

**ORGANIZATIONAL LEARNING ACTIVITIES
IN HIGH HAZARD INDUSTRIES:
THE LOGICS UNDERLYING SELF-ANALYSIS**

John S. Carroll

SS Working Paper No. 3936

January 1997

Organizational Learning Activities in High-Hazard Industries:
The Logics Underlying Self-Analysis

John S. Carroll

Massachusetts Institute of Technology Sloan School of Management

Address for reprints: John S. Carroll, MIT Sloan School, 50 Memorial Drive, Cambridge, MA 02139,
USA.

Organizational Learning Activities in High-Hazard Industries:

The Logics Underlying Self-Analysis

John S. Carroll

Massachusetts Institute of Technology Sloan School of Management

ABSTRACT

Organizational learning takes place through activities performed by individuals, groups, and organizations as they gather and digest information, imagine and plan new actions, and implement change. I examine the learning practices of companies in two industries -- nuclear power plants and chemical process plants -- that must manage safety as a major component of operations and therefore must learn from precursors and near-misses rather than exclusively by trial-and-error. Specifically, I analyze the logics underlying incident reviews, root cause analysis teams, and self-analysis programs as revealed in administrative documents, interviews, on-site observations, and collaborative workshop discussions. In organizations with fragmentary, myopic, and disparate understandings of how the work is accomplished, there are likely to be more failures to learn from operating experience, recurrent problems, and cyclical crises. Enhanced learning requires a more comprehensive set of lenses through which to understand past experience and to see new possibilities. I discuss some examples of companies trying out new ways of thinking and acting.

Organizational Learning Activities in High-Hazard Industries:

The Logics Underlying Self-Analysis*

John S. Carroll

Massachusetts Institute of Technology Sloan School of Management

INTRODUCTION

At one nuclear power plant, the Station Manager explained the plant's excellent safety record by saying, "What we do around here doesn't always make sense" (Carroll & Perin, 1993). What he meant is that the ways people achieve and maintain good performance are not easy to articulate within well-understood principles of organization and management. In short, he found it difficult to explain these activities and practices to his bosses (indeed, he might have to hide information to avoid trouble).

But why doesn't it make sense? Why should effective behaviors and activities not be explicable and perhaps not discussible (cf., Argyris and Schon, 1996)? The central argument of this paper is that the difficulty lies in the available "mental models" (e.g., Senge, 1990) or understandings of organizations, people, and technologies. When those mental models legitimate only certain types of behaviors, and exclude whole classes of effective behaviors, then there is need to broaden the models. When different knowledge bases and viewpoints cannot be negotiated across levels of hierarchy and occupational specialties, then organizations cannot make sense of events (Weick, 1995) in ways that support effective learning.

For example, during a visit to a nuclear power plant, a member of our research team asked to see the organization chart, which we use regularly as a way of getting some background context. The chart was dutifully handed over, but the respondent added that the organization chart did not represent the way people really work at the plant. Instead, he drew a set of intersecting circles very different in look and feel from an organization chart (see Figure 1), suggesting that people have to communicate and work across boundaries in more of a network or system. Later that year, at a meeting of the sponsor representatives from our research program[1], I presented these drawings as a way to illustrate a more organic form of organization different from the familiar “machine bureaucracy.” During the next coffee break, a Vice President from another nuclear power utility approached me and expressed concern over the intersecting circles, saying “once you let people out of their boxes, there will be chaos.”

Insert Figure 1

These vignettes suggest that employees at different levels in the hierarchy can have different understandings of how a plant operates, and as a result they may not communicate easily (cf. Schein, in press). The source of the intersecting-circles description was someone actively engaged in everyday operations at the worker level. The source of concern about lack of control and letting people out of their boxes was a Vice President located at corporate headquarters two hours drive from their nuclear power plant. The Station Manager who said “what we do around here doesn’t always

make sense” is located in the middle, at the intersection of the plant and the corporation: he understands what the plant does in order to succeed yet he perceives that his executive bosses might not accept those activities as legitimate.

The above examples highlight an important issue for organizational learning: what does it mean for an “organization” to “learn” or “know” something, apart from the knowledge of individuals within the organization? Organizational knowledge is embodied in physical artifacts (equipment, layout, data bases), organizational structures (roles, reward systems, procedures), and people (skills, values, beliefs, practices) (cf., Kim, 1993; Levitt and March, 1988; Schein, 1992). Different parts of the organization (e.g., plant vs. corporate) “know” different things about how work is done. This is a necessary feature of complex organizations where no one person can know enough about every specialty, nor appreciate both the “big picture” and the details, yet there must be sufficient coherence or alignment among these parts to operate effectively. Although organizations cannot learn without people, organizational knowledge exists in the interdependencies among these supporting and conflicting repositories of information.

In the remainder of this paper, I lay out a framework for organizational learning. Organizational learning is a vast topic (e.g., Argyris and Schon, 1996; Cohen and Sproull, 1996; Levitt and March, 1988) and this paper restricts our focus in several ways. First, I am less concerned with *what* is learned and *who* does the learning than with *how* that learning is accomplished. This focuses the discussion on specific learning activities. Second, I am more interested in sources of learning from analysis of an organizations’ own experiences and internal debates rather than learning by imitation, diffusion of

technical innovations, benchmarking, exchange of best practices, or other external sources of information. Third, I draw examples from two industries -- nuclear power plants and chemical process plants -- that have to manage hazardous operations. I argue that this makes their learning processes particularly challenging, because they cannot learn solely by trial-and-error due to the catastrophic costs of severe accidents (March, Sproull, and Tamuz, 1991; Weick, 1987). The analysis of incident reviews, root cause analysis, and self-analysis programs leads to a discussion of new conceptual tools that can help organizations rethink their learning activities.

