advocating a “questioning attitude” and “safety culture” directed at acknowledging doubt (Schulman,

1993), increasing awareness or mindfulness (Weick et al., 1999), respecting the contributions of others in an atmosphere of trust (Edmondson, 1999), and placing a positive value on teamwork and learning.

But openness cannot be mandated; it must be enacted. Weick et al. (1999) emphasize the importance of creating opportunities for diverse viewpoints to engage in conversations; for “moments of juxtaposition” (Weick & Westley, 1996) when participants can forget what each knows and learn through social interaction. Trust was built at Millstone when courageous pioneers (including whistleblowers) took personal risks and managers reacted favorably to new problems, which allowed openness to spread in a virtuous cycle.

The open stage at Millstone was also characterized by an awareness of people as different from machines. The ability to acknowledge emotions, conflicts, and different perceptions that underlie work relationships and political contests allowed for discussion of the human side of organization. This challenges deep assumptions about human nature and rationality (cf. Schein, 1992). Of course, many Millstone managers were uncomfortable and initially incompetent in this domain (Hirshhorn, 1993), but openness to its importance allowed a variety of mechanisms to emerge that institutionalized a new set of assumptions, new organizational practices, and a higher level of people skills. Over time, managers and employees learned by doing and through feedback from colleagues and coaches.

However, it is difficult to sustain an organization at the open stage unless people develop significant skills at producing usable knowledge that demonstrably improves performance and maintains legitimacy for learning. Individuals and groups may be able

to maintain openness as a desired value, especially with the support of senior management. But power differentials and the “entrenchment of the coercive logic suggest that the space for the emergence of enabling formalization is quite limited” (Adler & Borys, 1996, p. 82). Without demonstrable results, the open stage is vulnerable to forces that resist change (Schein, 1992) and push the organization to return to the familiarity and predictability of the constrained stage. As Hackman and Wageman (1995) note, “in too many TQM programs, moreover, it is the difficult-to-implement portions of the program that are being finessed or ignored and the rhetoric that is being retained” (p. 338). A problem investigation can be an opportunity to confront assumptions and juxtapose differences, or a ritual that justifies blaming the usual suspects and making familiar but ineffective responses. At Millstone, competitive pressure for efficiency or other stresses may encourage a return to constrained, centralized learning practices (cf. Staw, Sandelands, & Dutton, 1981).

#### Deep Learning Stage

Organizational learning capabilities reach a fourth stage when openness to learning becomes linked to disciplines for learning. 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. Managers assume that it is too costly to question and analyze everything, but developing stronger learning capabilities can reduce the costs of learning. We have some evidence of organizations struggling toward this deep learning stage in the Charge Heater case and the earlier Du Pont story.

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), as we saw in examples of the constrained stage, rather than to challenge deep assumptions with rigorous and systemic thinking. The same tools can often be used for controlling *people* and limiting their learning rather than controlling *processes* to provide more opportunities for learning (Adler & Borys, 1996; Bohn, 1994; Sitkin et al., 1994). The Charge Heater Fire investigation illustrates that actions and assumptions must be rethought 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. Applied rigorously, analytical approaches helped equalize status across disciplines and hierarchical levels through common reliance on shared reasoning, thus creating a shift from authority to expertise and involvement of diverse viewpoints (Weick et al., 1999). The tools and the learning activities became an opportunity to have new conversations, enact new behaviors, develop new skills, and build new relationships that also changed the culture.

Deep learning capabilities enable organizations to develop and share systemic mental models that represent the interdependencies or coupling of specific components and practices. Systems are difficult to understand, particularly when they have many parts, many linkages among the parts, and delays between causes and effects (Repenning & Sterman, 2000; Senge & Sterman, 1991). Consider the analogy of newcomers to rowing who typically maintain a rigid grip on their rowing oar because it gives them the *feeling* of control. However, a rigid grip decreases absorption of the shock of uneven

waters, thereby decreasing *actual* control. Managers often assume that “heavy-handed” incentives and authoritative micromanaging is the way to assure control. However, this may drive noncompliance out of sight, thus immediately increasing their feeling of control, but simultaneously increasing hidden noncompliance (cf., Reason, 1997) and making the system more difficult to manage in the long term (Hynes & Prasad, 1997).

Experience with deep learning cycles increases tolerance for short-term difficulties and resource shifts away from production toward learning. Managers assume that good decisions lead to measurable improvements, and that their task is to avoid problems or fix them (and punish the culprit) when they emerge. A systemic view (e.g., Senge, 1990; Senge & Sterman, 1991) suggests that changes take time to unfold and that performance measures typically get worse before they get better (since resources are shifted away from immediate needs). In a system, there are no “root causes” or “first causes” since causes are also effects. Problems are not simply someone’s fault, but rather a feature of the system. However, this does not leave individuals powerless or remove responsibility (as control-oriented managers fear). A deeper understanding of the system reveals leverage points, suggests new interventions, and allows a richer conversation about the implications of interventions.

