Chapter 15

FUNDAMENTAL SURPRISE

> An apocryphal story about the well-known lexicographer, Noah Webster:
>
> "One day, he arrived home unexpectedly to find his wife in the arms of his servant. 'You surprised me,' said his wife. 'And you have astonished me,' responded Webster."
>
> from Lanir (1986)

Notice Noah Webster's position: the surprising event calls into questions basic assumptions about his relationship to others. The event's surprise is fundamental because it opens a window for learning for a period of time. What Noah learns in that interval can vary. He may re-examine and modify his concepts about his world; he may learn new ways to act relative to others.

Some accidents, such as the destruction of the space shuttle *Challenger*, have this same kind of effect for stakeholders. These events question our model of how the system works, how it usually succeeds, and how it is vulnerable to failure. Celebrated accidents like the Three Mile Island reactor accident and political events such as the launch of Sputnik for the U.S., and the Yom Kippur War for Israel have this character. This apocryphal story about Noah Webster and these celebrated accidents illustrate a basic pattern in the reactions across organizations to failure first described by Lanir (1986) and termed fundamental surprise.

Before the fact, such events or accidents are largely inconceivable to the relevant operational, design and regulatory communities. As a result,

Wagenaar and Groeneweg (1987) and Lanir (1986) have termed these accidents as "impossible", in the sense that the event falls outside of the previous languages of description used in that setting. It is when a set of people in the various stakeholder communities realizes that an event falls outside the standard view that the fundamental surprise sequence begins. The challenge of fundamental surprise is to use them as sources of information for expanding the language of description of

- inherent hazards,
- paths to failure,
- forms of expertise in context, and
- the basis for success.

On the other hand, there is no crisis for Noah's wife – she can explain away the anomaly and simply tune her early warning mechanisms for tracking Noah. And Noah may come to view the surprising event and to resolve the dissonance created by the event in the same way – narrowing the lessons until the surprise is resolved through limited re-tuning of local factors.

Fundamental surprise is a pattern in the reactions to a failure that moves through three stages.

> *1. Opening up to Learning – Fundamental Surprise*
> First, the surprise event or events occur. It is not simply an anomalous event, costly failure, or tragedy. The event challenges previous views, demands an explanation, and leads participants and stakeholders to question basic assumptions. In this sense stakeholders are in crisis.
>
> Dissonance between belief and experience provides the opportunity to learn about safety for a period of time. The possibility for learning occurs because people are willing to question assumptions and current directions and they are willing to spend energy to search for how to explain the surprise.
>
> *2. New Views and Concepts -- Learning How to Learn*
> The fundamentally surprising event opens a window on how the system in question actually succeeds and how it is vulnerable to failure. Through this window we can look to see the internal workings of the system that produced the accident.

New concepts are imported or arise to explain the inherent hazards, paths to failure, forms of expertise in context, and the basis for success. But seeing the how the system works in new ways is very likely to challenge what had seemed intuitive or obvious previously.

Exploiting the opportunity is itself difficult and demanding. For example, as learning begins and starts to grow, countervailing forces also arise which lead stakeholders to resolve the dissonance by repairing and restoring the traditional viewpoints and assumptions. In this process people try to turn the fundamental surprise into a narrow anomaly that requires little or no fundamental change in view or change to future operation.

3. Resolution and Consolidation

Fundamental surprise is a dissonant state that calls for resolution and re-establishing a new equilibrium. Some changes are absorbed and the organization or industry settles back into a new equilibrium as it balances different subgoals related to production and to safety. The questions are: what has been learned and acted on? Are they productive, long term approaches to improve the ability to create safety?

An Example of Fundamental Surprise: The Three Mile Island Nuclear Power Accident

On March 28, 1979, the U.S. nuclear industry and technologists were rocked by the Three Mile Island accident (TMI). The consternation that resulted was due to more than the fact that it was the worst nuclear accident up to that time or the radiological consequences per se. The TMI accident was more than an unexpected progression of faults; it was more than a situation planned for but handled inadequately; it was more than a situation whose plan had proved inadequate.