ORGANIZATIONAL LEARNING FRAMEWORK

In order to understand organizational learning *outcomes* such as performance improvements, it is important to examine the learning *processes* that are enacted within work *activities* to produce those outcomes, and the learning *resources* that provide capabilities for those activities (cf. resource-based theory of the firm, Penrose, 1968).

These are given schematic form in Figure 2.



Many researchers characterize learning as a generic feedback cycle (e.g., Argyris and Schon, 1996; Kim, 1993; Kolb, 1984; Daft and Weick, 1984; Schein, 1987), and I continue that approach by describing four ordered learning processes that take place within and across individuals, groups, departments, organizations, and institutions. These four processes are:

- 1) Observing - noticing, attending, heeding, tracking;
- 2) Reflecting - analyzing, interpreting, diagnosing;
- 3) Creating - imagining, designing, planning; and
- 4) Acting - implementing, doing, testing.

This learning process cycle takes place at individual, group, departmental, organizational, and institutional levels as various kinds of work activities are carried out, for example, in efforts to verify one's own and others' work, plan of the day meetings, incident reviews, post-job critiques, peer visits among plants, exchanges of good practices, and so forth. Each activity requires resources -- such as time, information, know-how, tools, and procedures -- that are continually developed, depleted, renewed, and changed, thus changing learning capabilities along with achieving learning outcomes.

In high-hazard industries such as nuclear power plants, chemical production plants, offshore oil drilling, and aviation, the nature of the technology and the twin goals of safety and profitability define a unique set of organizational and managerial issues. By comparison with typical manufacturing or service industries, there are greater management challenges to maximize actors' abilities to avoid errors, recover quickly from problems, and learn efficiently from precursors and incidents (LaPorte and Consolini, 1991; Roberts, 1990; Sitkin, 1992; Weick, 1987). Nuclear power plants are the prototypical high-hazard industry because of the potential for catastrophic consequences and the public dread of radiation (Perrow, 1984; Slovic, 1987; Weart, 1988).

To maximize safe performance, designers rely on a defense-in-depth strategy. With layers of barriers, designers anticipate possible threats, and to assure that these

systems are "fail safe," critical safety functions performed by equipment and people are duplicated. The safety strategy results in a complicated and highly interdependent arrangement of dozens of major systems, hundreds of subsystems, and tens of thousands of components whose operation, maintenance, and design changes are governed by written procedures and layers of oversight. Such complexity obscures the impact of particular actions, and the invisibility of latent defects (Reason, 1990) masks the state of the entire system (Perrow, 1984; Turner, 1978): backup systems can be inoperable, oversight can be ineffective, and other vulnerabilities can exist for years without discovery. For example, in 1979 the Davis Besse nuclear power plant management discovered that their auxiliary feedwater system had been inoperable for many years when the main feedwater failed.

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). The Three Mile Island (TMI) event in 1979 constituted a "fundamental surprise" (Lanir, 1986) for the nuclear power industry that cast doubt on that belief. Neither the equipment nor the people functioned as predicted. A stuck-open valve, coupled with ambiguous indicators of equipment status, additional faulty equipment, and operators who had been trained to take particular actions based on an incorrect model of the system, all led to cooling water being drained from the uranium in the reactor so that

the fuel was exposed and partially melted. Because of this combination of problems, the utility lost the use of a billion-dollar unit. The logic of design was shown to be flawed: the complexity of the production system exceeded the capacity of a priori design strategies.

TMI demonstrated the need to complement design logics with operating logics, including learning-through-practice (Kemeny, 1979; Rogovin and Frampton, 1980). The information needed to prevent the TMI event had been available from similar prior incidents at other plants, recurrent problems with the same equipment at TMI, and engineers' critiques that operators had been taught to do the wrong thing in particular circumstances, yet nothing had been done to incorporate this information into operating practices (Marcus, Bromiley, and 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, 1979, p. 192).

Learning is now a central activity in the nuclear power industry. In response to the reports analyzing TMI, and under pressure of further regulatory action, the U.S. nuclear power industry established the Institute for Nuclear Power Operations (INPO) to promote safety and reliability through external reviews of performance and processes, training and accreditation programs, events analysis, sharing of operating information and best practices, and special assistance to member utilities (see Rees, 1994, for a highly favorable

reading of INPO's role). The International Atomic Energy Agency (IAEA) and World Association of Nuclear Operators (WANO) share these goals and serve similar functions worldwide.

In high-hazard industries, priority is placed on finding and fixing problems before they become more serious. This emphasis is reinforced by regulatory practices that consider it important for plants to identify their own problems before the regulators discover them (and before they are discovered through a serious accident). Each plant has ways to reflect on its own operating experience in order to identify problems, interpret the reasons for these problems, and select corrective actions to ameliorate the problems. In the next two sections of the paper, I will analyze closely the incident review program in a nuclear power plant and the root cause analysis of repetitive performance problems in a chemical process plant. These analyses reveal the kinds of resources and capabilities that underlie organizational learning.