#### A Framework for Organizational Learning Stages

The four organizational learning stages can be thought of as a progression, but the stages can also be examined for underlying dimensions and symmetries. In Figure 4, we organize the four stages into a 2 X 2 table representing two dimensions: (1) single- and double-loop learning and (2) improvising and structuring. As we have discussed earlier in this paper, single-loop learning adjusts goal-oriented actions based on feedback to

better achieve the same goal (Argyris, et al., 1985). In double-loop learning, a deeper inquiry surfaces and challenges underlying assumptions and values regarding the selection of that goal (see Figure 1). These distinctions are related to exploitation and exploration (March, 1991) and to incremental and radical learning (Miner & Mezas, 1996). Improvising is a process of acting intuitively into an emerging situation rather than following structured procedures or plans (Weick, 1998), drawing on imagination and finding interpretations that can be integrated meaningfully into unfolding action (cf. Crossan et al., 1999). Structuring is a collective achievement of consistency and predictability by integrating knowledge into routines and cultural meanings that are institutionalized throughout the organization (Crossan et al., 1999).

\*\*\* Insert Figure 4 \*\*\*

The dimensional analysis captures some of the interplay between control, represented by structuring and improvisation, and learning, represented by single- and double-loop learning. The progression of stages we have observed in our studies seems to zigzag through the dimensions. In particular, the transition from constrained to open involves changing *both* dimensions, moving from structured single-loop learning to improvised double-loop learning. Perhaps this is another reason why the transition is so difficult and organizations adhere so strongly to the constrained stage.

There will undoubtedly be arguments about whether the fourth stage should emphasize more rather than less improvisation. Weick et al. (1999) appear to argue for considerable improvisation in high-reliability organizations. However, most of their examples involve fast-paced activities with structured tasks and well-practiced participants (e.g., aircraft carriers, jazz musicians). We came to our conclusions from a

small number of case studies of reflective practices (e.g., problem investigations) that engage multidiscipline teams in the analysis of operational surprises and disappointments. The balance between, or integration of, structure and improvisation, or learning and control, could be different in such settings. However, we reiterate that organizations at the fourth stage have come through the earlier stages and are exhibiting behaviors and capabilities from the earlier stages. Thus, the deep learning stage *adds* structured, disciplined learning capabilities onto the learning values and improvisational capabilities of the open stage (as well as the capabilities of the local and constrained stages).

Levels of learning extend beyond the organization as well. The nuclear power industry has built extensive capabilities for exchanges across plants and industry-wide learning. More opportunities exist for learning across industry, including conferences, workshops, and consultants who bridge industries. For this reason, the fourth organizational stage of deep learning is not an ending point, since the organization itself can be thought of as *local* within industry, value chain, professional, and national boundaries that can be bridged with new learning capabilities. Thus, an organization may cycle between exploration and exploitation not only within the organization (Crosson &Hurst, 2001) but also across multiple organizations.

## **DISCUSSION AND CONCLUSIONS**

### Stages of Learning

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 constrained 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 striving to learn more and to challenge common assumptions.

The progression from constrained learning to open learning seems particularly difficult. Expression of minority viewpoints creates conflict, as in the case of Millstone. Courting ambiguity and loosening structure can provoke anxiety. Powerful stakeholders may be threatened if their assumptions, skills, and status are challenged. Anxiety and fear are impediments to learning and change, whereas a desire to learn (prompted by intrinsic interest or external threats) and trust and support from others engages learning and imagination (Amabile, 1996; Edmondson, 1999; Schein, 1992).

The new learning initiatives we studied required significant management support, often in reaction to external pressure building from years of poor performance or safety problems. Senior managers typically created an opportunity for new learning, but that opportunity had to attract internal and external participants, such as the problem investigation experts from a sister company in the Charge Heater case or the Director of the Employee Concerns Program and dozens of other Millstone managers and employees who gradually shaped a safety conscious work environment. Our results suggest that, to some degree at least, attitudes and values favorable to learning emerged before skills and practices that supported deeper learning. The opportunity and desire to learn began to stimulate acquisition of new skills that required significant commitment, discipline, and learning-in-action.

In portraying the stages as both a linear sequence and a 2X2 table, we leave open the question of whether these are “stages” that occur in a fixed sequence or something more like capabilities that are distributed in a complex and dynamic way across location and time. While issuing the obligatory call for further research, we also maintain that both frameworks are valid. Organizations exhibit all the capabilities we have discussed, in various mixtures and degrees, but some capabilities are more difficult to develop and sustain, and therefore tend to flower later. Organizations could under-invest or over-invest in capabilities and practices at any stage, thereby failing to meet environmental demands or capitalize fully on their resources. Although our examples point to the need to challenge assumptions and learn more deeply, too much double-loop learning could waste resources and destabilize the organization (Rudolph & Repenning, in press).