The TMI accident constituted a fundamental surprise in that it revealed a basic incompatibility between the nuclear power industry's view of itself and reality. Prior to TMI, the industry could and did think of nuclear power as a purely technical system where all problems

were in the form of some engineering technical area or areas, and the solutions to these problems lay in those engineering disciplines. TMI graphically revealed the inadequacy of that world view because the failures were in the socio-technical system and not due to pure technical factors (a single equipment or mechanical flaw) or to a purely human failure (gross incompetence or deliberate failures).

Prior to TMI, the pre-planning for emergencies consisted of considering large equipment failures; however, it did not consider a compounding series of small failures interacting with inappropriate human assessments of the situation and therefore erroneous actions. Prior to TMI, risk analysis also focused on large machine failures, not on the concatenation of several small failures, both machine and human. The kind of interaction between human and technical factors that actually occurred was inconceivable to the nuclear industry as a whole prior to TMI.

The post-TMI nuclear industry struggled to cope with, and adjust to, the revelations of TMI. The process of adjustment involved the phases associated with fundamental surprise described by Lanir. First, the surprise event itself occurs. Reactions spilled over the boundaries of the event itself to include issues that had little to do with the triggering event—crises. Second, these crises provide the opportunity for learning about vulnerabilities and innovating new directions. Finally, some of these changes are absorbed and a new equilibrium is reached.

The immediate investigations of the TMI accident focused heavily on the mutual interaction between technical systems and people. The proposed changes that resulted from these investigations addressed the basic character of the joint human-machine system. These included providing new kinds of representations of the state of the plant, restructuring the guidance for board operators on how to handle abnormal conditions, and restructuring the organization of people in various facilities and their roles in handling different problems created by accidents.

However, in the process of carrying through on these and other “lessons learned” the U.S. nuclear industry shifted direction and treated the accident as if it was nothing more than a narrow anomaly. They began to focus on localized and purely technological solutions—what could be termed the fundamental surprise error, after

Lanir's analysis (cf., Reason, 1990).[36] This occured despite the fact that the revelations of TMI continued to re-occur in other major incidents in the U.S. nuclear industry (e.g., the Davis-Besse nuclear power plant incident, see US NRC, 1985) as well as in other risky technological worlds. While the post-TMI changes clearly improved aspects of the socio-technical system through such things as new sensors, new analyses of possible accident conditions, new guidance on how to respond to certain accident conditions, and changes in emergency notification procedures, the basic socio-technical system for operating plants and responding to failures did not change (Moray and Huey, 1988).

Learning from Fundamental Surprise

One of the themes reverberating throughout this book is that human error represents a symptom rather than a cause. In this view system failure, or anticipations of system failure, are, in part, a form of information about the functioning of the system in which people are embedded (Rasmussen, 1986). Organizational learning depends on using the information in these events to go behind the label human error and to see ways improve human-machine and socio-technical systems of work.

Despite terrible consequences, fundamentally surprising events offer an opportunity for learning and change because the event disrupts the usual internal political alignments of organizations. There is a profound sense among many that the usual concepts and policies are insufficient to cope with what has happened. Whether this sense of surprise can be translated into learning and effective change is, of course, the central issue.

It is clear that the immediate aftermath of a serious failure produces an atmosphere of inquiry and frees up resources normally dedicated to production can be brought to bear on the accident and its consequences. The lines that divide participants, management, regulators, and

---

[36] In general, the "fundamental surprise error" is re-interpreting a fundamental surprise as merely a situational surprise which then requires no response or only a limited response. For the context of complex system failures, research results indicate that a specific version of this error is re-interpreting a human-machine system breakdown as being due to purely human factors.

victims from each other momentarily thin. As one National Transportation Safety Board observer put it, "when the vividness of the tragedy is fresh in everyone's mind, and the broken wreckage is still smoldering...people....have only one pressing goal, and that is to determine precisely what happened and see that it does not happen again." (NTSB Report RP-84/01, Accident Investigation Symposium, Springfield, Virginia April 26-28, 1983, page 8). This period of cooperation and focus makes it possible to ask questions that are not usually asked, gather data not usually gathered, and inquire into issues not usually open to inquiry.