AN INCIDENT REVIEW PROGRAM AT A NUCLEAR POWER PLANT

Incident reviews are important vehicles for self-analysis, knowledge sharing across boundaries inside and outside specific plants, and development of problem resolution efforts. Incidents vary from seemingly minor slips with no outward manifestations to highly consequential damage to equipment, safety system actuations, unplanned shutdowns, and injuries to employees that must be reported to the regulator. Plants (and regulators) classify these incidents into categories and expend different levels of resources for analyzing and correcting problems. Special teams are formed to examine serious

incidents, repeat incidents, or negative trends in greater depth, and may involve external teams from corporate headquarters, regulators, INPO, IAEA, or WANO.

At an industry level, dissemination of incident reports and analyses is part of the legacy of Three Mile Island and the mission of INPO. Both INPO and the NRC issue various letters and reports to make the industry aware of incidents as part of operating experience feedback, as does IAEA's Incident Reporting System. INPO, IAEA, and WANO arrange for teams of experts to visit plants to provide various kinds of assistance. For example, IAEA responds to member countries' requests for an international team of experts to visit a plant in order to reanalyze a small number of important recent incidents and gain new ways of understanding them.

Despite regulatory requirements, dissemination of best practices, and other mechanisms that serve to establish and improve incident review programs at each plant, each plant seems to carry out incident reviews in its own way. Thus, for example, there are regulatory requirements for reporting more serious incidents, but the threshold for reporting, documenting, and analyzing less serious incidents and whether this is carried out by special staff or regular line employees, and by individuals or groups, is left up to each plant. As illustrated below, plant employees may not recognize fully the assumptions and logics built into the incident review program which may limit the learning process.

The Carver Incident Review Program

At one nuclear power plant, which we call by the pseudonym of "Carver," administrative documents describe the incident review program as a search for "root cause": "the primary or direct cause(s) that, if corrected, will prevent recurrence of performance

problems, undesirable trends, or specific incident(s).” The documents give additional detail about types of causes:

If the incident involved personnel error, the description should include determination of cognitive personnel error (unintentional action perceived to be correct prior to initiation) or procedure error (error in approved procedure)... If the cause(s) is an equipment failure, the description of the root cause should include the failure mode or mechanism (e.g., valve not included in Preventive Maintenance program, therefore stem not being lubricated, and subsequently corroded and broke).

A list of root cause categories is given:

component failure	work control
man-machine interface	work planning, procedure, documentation
communication	work practices, techniques
training	external cause
work organization	environment
work schedule	unknown, needs more investigation

Within each of the above cause categories, there are subcategories and examples given. However, the detail and extensiveness of these subcategories varies by category: the “component failure” category has pages of subcategories and examples, whereas the category of “work organization” has only a few subcategories. Finally, the documents specify that for review and approval: “The Root Cause Team will discuss the results of

the investigation as soon as possible with the responsible manager. The Team should normally have the concurrence of the manager in determining the corrective action(s).”

The Logic of the Incident Review Program

Although the Carver program appears very sensible and is undoubtedly similar to many other programs in this industry and others, there are some issues that are hidden from view by the assumptions and logics that underlie the program. First, the very concept of “root cause” seems to focus attention on a single cause rather than an exploration of multiple causes or chains of events. This has been called “root cause seduction” (Carroll, 1995) because the idea of a singular cause is so satisfying to our desire for certainty and control. In the Carver documents, recognition is given in an offhand way that there may be cause(s), but the expectation is that a primary cause will be found. This expectation is even built into the computer system that records root cause reports for tracking and analysis which limits data entry to a single root cause. Interestingly, at this plant there was a difference of opinion regarding the proper number of causes to include in a root cause report. The majority believed that there should be only a few root causes, two or three at most, in order to avoid “diluting” the impact of the report. The minority believed that a more extensive list of causes would be helpful for learning and improvement.

Second, some kinds of causes are given more attention than others, based on how well they are understood. The “component failure” designation has extensive detail, serves as the source of many illustrative examples, and has special requirements for interfacing with component problem tracking systems. Procedures and individual human error are also called out for attention by serving as the source of examples and detailed

instructions. These causes tend to be more immediate to the emergence of problems (at the “sharp end” of systems, Reason, 1990) and better understood as causes and for selecting corrective actions, i.e., we know how to “fix” these problems (solution-driven search, Carroll, 1995). In contrast, more systemic and organizational causes that are further back in time (not as “direct”) are left more vague; the documents express a lesser understanding and lesser salience of these causes by their lack of detail, lack of examples, and lack of familiar corrective actions. Reliance on ready-made and available solutions can become a trap when “band-aids” are used to fix symptoms or offer symbolic value but fail to address underlying causes.

The fixing orientation is consistent with American managers' desires for certainty and action. People are not encouraged to look at the linkages among problems and issues, because it creates a feeling of overwhelming complexity and hopelessness in the face of pressure to do something quickly. For example, one engineering executive at a U.S. nuclear power plant commented that, "it is against the culture to talk about problems unless you have a solution." The question is whether this approach works successfully with complex, ambiguous issues that lack ready answers, or with diffuse organizational and cultural processes that are poorly understood, or whether a different approach is needed.

Third, the review and approval process forces the Root Cause Team to negotiate with line management over the content of the report. The virtue of this requirement is that line management must take responsibility for implementing change, and they must therefore have opportunity to provide input and commit to the new actions. Everyone

recognizes that it is easier to produce reports and analyses than it is to create effective change. However, the danger of this approach is that needed change may get buried in politics: the power resides in line management, who may fail to acknowledge issues that reflect badly on them, diminish their status or power, or have solutions that are risky to their own agendas. The anticipation of resistance from line management can lead to sanitizing the report, which is a kind of “acceptability heuristic” (Carroll, 1995; Tetlock, 1985).