#### Boundary Activities and Trustful Interrelating

A consistent theme underlying the empirical examples is the importance of boundary spanning or bridging activities (Ancona & Caldwell, 1992; Cook & Brown, 1999; Yan & Louis, 1999). Problem investigations are opportunities to enlarge communities of practice (Brown & Duguid, 1991), develop new vocabularies of thought and action, and address common problems. Boundary objects such as reports and cause and effect diagrams facilitated productive conversations. Team and organizational learning depended on access to information and diversity of representation on the team.

Direct management involvement on the teams, which was more likely at the industry-leading plant, was important for change. But management involvement can be constraining, as in the mistrust between managers and employees at Millstone and the defensiveness of managers that partially inhibited problem investigation teams. The US

Army After Action Review process, for example, has senior officers take the facilitator role in discussions of problems and successes in the recent action (Garvin, 2000). The demeanor of the officers and the strict separation of information in this process from the personnel evaluation process are important for building trust among the troops.

The empirical studies illustrate that “learning occurs at several different but interrelated levels at the same time” (Levinthal & March, 1993, p. 100). Plant managers and other senior managers established new values and set new conditions at some of the plants, but numerous individuals and small groups improvised new learning practices as well as new insights. The team reports became a record of organizational learning, and the conversations around the reports were an opportunity to bridge differences between managers and team members. Team members can then bring new respect for other groups and new relationships back to their groups (Gruenfeld et al., 2001).

In conclusion, our empirical studies and emerging model of organizational learning are beginning to show how organizational learning can link with managerial control and structure in a variety of ways. We agree with Perrow (1979), Adler and Borys (1996), and others that bureaucracy and structure are not inherently destructive to innovation and learning, although mental models that link structure with control can jeopardize intrinsic motivation and full engagement. We believe that the development of a full range of learning capabilities requires connections among action, reflection, and emotion, in *energizing* knowledgeable action (implementation) based on actionable knowledge (sensemaking) (Argyris, et al, 1985; Crossan, et al., 1999; Weick et al., 1999). Future research will undoubtedly put more flesh on the bones of these ideas, and contribute alternative ways to think about organizational learning.

## REFERENCES

- Adler, P. S. (1993). The 'learning bureaucracy': New United Motor Manufacturing, Inc. *Research on Organizational Behavior, 15*, 111-94.
- Adler, P. S. & Borys, B. (1996). Two types of bureaucracy: Enabling and coercive. *Administrative Science Quarterly, 41*, 61-89.
- Amabile, T. M. (1996). *Creativity in context*. Boulder, CO: Westview.
- Aldrich, H. (1999). *Organizations evolving*. Thousand Oaks, CA: Sage.
- Ancona, D. G. & Caldwell, D. F. (1992). Bridging the boundary: External activity and performance in organizational teams. *Administrative Science Quarterly, 37*, 634-55.
- Argote, L. (1999). *Organizational learning: Creating, retaining and transferring knowledge*. Norwell, MA: Kluwer.
- 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. (1996). Unrecognized defenses of scholars: Impact on theory and research. *Organization Science, 7*, 79-87.
- Argyris, C., Putnam, R., & Smith, D. M. (1985). *Action science: Concepts, methods and skills for research and intervention*. San Francisco: Jossey-Bass.
- Argyris, C. & Schön, D. (1996). *Organizational learning II: Theory, method, and practice*. Reading MA: Addison-Wesley.
- Bandura, A. & Walters, R. H. (1963). *Social learning and personality development*. New York: Holt.