Not all stakeholders are in the same position relative to the surprising event. Some are closer; others more distant in knowledge, point of view, and in experiencing the consequences. Distance from the work context where an accident occurs appears to alter what is learned from incidents and accidents. Those people who are at the epicenter of high consequence accidents are usually devastated and entirely caught up in the consequences and reactions to failure. Conversely, people who are far distant from the epicenter are too divorced from the complex context of technical work to learn anything meaningful from the sterile, skeletal descriptions of the events that reach them.

But near the epicenter are people who have both a detailed understanding of the context of work yet are sufficiently distant from the consequences. This mixture seems to be important in producing fundamental surprise. Because their technical work corresponds closely to conditions of work at the epicenter, the event has direct relevance. Because they are some distance from the epicenter, their attention is not captured by the need to react to the event itself and they have an opportunity to extract useful meaning, i.e. to learn about safety.

Avoiding Learning

When an organization experiences an accident, the real, tangible and tragic consequences throw up so much debris that learning is difficult. Forms of debris include:

- negative consequences are emotional and distressing for all concerned,

- failure generates pressure from different stakeholders to resolve the situation,
- a clear understandable cause and fix helps stakeholders move on from a tragedy, especially when they continue to use or participate in that system,
- managing financial responsibility for ameliorating the consequences and losses from the failure,
- desire for retribution from some stakeholders and processes of defense against punitive actions,
- confronting dissonance and changing concepts and ways of acting is painful and costly in non-economic senses.

It is equally clear that any period of cooperation and openness may be short lived. As litigation and regulatory issues begin to loom and as the freshness of tragedy recedes, the dividing lines across stakeholders re-thicken and become distinct. Views about what the accident means are re-shuffled in ways that are increasingly self serving.

The dissonance of a fundamental surprise can also be resolved by repairing and restoring the traditional viewpoints and assumptions. The fundamental surprise is denied by narrowing the attribution of "cause" to local factors with well bounded responses. We call this response the "fundamental surprise fallacy." Attributing the surprise to narrow local factors absolves the organization or industry of any need to confront difficult and costly changes.

Typically the fundamental surprise fallacy takes the form of attributing the source of the surprising event to others who are not as well motivated, knowledgeable, or careful as we are – this couldn't happen to me. In effect, we transfer the issue and the difficulties to other parts of the system.

The label "human error" is a good example of a narrow interpretation that avoids confronting the challenges raised by a fundamentally surprising event. If the source of the incident is "human error," then only local responses are needed which do not change the larger organization or system. Curing "human error" in this "local" sense only requires sanctions against the individuals involved, injunctions to try harder or follow the procedures more carefully, or some remedial adjustments in the training programs. Stopping with the label "human error" is a defense mechanism that safeguards the person or

organization from the dissonance of the fundamental surprise, precludes learning, and prevents constructive forward movement.

For the technologist, the conclusion "human error", i.e., that the people in the system are an unreliable component, is a very comfortable notion. This leads to the idea that "just a little more technology will be enough" (Woods, 1991)—that purely technological responses without consideration of human-machine systems or larger organizational factors can produce high reliability organizations. As we have seen in Part III this is a fallacy. Rationalizations like this block the opportunity to learn from the fundamental surprise.

The characteristic signature of accidents in complex systems – multiple latent contributors -- confound our ability to learn from them. Singling out a single cause from among the many factors playing roles in accidents is a social and psychological process rather than a technical, objective one, as we have seen in Chapter 13. Contributions nearby in time and space always receive intense scrutiny. Narrowing in on single factors simplifies the story of the event and blocks learning about the other contributors and how they come together to produce the conditions for failure.

The availability of multiple contributors allows different stakeholders to single out as the critical factor some contributor which falls within the scope of authority of other people, groups, or organizations. Therefore, it is those other people who need to change or to pay the costs of change.