If we analyze the Carver program more deeply, the underlying cultural values and assumptions begin to emerge (cf. Schein, 1992). The culture of nuclear power plants, as in most technological organizations, emphasizes the importance of avoiding problems through engineering design and managerial controls and, when necessary, fixing problems. Typically, people who operate the technology are seen as a potential source of trouble that must be controlled by designing people out of the system or providing procedures and training to minimize human error. A joke is told about the new control room staffing plans for the advanced design nuclear power plants. It consists of an operator and a dog. The operator has one job: feed the dog. The dog has two jobs: keep the operator company and, in case of emergency, keep the operator away from the control panel.

There is a presumption that organizations are like machines whose problems can be decomposed into parts, the causes identified, and fixes put in place. The “fixing” orientation looks for linear cause-effect relationships, simplifies problems by decomposing them into well-understood components, and applies specialized knowledge to create technical solutions. This is most clearly represented in Probabilistic Safety Analyses

(PSA) which vividly reveal the expectation that events are enumerable, predictable, and additive, like a complex wiring diagram. Although extremely useful for some purposes, PSAs do not mirror the actual plant: for example, serious accidents nearly always involve being outside the “design basis” or modeling assumptions in PSA (Rasmussen, 1990).

Rochlin and von Meier (1994) associate these assumptions and ways of thinking with the engineering occupational subculture. Indeed, safety is most often conceptualized as an engineering specialty, and most of the literature on safety focuses on equipment and other technical concerns. In contrast, the operating subculture (of occupations that carry out hands-on, real-time functions) has a less formal, more organic, and more dynamic view of the plant. For example, Plant Managers may be more likely to write about safety from a social and organizational perspective (e.g., Chiba, 1991). Schein (in press) suggests that there are typically three subcultures in organizations: an operator culture or line organization that considers work to involve interconnected systems and cooperation among people; an engineering subculture that values technical, error-free solutions; and a CEO subculture that focuses on the financial bottom line.

I held a workshop held at the Carver plant that explored the incident review process, focusing on analysis of a particular recent incident in which a hot water pump was taken out of service to repair an oil leak and then returned to service without the motor being reconnected. The root cause analysis report listed as the root cause that the Electrical Foreman, who is administratively responsible for reading the work order, had failed to verify the work complete before giving up clearance to run the pump. Workshop participants included employees from engineering, operations, maintenance, and other

groups, and varied in hierarchical level. Although this was a visible recent incident at the plant that everyone had heard about, it became evident that different participants knew bits of additional information that generated a lively exchange of details, history, and viewpoints.

With minimal prompting to look for causes beyond the immediate circumstances of the incident, workshop participants were readily able to focus on why the Foreman and several others had failed to prevent the incident, and to draw insights about organizational and cultural issues. This included: (a) the complexity revealed in the numerous hand-offs and layers of management, (b) the use of the same work order form to handle changed job content, (c) the perceived time pressure from Daily Planning that was unnecessary, (d) the fact that the Electrical Supervisor was standing over the Electrical Foreman holding the work order so the Foreman assumed the Supervisor had checked, (e) the Supervisor was new in the job and rank and politeness may have interfered with their communication, and (f) the ways in which specialties optimize their own work without considering their impact on overall plant goals. Each of these issues held lessons for the plant and suggestions for ways to improve work practices.

ROOT CAUSE ANALYSIS AT A CHEMICAL PLANT

Chemical plants, like nuclear power plants, are continuous process technologies with significant hazards arising from toxic chemicals, high temperatures, explosions, etc. At one chemical process plant, management was concerned enough about a long history of costly operating problems, and enlightened enough, to request for the plant a root cause analysis intervention that had begun recently to be offered by corporate employees.

The Root Cause Analysis Intervention

The Root Cause Analysis intervention involved several trainers and two dozen employees assembled for two weeks of effort; half the employees were from the plant, and half had traveled from other operating units to participate in the intervention. These employees were selected to span several different levels of hierarchy and kinds of functional specialties ranging from experienced engineers to nontechnical administrative staff. The intervention consisted of systematic training on how to analyze performance problems and develop a root cause report, taught in the context of analyzing actual problems at the plant and making recommendations to plant management and other important managers (this was a visible event in the company).

This process began with an analysis of all the production losses that could be identified and classified, and selection of six problems that had real “bottom-line” importance to the plant. Each problem was assigned to one of six teams from the group, selected to have both insiders and outsiders, technical and nontechnical members. Each team’s task was to analyze their problem and develop a report to management. The teams received instructions each day in concepts and methods, coordinated to be most useful to them as they proceeded in their analysis. Teams regularly reported their progress to the whole group including tours of the workspace of each team (teams were encouraged to “cheat” or learn from each other). After a week of training and work on the problems, they returned to their jobs for a week and percolated their thinking. They then reconvened for the second week with continued training and analysis culminating in a formal round of

presentations from each team made to each other and to plant management on the last day of the workshop.