- Barley, S. R. & Kunda, G. (2001). Bringing work back in. *Organization Science*, 12, 76-95.
- Barker, J. R. (1993). Tightening the iron cage: Concertive control in self-managing teams. *Administrative Science Quarterly*, 38, 408-37.
- Baron, J. & Hershey, J. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, 54(4), 569-79.
- Bartunek, J. M. (1984) Changing interpretive schemes and organizational restructuring: The example of a religious order. *Administrative Science Quarterly*, 29, 355-372.
- Berlyne, D. E. (1960). *Conflict, arousal and curiosity*. New York: McGraw-Hill.
- Bohn, R. E. (1994). Measuring and managing technological knowledge. *Sloan Management Review*, Fall, 61-73.
- Bourrier, M. (1999). Constructing organizational reliability: The problem of embeddedness and duality. In J. Misumi, B. Wilpert, & R. Miller (Eds.), *Nuclear safety: A human factors perspective*. Philadelphia: Taylor & Francis.
- Brown, J. S. & Duguid, P. (1991). Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science*, 2, 40-57.
- Campbell, D. T. (1969). Variation and selective retention in social-cultural evolution. *General Systems*, 14, 69-85.
- 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., Hatakenaka, S., & Rudolph, J. W. (2001). Problem Investigation Teams in High-Hazard Industries: Creating and Negotiating Knowledge for Organizational Improvement. *Management Science*, under review.
- Carroll, J. S., Rudolph, J. W., Hatakenaka, S., & Boldrini, M. (in press). The Difficult Handoff from Incident Investigation to Implementation. Paper presented at the New Technologies and Work Conference, Bad Homburg, Germany, May, 1999. To be a book chapter in an edited book from the conference contributions.
- Carroll, J. S., Rudolph, J. W., Hatakenaka, S., Wiederhold, T. L., & Boldrini, M. (2001). Learning in the Context of Incident Investigation Team Diagnoses and Organizational Decisions at Four Nuclear Power Plants. In E. Salas & G. Klein (eds.) *Linking expertise and naturalistic decision making*. Mahwah, NJ: Lawrence Erlbaum, pp. 349-65.
- Carroll, J. S., Sterman, J., & Marcus, A. A. (1988). 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.
- 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.
- Cole, R. E. (1998). Learning from the quality movement: What did and didn't happen and why? *California Management Review*, 41, 43-73.
- Cook, S. D. N. & Brown, J. S. (1999). Bridging epistemologies: The generative dance between organizational knowledge and organizational knowing. *Organization Science*, 10, 381-400.
- Corcoran, W. (2001). Personal communication, August.
- Crossan, M. M. & Hurst, D. (2001). *Strategic renewal as improvisation: Reconciling the tension between exploration and exploitation*. London, Ontario: Richard Ivey School of Business, unpub. ms.
- Crossan, M. M., Lane, H. W., & White, R. E. (1999). An organizational learning framework: From intuition to institution. *Academy of Management Review*, 24, 522-37.
- Crossan, M. M., Lane, H. W., White, R. E., & Djurfeldt, L. (1995). Organizational learning: Dimensions for a theory. *International Journal of Organizational Analysis*, 3, 337-60.
- 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.
- Davoudian, K., Wu, J., & Apostolakis, G. (1994). Incorporating organizational factors into risk assessment through the analysis of work processes. *Reliability Engineering and System Safety*, 45, 85-105.
- Deming, W. E. (1986). *Out of the crisis*. Cambridge, MA: MIT Center for Advanced Engineering Study.
- Dougherty, D. (1992). Interpretive barriers to successful product innovation in large firms. *Organization Science*, 3, 179-202.
- Dougherty, D. & Hardy, C. (1996). Sustained product innovation in large, mature organizations: Overcoming innovation-to-organization problems. *Academy of Management Journal*, 39, 1120-53.
- Dutton, J. E. (1993). The making of organizational opportunities: An interpretive pathway to organizational change. *Research in Organizational Behavior*, 15, 195-226.
- Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44, 350-383.
- Fiol, C. M. & Lyles, M. A. (1985). Organizational learning. *Academy of Management Review*, 10, 803-813.
- 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. & Lipshitz, R. (1992). Teaching people to shift cognitive gears: Overcoming resistance on the road to Model II. *Journal of Applied Behavioral Science*, 28(1), 118-136.
- Frisch, D. & Baron, J. (1988). Ambiguity and rationality. *Journal of Behavioral Decision Making*, 1, 149-57.
- Gardner, W. L. & Martinko, M. J. (1996). Using the Myers-Briggs Type Indicator to study managers: A literature review and research agenda. *Journal of Management*, 22, 45-83.
- Garvin, D. A. (2000). *Learning in action: A guide to putting the learning organization to work*. Boston, MA: Harvard Business School.
- Glynn, M. A., Lant, T. K., & Milliken, F. J. (1994). Mapping learning processes in organizations: A multi-level framework linking learning and organizing. *Advances in Managerial Cognition and Organizational Information Processing*, Vol. 5. Greenwich, CT: JAI Press, pp. 43-83.
- 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.
- Guillen, M. F. (1994). *Models of management: Work, authority, and organization in a comparative perspective*. Chicago: U. Chicago.
- Hackman, J. R. & Wageman, R. (1995). Total quality management: Empirical, conceptual, and practical issues. *Administrative Science Quarterly*, 40, 309-42.

- Hannon, J. N. et al. (1996). *Millstone Independent Review Group, Handling of Employee Concerns and Allegations at Millstone Nuclear Power Station Units 1, 2, & 3 from 1985 - Present*. October 24.
- Hannon, M. T. & Freeman, J. The population ecology of organizations. *American Journal of Sociology*, 82, 929-66.
- Higgins, E. T. (1998). Promotion and prevention: Regulatory focus as a motivational principle. *Advances in Experimental Social Psychology*, 30, 1-46.
- Hirschhorn, L. (1993). Hierarchy versus bureaucracy: The case of the nuclear reactor. In K. H. Roberts (ed.), *New challenges to understanding organizations*. New York: Macmillan, pp. 137-50.
- Huber, G. P. (1991). Organizational learning: The contributing processes and the literatures. *Organization Science*, 2, 88-115.
- Hynes, T. & Prasad, P. (1997). Patterns of "mock bureaucracy" in mining disasters: An analysis of the Westray coal mine explosion. *Journal of Management Studies*, 34(4), 610-623.
- Isaacs, William (1999). *Dialogue and the art of thinking together*. New York, NY: Currency Doubleday.
- 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.
- Jackson, S. E. (1992). Team composition in organizational settings: Issues in managing an increasingly diverse work force. S. Worshel, W. Wood, & J. A. Simpson, Eds. *Group process and productivity*. Sage, Newbury Park, CA.