Because accidents are unique on the surface, even though they exhibit more general patterns, it is easy for stakeholders to see the event as simply unique. If it is unique, there is little need for change, or at least only local fine tuning steps are needed to address those exact conditions. Emphasizing the differences between the conditions of any given accident and current conditions also will allow a stakeholder to rationalize -- this can't happen now.

All of the above factors produce a strong incentive to stop at first stories and avoid the work needed to uncover second stories. Most organizations treat accidents as isolated, unique, unfortunate but limited events that have limited significance.

The scope of operations and the speed and efficiency of large complex systems is so great that overt failure may involve enormous amounts of damage or damage that is inherently difficult to repair Because this is so apparent, it is incumbent on all organizations involved with such systems to be obviously, overtly, and publicly committed to safety. Even a cursory review of the landscape of aviation, manufacturing, or medicine shows that every organization is critically concerned with maintaining its obvious commitment to safety.[37]

Because of the scrutiny that follows accidents, posturing about safety is especially pronounced after accidents occur. It is organizationally and institutionally satisfying to blame isolated human operator error for accidents because it limits the need for change. So long as the locus of failure is the human practitioner (and, it must be said, so long as that practitioners expertise in context remains a black box into which we cannot see) the occurrence of accidents says nothing about the system in which the operator functions. This means that the organization can continue operate as before, avoid the disturbing questions, and comfortably pursue its goals of production or efficiency. At most, the conclusion that human error is the cause of failure leads to ratcheting up the constraints on operators.

This produces a paradoxical effect: lodging the root cause of failure in isolated human operators makes it impossible to discern patterns of latent failures that lead to accidents, but it does satisfy the need for 'cover' for organizations. When failure is regarded as a property of human operators rather than larger systems, learning from accidents becomes nearly impossible.

Extending or Enhancing the Learning Opportunity

An important question for researchers is a better understanding of how the window of opportunity for learning can be extended or enhanced. We need new information and techniques relevant to the following:

---

[37] It is instructive to think of how many times a hospital or airline company has been willing to state publicly that safety is only one goal of many or that safety must sometimes be compromised in order to achieve other desired goals such as economies of operation or rates of production. None of the authors can recall any such public announcement by a responsible representative of any organization in any field at any time in the past two decades.

- What makes an event fundamentally surprising rather than merely anomalous?
- Can we influence when these windows open, for whom and for how long?
- Can we capture and re-use events to create the window of opportunity for learning without waiting for a system, group, organization, or industry to experience an accident?
- Is the tendency to re-interpret the surprise as a narrow anomaly and to re-establish conventional views inevitable?
- Can knowledge of how complex systems fail and knowledge of patterns in reactions to failure be used to accelerate the learning process?

Forms of Lessons

What does it mean to treat incidents, accidents, and disasters as information? What aspects of world views about hazard, paths to failure and sources of expertise in context need to be re-examined following a surprising event?

What is fundamentally surprising about the event is that it challenges conventional views of how safety is achieved:

- One must look for sources of success and failure in the larger system, the socio-technical system and the human-machine system, not in single factors.
- Fundamentally surprising accidents continue to point to the interaction of people, technology, and the larger organization in which practitioners at the sharp end are embedded.
- Understanding how people create safety and how these processes break down, instead of attributing failure to unreliable human component.
- Stakeholders must confront how hindsight biases their review of past events and blocks their ability to accurately gauge how the system works.

There are a few simple lessons to keep in mind as one pursues second stories. One, there are always multiple lessons to be learned because there are always multiple contributors. The multiple contributors do

not form a simply chain, rather there are always parallel factors at different levels of analysis with implications at different levels of the system in question. The response should not be remedial but systems oriented.

Two, there is a need for continuity of response. The typical response is to reach closure and close the books on the case. Instead, continuity implies an ongoing investment and building up an investment pool to support and sustain multiple responses (avoid one shot remedies). Continuity demands saving and broadcasting cases and the underlying patterns they represent to reinforce a culture of information flow and awareness. Continuity means monitoring of investment and change relative to the purposes they were intended to achieve – Are the changes effective? Are there side effects of the change? Is follow through needed?