Analysis of the Intervention

The root cause analysis training appears to offer a set of simple ideas and tools. For example, there is a careful separation of facts vs. beliefs or what is sometimes called the ladder of inference. Participants are exhorted to be factual, i.e., to anchor their analyses so that every argument can be demonstrated factually. A second tool is to use a list divided into “is/is not” parts. For example, if some motors are wearing out too rapidly, an extensive list of characteristics associated with the motors that are wearing out (“is”) is developed and compared against a list of characteristics associated with motors that are not wearing out (“is not”). A third technique is to keep asking “why?” with the expectation that this question will have to be answered six or more times (a classic Total Quality Management technique). Similarly, instructions are given for developing time lines, drawing flow charts, conducting interviews, and writing reports, lists of possible causes are described to help stimulate thinking and provide shared categories, and the root cause process itself is described in phases so that teams know where they are going, how to get there, how to assess their progress, and how to anticipate and cope with the frustration that often accompanies the process.

Although the training appears simple, the application of those ideas to complex and ambiguous problems is very difficult. The ideas and tools actually build into a discipline with rigorous logic and surprising leverage. Root cause analysis is not taught as a formula, but as a kind of practice: it is the emergent connection of principles to the

details of a particular context. Although directed at concrete problems that need solutions, the real value of root cause analysis training is achieved by changing the way people think about work and the way they seek information and communicate with others. This has the benefit of reducing the introduction of new problems and solving problems more quickly. Indeed, I was told that plants experienced substantial improvements following a root cause intervention *before* solutions to the analyzed problems were even implemented! We might hypothesize that becoming alert to a wider range of plant conditions, developing a questioning attitude, thinking in more systemic ways, communicating more readily with other groups, and believing that problems can be addressed effectively were sufficient to introduce changes in behavior apart from specific “fixes.”

Observation of the team interactions as they progressed during their two weeks of analysis revealed some interesting features of the intervention. First, analyzing any event deeply enough into the chain of causes reveals a lot about how the organization really works. Each of the six problems initially were defined and organized around equipment or particular points in a chemical process, such as a pH control problem or a high failure rate in centrifuge motors. As the investigations proceeded, the analysis led back into the organizational and cultural systems of the plant, such as operator practices for adjusting the equipment or taking equipment down for service whose idiosyncracies across shifts were unrecognized.

Second, as they gathered information about the presenting problem, its context and history, teams uncovered many other problems that turned out to be unrelated, but could

be listed and later addressed. In simply identifying the way a particular process worked and laying out a diagram, which typically involved consulting manuals and diagrams and interviewing knowledgeable employees, numerous examples were found where people thought they knew the chemistry of the system, the current physical layout of equipment, how people actually work with the equipment, etc., but they were mistaken. More importantly, issues surfaced around how designers and operators live in different "thought worlds" (Dougherty, 1990; Schein, in press) and rarely communicate; even different operators failed to tell each other (or to ask) what they needed to know. Incomplete and incorrect procedures were used until a serious problem cropped up. These work practices and ways of thinking affect many aspects of plant performance, so their recognition and change has broad potential impact.

Third, data accessibility is absolutely critical for effective root cause analysis, and effective operations in general. Investigations are done immediately in the nuclear power industry, with a rule-of-thumb that people start to forget important detail in two days. In the case of these root cause teams in the chemical plant, they were investigating events that were weeks, months, or years old. Key informants change jobs, forget details, and even the physical information changes. For example, for one of the team investigations it turned out that a key source of information was continuous process data stored in the plant computers. However, the computers store that information for only one week and then erase the information, maintaining in permanent storage only less informative averages. This particular team got very lucky when a repeat event occurred during their analysis.

Fourth, the composition of teams that crossed disciplinary and hierarchical boundaries conveys an implicit message that no one person or group knows everything needed to solve the problems of the plant. There is a necessary strength to bringing the information and viewpoints of multiple kinds of people to bear on problems, and there are additional payoffs in the form of information exchange, new working relationships, and mutual respect. For example, because the teams began with problems organized around technical features of the equipment, it was natural for team members with technical competence to take their normal role as “expert” and assume that they and only they can provide answers for the plant. In the early interaction of several teams, there were particular engineers whose behavior could only be characterized as arrogant: they assumed there was a technical answer and that nontechnical team members were unable to contribute to the discussion and could only waste their time. Not surprisingly, there were some unpleasant team interactions. Yet, over the course of the analysis, these engineers began to realize that obvious technical “fixes” were not emerging; as problems were traced back they turned out to have underlying causes in work practices and culture involving designers, operators, maintainers, and managers. For these problems, the engineers were not the experts and, in fact, might be myopic in their understanding of human behavior and need the insights of others on the team. This is a humbling but potentially very powerful experience.

DISCUSSION

Differences Between Nuclear Power Plants and Chemical Process Plants

Both nuclear power plants and chemical process plants struggle to avoid problems and, when problems occur, to take corrective actions based on their understanding of causes. There are considerable similarities in the incident review and root cause analysis approaches taken by these industries, partly because they each pay attention to other industries' practices. However, the two examples in this paper are directed at rather different kinds of problems. In the nuclear power plant, the incident review process is engaged typically by a single incident, which may range from a serious event involving damage to equipment, personnel injury, or unplanned shutdown to what appears to be a very minor mistake with no consequences, such as a fire door improperly left open, tools left lying around, or maintenance work carried out without proper clearance or "tagging." Attention is paid to these small incidents not only because of financial or safety consequences, but because of the potential consequences driven by a belief that the accumulation of small incidents increases the probability of a major problem such as Three Mile Island. In the chemical plant, the typical problem was an intermittent but repetitive deviation from a designed chemical or physical process that cost the company thousands or tens of thousands of dollars for each occurrence. Each single incident was tolerated and small adjustments were made locally by a particular operator, engineer, or maintenance worker. The combined impact of these incidents was hardly considered until the root cause intervention.