- Jasper, J. M. (1990). *Nuclear politics: Energy and the state in the United States, Sweden, and France*. Princeton, NJ: Princeton U.
- 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.
- 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.
- Kepner, C. H. & Tregoe, B. B. (1981). *The new rational manager*. Princeton, NJ: Princeton.
- 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.
- Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology*, 32, 311-28.
- Lanir, Z. (1986). *Fundamental surprise*. Eugene, OR: Decision Research.

- Lave, J. & Wenger, E. (1991). *Situated learning: Legitimate peripheral performance*. Cambridge, MA: Cambridge University.
- Levinthal, D. A. & March, J. G. (1993). The myopia of learning. *Strategic Management Journal*, 14, 95-112.
- 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. & Simon, H. (1958). *Organizations*. New York: Wiley.
- 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.
- Miner, A. S. & Mezias, S. J. (1996). Ugly duckling no more: Past and futures of organizational learning research. *Organization Science*, 7, 88-99.
- Mirvis, P. H. (1996). Historical foundations of organizational learning. *Journal of Organizational Change Management*, 9, 13-31.
- Moreland, R. L. & Myaskovsky, L. (2000). Exploring the performance benefits of group training: Transactive memory or improved communication? *Organizational Behavior and Human Decision Processes*, 82, 117-33.

- 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.
- Myers, I. B. & Myers, P. B. (1980). *Gifts differing*. Palo Alto, CA: Consulting Psychologists.
- Nelson, R. & Winter, S. (1981). *An evolutionary theory of economic change*. Cambridge, MA: Harvard U.
- Nisbett, R. E. & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- Nonaka, I. & Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. New York: Oxford University Press.
- Nutt, P. C. (1999). Surprising but true: Half the decisions in organizations fail. *Academy of Management Executive*, 13, 75-90.
- Oliver, C. (1992). The antecedents of deinstitutionalization. *Organization Studies*, 13, 563-88.
- O'Reilly, C. A. (1978). The intentional distortion of information in organizational communication: A laboratory and field approach. *Human Relations*, 31, 173-93.
- O'Reilly, C. A. & Chatman, J. A. (1996). Culture as social control: Corporations, cults, and commitment. *Research in Organizational Behavior*, 18, 157-200.
- Pan, S. L. & Scarbrough, H. (1999). Knowledge management in practice: An exploratory case study. *Technology Analysis and Strategic Management*, 11, 359-74.

- Perin, C. (1998). Operating as experimenting: Synthesizing engineering and scientific values in nuclear power production. *Science, Technology, and Human Values*, 23, 98-128.
- Perrow, C. (1970). *Organizational analysis: A sociological view*. Belmont, CA: Wadsworth.
- Perrow, C. (1979). *Complex organizations: A critical essay (2<sup>nd</sup> ed.)*. Glenview, IL: Scott Foresman.
- Perrow, C. (1984). *Normal accidents*. New York: Basic Books.
- Perrow, C. (1994). Accidents in high-risk systems. *Technology Studies*, 1, 1-19.
- Perron, M. J. & Friedlander, R. H. (1996). The effects of downsizing on safety in the CPI/HPI. *Process Safety Progress*, 15(1), 18-25.
- Pfeffer, J. (1982). *Organizations and organization theory*. Boston: Pitman.
- Popper, M. & Lipshitz, R. (1998). Organizational learning mechanisms: A structural and cultural approach to organizational learning. *Journal of Applied Behavioral Science*, 34, 161-179.
- Powell, W. W. & DiMaggio, P. J., eds. (1991). *The new institutionalism in organizational analysis*. Chicago: U. Chicago.
- Pugh, D. S. & Hickson, D. J. (1976). *Organizational structure and its context: The Aston Program*. Lexington, MA: D. C. Heath.
- 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. (1990). The role of error in organizing behavior. *Ergonomics*, 33, 1185-1190.
- 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.
- Reason, J. (1997). *Managing the risks of organizational accidents*. Brookfield, VT: Ashgate.
- Rees, J. V. (1994). *Hostages to each other: The transformation of nuclear safety since Three Mile Island*. Chicago: U. of Chicago.
- Repenning, N. P. & Sterman, J. D. (2000). Getting quality the old-fashioned way: Self-confirming attributions in the dynamics of process improvement. In R. Scott & R. Cole (eds.), *Improving research in total quality management*. Newbury Park, CA: Sage, pp. 201-35.
- Roberts, K. H. (1990). Some characteristics of one type of high reliability organization. *Organization Science*, 2, 160-176.
- 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.
- Rooke, D. & Torbert, W. R. (1998). Organizational transformation as a function of CEO’s developmental stage. *Organization Development Journal*, 16, 11-28.
- Royal Society (1992). *Risk: Analysis, perception, and management*. London: author.