Third, the breadth of learning and response should extend well beyond the failed device, person, unit, or organization to examine how the case contains patterns at some level of analysis that is relevant to all in whatever role or organization (avoid distancing through differencing -- they are different on dimension x, therefore it doesn't have lessons for me).

Learning How to Learn about Safety

In one or another role, we have observed motivated, intelligent individuals and organizations inside nuclear power, aviation, and healthcare grappling with safety, accidents, research, regulations, technology and tradition. A few generalization capture what we have seen about the process of learning from incidents or accidents.

1. *Learning requires dissonance between belief and experiences.*
So long as beliefs (about how the system works, about the roles of people, about the ways in which people work, about technical work, etc.) are consistent with experience, learning is limited. When events occur that challenge existing beliefs (e.g. "impossible" accidents, unanticipated bad side effects of "safety neutral" organizational change), the disturbance in the [manager's, regulator's, operator's and designer's] structure of belief opens opportunities for fundamental learning.

2. *Learning about safety is not continuous but occurs at intervals.*
Safety comes to the foreground only at certain moments, usually after accidents, in connection with internal strife, in response to new challenges (e.g. production pressures), or associated with moments of organizational or technological change.

3. *Not everyone learns at the same time.*
Especially in large systems, learning occurs differently in different parts of the organization. A corollary is that some parts of organizations may learn quickly, especially if there is a perceived relationship between the need to learn about safety and pressures on their current, local work.

4. *Learning about safety is not permanent.*
What is learned can be forgotten. Knowledge can become stale, inert, and inaccessible. Continuing experience at the sharp end is valuable to prevent forgetting.

5. *Learning about safety requires close contact with failure and also the distance needed for reflection.*
Learning about safety appears to require close, "hands on" experience with failure. (It is interesting that most operator cultures value such experience and seek to make certain that apprenticeship training involves such experience.) But learning about safety is not simply a matter of experience. It also requires the opportunity to step back and reflect on the larger lessons that are exemplified by experience, to see connections, and to interpret. This reflection requires guidance and contact with the research results.

6. *Learning inherently involves exploring the "second stories" that lie behind accidents and failure.*
The widely known, easily told "first stories" of accidents are always incomplete and misleading in ways that make them impossible to use for progress on safety. The deeper, more detailed stories that characterize the way failures arise from multiple contributing factors, competing demands, process complexity, usual success, powerful adaptations, and manifold hazard are the basis for progress on safety. Technical work and technical failure are intimately bound up with complexity and conflicting demands of the domain of practice and these are the subject of the productive "second stories" that can be (but seldom are) discovered and told.

7. L*earning recapitulates the sequence of research that comprises the* NEW LOOK.
Certain components of the NEW LOOK are essential groundwork for others. Although there are variations, the primary sequence of learning that we observe is one where individual and organizational learning retraces the history of the research that makes up the NEW LOOK. It is also clear that people get sidetracked along the learning path at various places for a variety of reasons.

8. *Learning about safety exposes organizational stress.*
Virtually all components of the NEW LOOK expose the underlying rifts, disagreements, and mixed character of the organization itself. This threatens some and makes many uncomfortable. It exposes tensions between different parts and levels of the organization or across organizations in an industry. But these conflicts are central to learning about safety because they are the sources of conflicting demands and resource limitations that constrain practitioners' and management's ability to create safety.

9. *Learning about safety begins with learning that people create safety*.
People are constantly creating safety, at all levels. People are the adaptive element of systems. At the sharp end, the story of creating safety is mainly the story of adaptation to cope with complexity. Practitioner's strategies are failure sensitive; people are aware of the potential for failure in their world. They are also aware of many of the mechanisms that lead to failure and constantly devising strategies to forestall failure. Investments that extend the robustness of these strategies are characteristic of successful organizational. At the blunt end, the story of safety is mainly about how to shape the world so that the sharp end adaptations are effective.