The nuclear power technology has highly coupled systems of equipment that interact in complex ways; it is the complexity of minor problems cascading in combination into major problems that worries the industry. As they say, the serious problems occur

when "all the holes line up." Intense public scrutiny is placed on the nuclear power industry because of the potential for catastrophe (Chernobyl) and dread of radiation (Slovic, 1987; Weart, 1988). The only way to prevent this is to prevent even the "minor" problems and to learn to recover quickly from problems so the plant is not in an unusual configuration and vulnerable to subsequent errors and surprises.

Since TMI, the nuclear power industry has worked very hard to correct equipment problems and train operators intensively. Equipment that could not be fixed was replaced or redesigned. Many more employees were hired to do more designing, more maintenance, have additional operators on every shift with engineering backgrounds and responsibility for the big picture, quality control, regulatory compliance, and management. One plant that ran with 150 employees in 1975 had 950 employees in 1992, which in the context of deregulation has created enormous new financial pressures. The result is that fewer problems are now occurring (but more attention is paid to less serious ones) but those that do occur are more difficult to understand, involve combinations of causes, almost always have human errors and programmatic issues, and have far less objective data associated with them. To address these problems, general problem solving skills, common sense, and the collection of technical information from within the plant are not enough. Everyone is struggling with how to understand and make improvements to organizational and cultural factors.

Systems Thinking and Organizational Learning

In any industry, well-intentioned, commonplace solutions can fail to help, have unintended side effects, or even exacerbate problems. When problems are understood in linear cause-

effect terms that results in a search for “fixes,” it is common to find a “fixes that fail” scenario (Senge, 1990), as exemplified in Figure 3. Blaming and disciplining particular individuals, intended to encourage accountability, can create an environment in which people do not report problems. For example, when the Federal Aviation Administration provided immunity from prosecution for pilots who reported near midair collisions, the number of reports tripled; when immunity was retracted, the number of reports dropped by a factor of six (Tamuz, 1994). Under such conditions, problems cannot be addressed early, trending is incomplete, and the result can be more problems. Similarly, increased procedural detail and monitoring of compliance, intended to ensure the quality of work, can be perceived as mistrust and regimentation. This may result in loss of motivation, blind compliance to procedures that may still be incomplete, malicious compliance when workers know the right thing to do but also know that only rote compliance is safe from disciplinary action, the departure of skilled workers who find more interesting work elsewhere, and, ultimately, more problems.

Insert Figure 3

From this viewpoint, the incident review process should have learning rather than fixing as its goal. For example, incidents would not be approached with the expectation of finding the single “root” cause of the problem, nor is there a “solution” to the problem. Instead, the incident becomes an occasion to identify and discuss issues, to encourage new insights, and to explore possibilities for change and their consequences. I believe that an

important reason for the success of benchmarking, total quality management, quantitative indicator tracking, and business process reengineering is that the veneer of quantitative modeling legitimates their hidden function of promoting discussion and collaborative learning. This suggests why, all too often, these approaches become new "fixes" that apply limited solutions with disappointing results. Antidotes to this myopia depend upon broader participation and discussion among specialized groups, and can be facilitated by new conceptual lenses (theories) and modeling tools to organize dynamic interdependencies (Senge and Sterman, 1991).

One example of the value of a learning orientation to promote collaboration of theorists, researchers, and operating personnel comes from Du Pont Chemicals (Carroll, Sterman, and Marcus, in press), whose chemical process plants were plagued with equipment failures. In the context of company-wide cost-reduction efforts, a benchmarking study showed that Du Pont spent more than its competitors on maintenance, yet had worse equipment availability. A culture of reactive fire-fighting had developed, with workers regularly pulled off jobs to do corrective maintenance. Responding to the benchmarking study, a series of cost-cutting initiatives were undertaken that had no lasting impact. Finally, one team questioned the basic assumption that reducing maintenance costs could help reduce overall manufacturing costs; they thought that the effects of maintenance activities were tightly linked to so many aspects of plant performance that no one really understood the overall picture.

Du Pont was able to improve maintenance only after a collaborative conceptual breakthrough. An internal team developed a dynamic model of the system of relationships

around maintenance (a "modeling for learning" exercise with the assistance of a researcher/consultant, Senge and 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 programs and policies. Having a broad range of employees with a system-wide understanding of the relationships between operations, maintenance, quality, and costs laid the groundwork for a successful pump maintenance pilot program.

Mobilizing Change

There is a crucial transition between analysis and implementation that is difficult to manage. Many nuclear power plants complain that they have wonderful written analyses but nothing changes. There are difficult "hand-offs" between the root cause team and the line managers who must take action. There are problems in recommending quick fixes that do not address the real causes; sometimes that is what management is demanding of the teams. There are shortages of resources and an unwillingness to commit to longer-term efforts in the face of daily demands. There is resistance to change and cynicism about the "program of the month." Sometimes it takes a major personnel change (such as firing a plant manager or vice president) to unfreeze people and open the space for new actions.