- Rudolph, J. W. & Reppenning, N. P. (in press). Disaster dynamics: Understanding the role of quantity in organizational collapse. *Administrative Science Quarterly*.
- 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.
- Sagan, S. D. (1993). *The limits of safety: Organizations, accidents, and nuclear weapons*. Princeton, NJ: Princeton University Press.
- Schaaf, T. W. Van der, Lucas, D. A., & Hale, A. R. eds. (1991). *Near miss reporting as a safety tool*. Butterworth-Heinemann, Oxford
- 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.
- Senge, P. (1990). *The fifth discipline*. New York: Doubleday.
- Senge, P. & Sterman, J. D. (1991). Systems thinking and organizational learning: Acting locally and thinking globally in the organization of the future. In T. Kochan & M. Useem (Eds.), *Transforming organizations*. Oxford: Oxford University Press, pp. 353-370.
- Shrivastava, P. (1987). *Bhopal: Anatomy of a crisis*. New York: Basic Books.
- 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.
- Star, S. L. (1989). The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving. In M. Huhns & L. Gasser (Eds.), *Readings in distributed artificial intelligence*. Menlo Park, CA: Morgan Kaufman.
- Staw, B. M., Sandelands, L. E., & Dutton, J. E. (1981). Threat-rigidity effects in organizational behavior: A multilevel analysis. *Administrative Science Quarterly*, *26*, 501-24.
- Tetlock, P. E. (1983). Cognitive style and political ideology. *Journal of Personality and Social Psychology*, *45*, 118-126.
- Thorndike, E. L. (1911). *Animal intelligence*. New York: Macmillan.
- Torbert, W. R. (1972). *Learning from experience: Toward consciousness*. New York, Columbia University Press.
- 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.
- Turner, B. A. (1978). *Man-made disaster*. London: Wykeham.
- US Nuclear Regulatory Commission (1975). *Reactor safety study: An assessment of accident risks in US commercial nuclear power plants*. Washington, DC: US Nuclear Regulatory Commission Report, WASH-1400, NUREG-75/014.
- Vaughn, D. (1996). *The Challenger launch decision*. Chicago: U. Chicago.

- von Hippel, E. (1994). "Sticky information" and the locus of problem solving: Implications for innovation. *Management Science*, 40(4), 429-439.
- Webster, D. M. & Kruglanski, A. W. (1994). Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology*, 67, 1049-62.
- 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. & Roberts, K. H. (1993). Collective mind in organizations: Heedful interrelating on flight decks. *Administrative Science Quarterly*, 3, 357-381.
- Weick, K. E. & Sutcliffe, K. M. (2001). *Managing the unexpected: Assuring high performance in an age of complexity*. San Francisco: Jossey-Bass.
- Weick, K. E., Sutcliffe, K. M. & Obstfeld, D. (1999). Organizing for high reliability: Processes of collective mindfulness. *Research in Organizational Behavior*, 21, 81-123.
- Weick, K. E. & Westley, F. (1996). Organizational learning: Affirming an oxymoron. In S. R. Clegg, C. Hardy, & W. R. Nord (eds), *Handbook of organization studies*. Thousand Oaks, CA: Sage, pp. 440-58.
- Westrum, R. (1988). *Organizational and inter-organizational thought*. Paper presented at the World Bank Conference on Safety Control and Risk Management.
- White, K. B. (1984). MIS project teams: An investigation of cognitive style implications. *MIS Quarterly* 95-101.

- White, P. A. (1988). Causal processing: Origins and development. *Psychological Bulletin*, 104, 36-52.
- 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.
- Williams, K. Y. & O'Reilly, C. A. (1998). Demography and diversity in organizations. *Research in Organizational Behavior*, 20, 77-140.
- Wood, W., S. Lundgren, J. A. Ouellette, S. Busceme, T. Blackstone. 1994. Minority influence: A meta-analytic review of social influence processes. *Psychological Bulletin*, 115, 323-45.
- Woods, D. D., Johannesen, L. J., Cook, R. I., & Sarter, N. B. (in press). *Behind human error: Cognitive systems, computers, and hindsight*. Wright-Patterson Air Force Base, OH: Crew Systems Ergonomics Information Analysis Center.
- Yan, A. & Louis, M. (1999). The migration of organizational functions to the work unit level: Buffering, spanning and bringing up boundaries. *Human Relations*, 52, 25-47.

Figure 1

Single- and Double-loop Learning

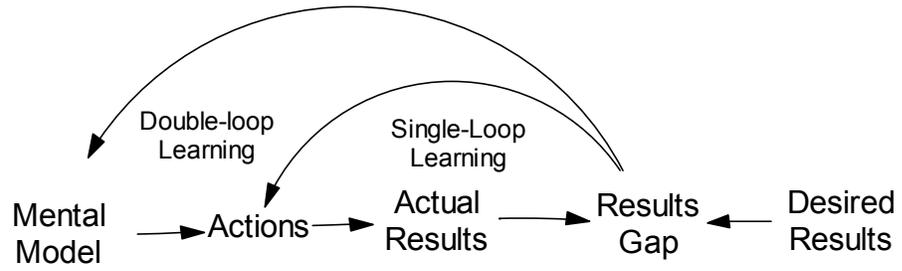
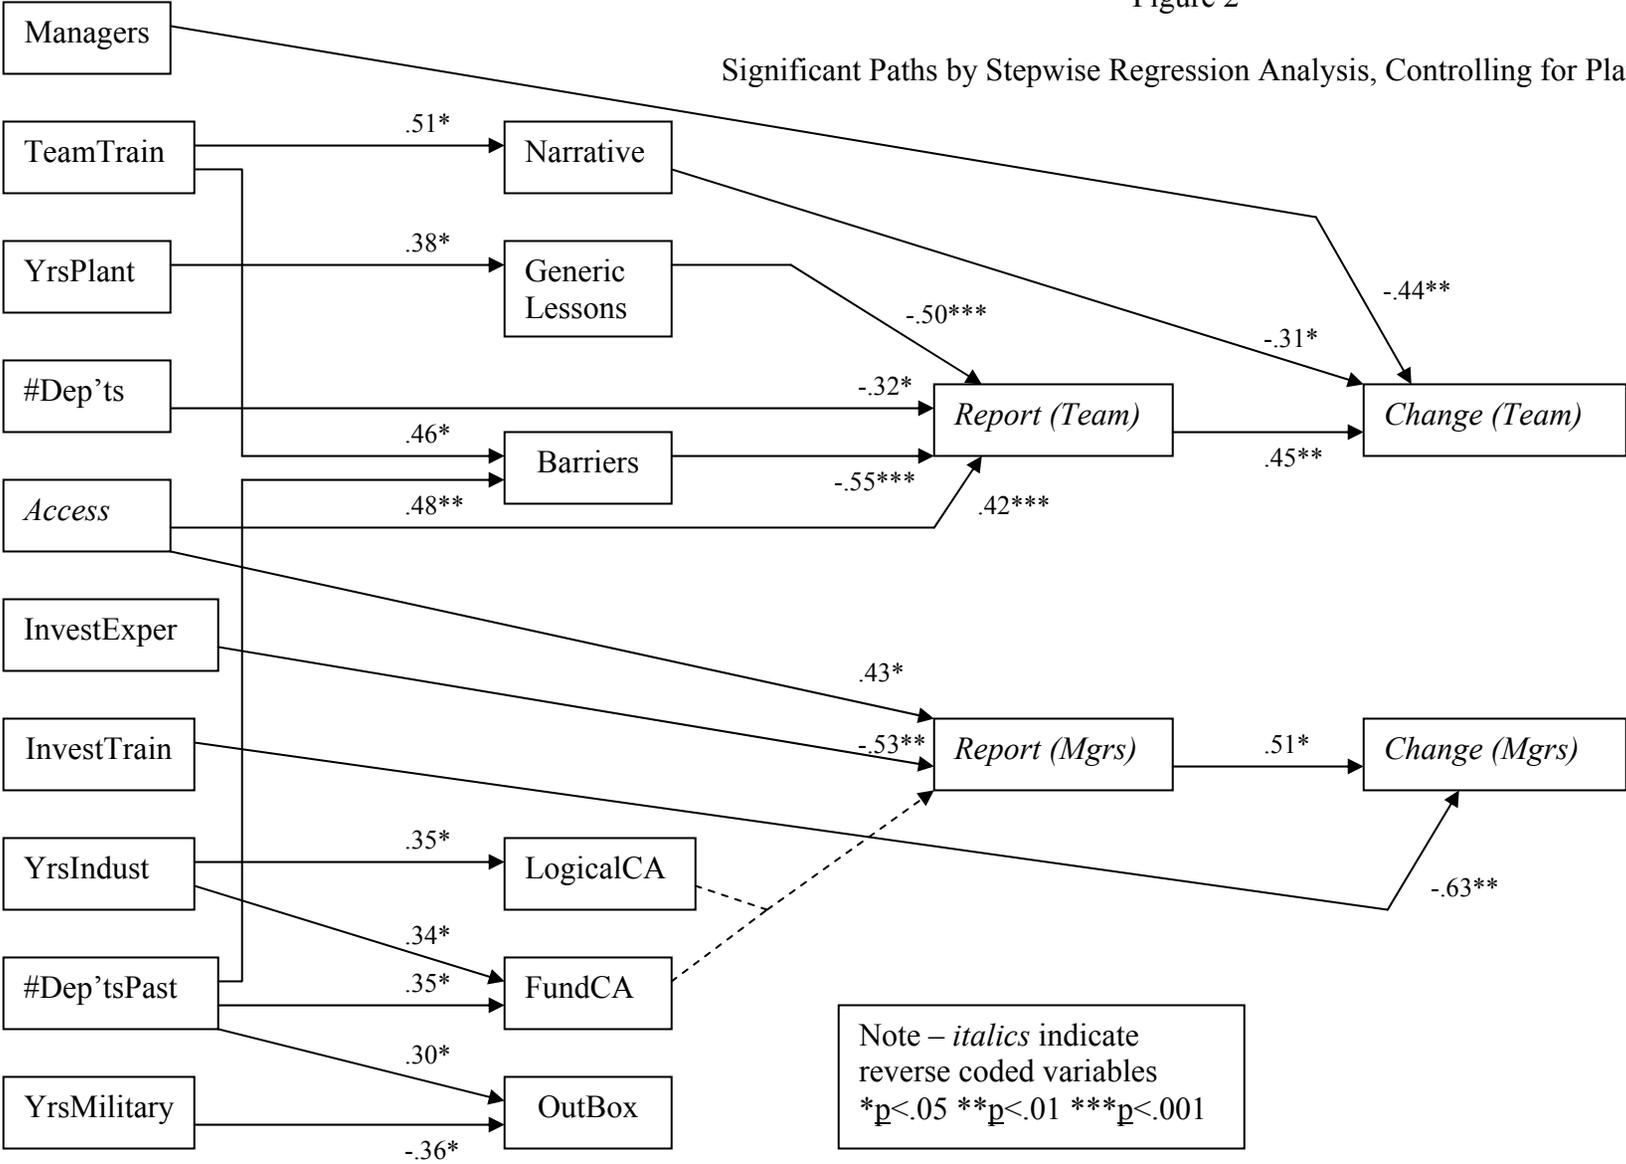