At one troubled nuclear power plant, senior management commissioned a Self-Assessment Team to work full-time for several months re-analyzing the reports from over 20 serious recent incidents, and then to act as "change agents" to jump start corrective

actions. The team developed a detailed set of causal categories and a list of issues, which were investigated further through semi-structured interviews with over 250 employees of the utility. From these data, the team created a list of 15 problems organized in three groupings: (1) Management Philosophy, Skills and Practices; (2) People Performing the Work; and (3) Problem Solving and Follow-Up. In particular, two items within the Management grouping were considered highest priority and most fundamental: supervisory practices and risk assessment and prioritization. The Team was careful to illustrate each problem statement with concrete examples of events and interview quotations, and to link the problem statements back to the original symptomatic events.

Each problem statement was then turned over to an action group comprising one or more "customer," "suppliers," executive sponsors, and team member change agents. The roles of customer and supplier, a relatively new terminology for this organization, are products of its quality improvement process. However, the problems in the Management grouping were *not* turned over to an action group; instead, management collectively took responsibility for these issues. Although this was described as a legitimate strategy to force management "buy-in" and ownership, it is also a way for management to close ranks and prevent other employees from exerting control. Subsequently, this initiative dissipated in the face of new problems and wholesale changes in top management.

CONCLUSIONS

This paper has described some learning activities, specifically the self-analysis of operating problems, in the nuclear power and chemical process industries. The analyses and examples illustrate that the logic of learning is fundamentally different from the logic of

fixing, and requires both discipline and support. People are rarely encouraged to look at the linkages among problems and issues, because it creates a feeling of overwhelming complexity and hopelessness. Especially when faced with pressure to do something quickly, it is difficult to argue for a slow and uncertain approach. Yet, given encouragement, resources, and helpful conceptual lenses, it is possible to consider multiple issues rather than a single root cause, to look beyond individuals who made errors to the organizations and systems that set them up for errors, and to raise issues about accepted practices and high status groups that are usually too powerful to be challenged. When that is done, we must also recognize that change will take time, commitment, and sustained resources. That is why it is so instructive to participate with companies that are trying to take a deeper look at themselves and to initiate lasting change, even though it is a lengthy and challenging process.

NOTES

*Many of the ideas in this paper came directly or indirectly from Dr. Constance Perin during our collaborative research. A version of this paper was presented under the title “Failures to Learn from Experience: An Analysis of Incident Reviews in Nuclear Power Plants and Chemical Process Plants” in the Symposium “High Stakes Learning: Making Sense of Unusual, High-Hazard Events” at the Academy of Management meetings, Cincinnati, August, 1996.

[1] The MIT International Program for Enhanced Nuclear Power Plant Safety was supported by a consortium of utilities, suppliers, vendors, and foundations from the U.S. and other countries, who provided financial resources, access to sites and staff, and collaboration.

REFERENCES

- Argyris, C., and Schon, D. (1996). Organizational Learning II: Theory, Method, and Practice. Reading Ma: Addison-Wesley.
- 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. and Perin, C. (1993). Organization and Management of Nuclear Power Plants for Safe Performance: 1993 Annual Report. Cambridge, MA: MIT Sloan School of Management.
- Carroll, J. S., Sterman, J. D., and Marcus, A. A. (in press). 'Playing the manufacturing game: How mental models drive organizational decisions'. In Stern, R. N. and Halpern, J. J. (Eds.), Debating Rationality: Nonrational Aspects of Organizational Decision Making. Ithaca, NY: Cornell University ILR Press.
- Chiba, M. (1991). 'Safety and reliability: A case study of operating Ikata nuclear power plant'. Journal of Engineering and Technology Management, 8, 267-278.
- Cohen, M. D. and Sproull, L. S. (Eds.) (1996). Organizational Learning. Thousand Oaks, CA: Sage.
- Daft, R. L. and Weick, K. E. (1984). 'Toward a model of organizations as interpretation systems'. Academy of Management Review, 9, 284-295.
- Dougherty, D. (1992). 'Interpretive barriers to successful product innovation in large firms'. Organization Science, 3, 179-202.

- Jasper, J. M. (1990). Nuclear Politics: Energy and the State in the United States, Sweden, and France. Princeton, NJ: Princeton U.
- Kemeny, J. G. et al. (1979). Report of the President's Commission on the Accident at Three Mile Island. New York: Pergamon.
- 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.
- Lanir, Z. (1986). Fundamental Surprise. Eugene, OR: Decision Research.
- LaPorte, T. R., and Consolini, P. (1991). 'Working in practice but not in theory: Theoretical challenges of high reliability organizations'. Journal of Public Administration Research and Theory, 1, 19-47.
- Levitt, B. and March, J. G. (1988). 'Organizational learning'. Annual Review of Sociology, 14, 319-340.
- March, J. G., Sproull, L. S. and 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.
- Penrose, E. (1968). The Theory of the Growth of the Firm, 4th ed. Oxford: Basil Blackwood.