Figure 2

Significant Paths by Stepwise Regression Analysis, Controlling for Plant



Note – *italics* indicate reverse coded variables  
 \*p<.05 \*\*p<.01 \*\*\*p<.001

Figure 3

The Four Stages of Organizational Learning

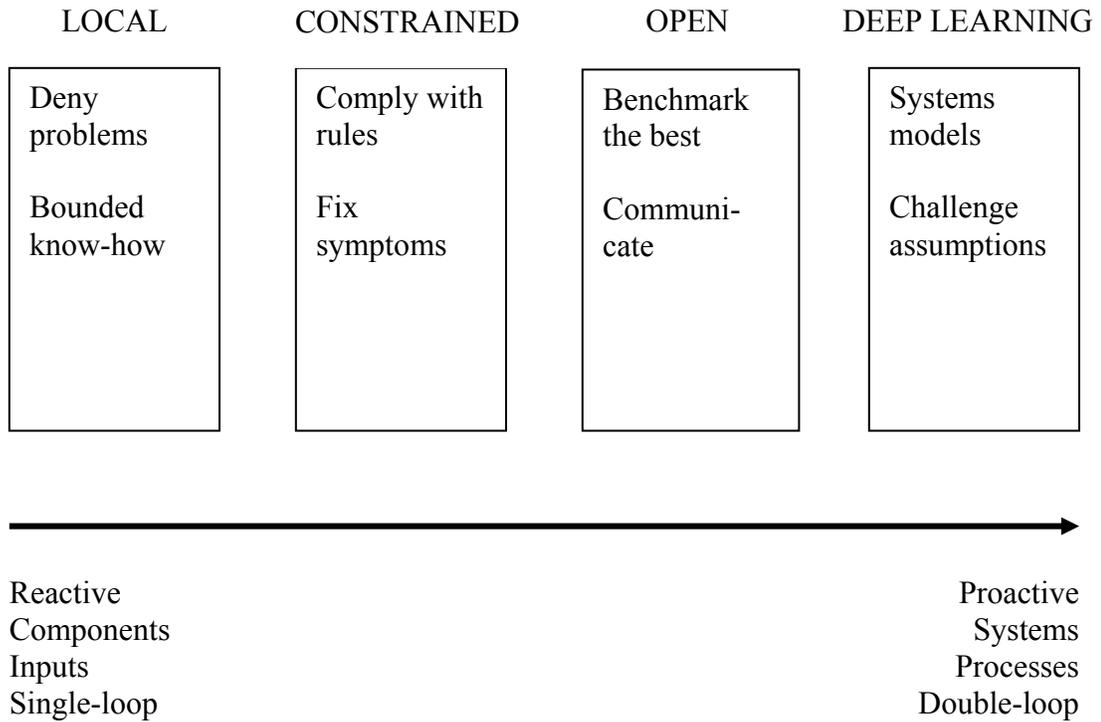


Figure 4

Framework For The Four Stages of Organizational Learning