- Perrow, C. (1984). Normal Accidents: Living with High Risk Systems. New York: Basic Books.
- Rasmussen, J. (1990). 'The role of error in organizing behavior'. Ergonomics, 33, 1185-1190.
- Reason, J. (1990). Human Error. New York: Cambridge U.
- Rees, J. V. (1994). Hostages to Each Other: The Transformation of Nuclear Safety Since Three Mile Island. Chicago: U. of Chicago.
- Roberts, K. H. (1990). 'Some characteristics of one type of high reliability organization'. Organization Science, 1, 160-176.
- Rochlin, G. I. and von Meier, A. (1994). 'Nuclear power operations: A cross-cultural perspective'. Annual Review of Energy and the Environment, 19, 153-187.
- Rogovin, M. and Frampton, G. T. Jr. (1980). Three Mile Island: A Report to the Commissioners and to the Public. Washington, D.C.: U.S. Nuclear Regulatory Commission.
- Senge, P. (1990). The Fifth Discipline: The Art and Practice of the Learning Organization. New York: Doubleday.
- Senge, P., and Sterman, J. D. (1991). 'Systems thinking and organizational learning: Acting locally and thinking globally in the organization of the future'. In Kochan, T. and Useem, M. (Eds.), Transforming Organizations. Oxford: Oxford University Press, pp. 353-370.
- Schein, E. H. (1987). Process Consultation, Vol. II: Lessons for Managers and Consultants. Reading, MA: Addison-Wesley.

- Schein, E. H. (1992). Organizational Culture and Leadership, 2nd ed. San Francisco: Jossey-Bass.
- Schein, E. H. (in press). 'The three cultures of management: Implications for organizational learning.' Sloan Management Review.
- Sitkin, S. (1992). 'Learning through failure: The strategy of small losses'. Research in Organizational Behavior, 14, 231-266.
- Slovic, P. (1987). 'Perception of risk'. Science, 236, 280-285.
- Tamuz, M. (1994). 'Developing organizational safety information systems for monitoring potential dangers'. In Apostolakis, G. E. and Wu, J. S.(Eds.), Proceedings of PSAM II, vol. 2. Los Angeles: University of California, pp. 71:7-12.
- Tetlock, P. (1985). 'Accountability: The neglected social context of judgment and choice'. In Staw, B. and Cummings, L.(Eds.), Research in Organizational Behavior, Vol. 1. Greenwich, CT: JAI Press.
- Turner, B. A. (1978). Man-Made Disaster. London: Wykeham.
- Weart, S. R. (1988). Nuclear Fear: The History of an Image. Cambridge, MA: Harvard University Press.
- 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.

Figure 1

Two Images of Organization

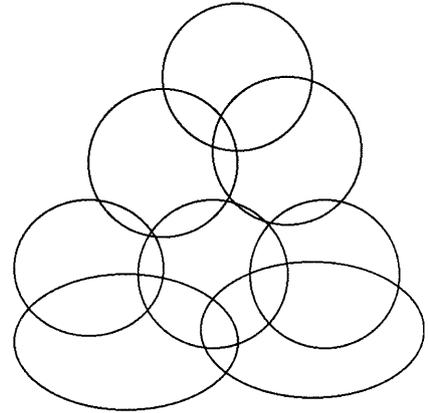
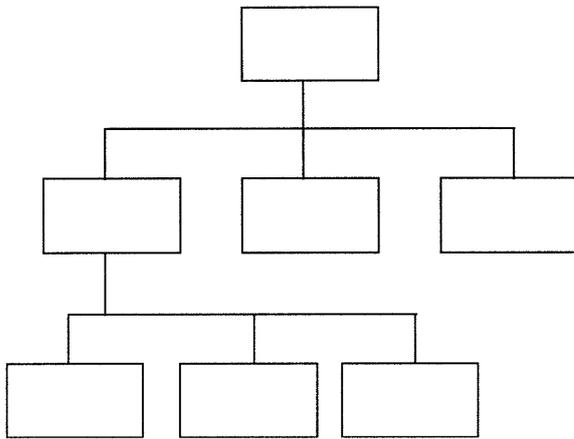


Figure 2

Organizational Learning Framework

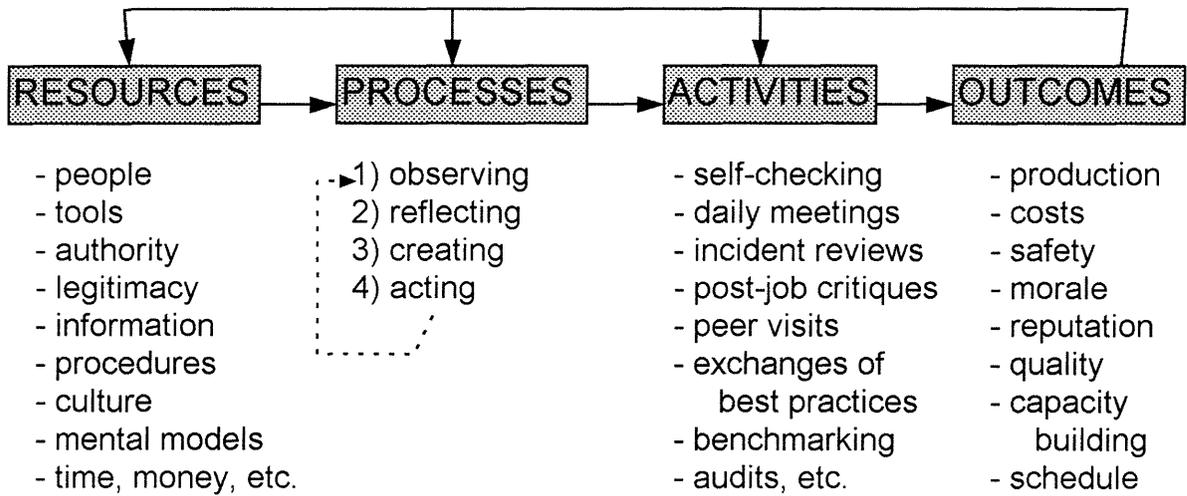


Figure 3

Fixes That Fail

