12-2008

Why Catastrophic Organizational Failures Happen

Michael A. Roberto
Bryant University

Why Catastrophic Organizational Failures Happen

Michael A. Roberto

Associate Professor, Trustee Professor of Management

Bryant University

Faculty Suite F-414

1150 Douglas Pike

Smithfield, RI 02917

Email: mroberto@bryant.edu
Phone: 401-232-6155

Introduction

On January 16, 2003, the *Columbia* space shuttle set out on its twenty-eighth flight into space, in what was characterized as a routine scientific mission. Two weeks later, the shuttle disintegrated while re-entering the earth's atmosphere, killing the seven astronauts on board. Soon thereafter, the Columbia Accident Investigation Board (CAIB) – a star-studded panel of experts on the safety of complex, high-risk systems – began trying to understand what had gone so terribly wrong. The investigative board determined the technical cause of the accident within several months. A piece of insulating foam had dislodged from the external tank of the shuttle during launch. That foam debris struck the leading edge of the vehicle's wing, puncturing a hole in it. During the shuttle's return to the earth's atmosphere at the end of the mission, extremely hot gases entered the interior of the wing through that hole, melting the structure from the inside out. That melting caused the break-up of the vehicle.

The investigators did not simply conduct a technical analysis of this catastrophic failure. They went on to evaluate the organizational systems, processes, and behaviors that enabled the tragedy to occur. They wanted to understand why NASA kept launching the shuttle despite a lengthy history of foam strike problems. They sought to determine why management had concluded that the astronauts were safe, despite some engineers' serious concerns about the foam strike. The board members noticed organizational problems similar to those uncovered during 1986 Challenger accident investigation, and they wondered why NASA had not corrected those problems in subsequent years. In the words of CAIB member Duane Deal, a retired Air Force General, the investigators sought to go "beyond the widget" during their analysis. They wanted

to understand the organizational causes of the catastrophic failure, not simply the technical cause (i.e., the widget that broke). General Deal summarized the board's findings: "The foam did it. The organization allowed it." (Deal, 2004: 44)

As the board conducted its investigation, many management scholars turned their attention to the *Columbia* accident as well. These researchers typically wanted to go "beyond the widget" to understand the human and organizational conditions that led to this tragedy. The scholars that studied the *Columbia* accident followed in a long tradition of research into catastrophic failures. Prior studies have examined incidents such as the 1977 Tenerife airliner collision, 1979 Three Mile Island nuclear power plant accident, 1986 *Challenger* explosion, and the 1994 friendly fire incident in the Iraqi No-Fly Zone. Why should management researchers be interested in these types of complex and unusual tragedies? According to William Starbuck and Moshe Farjoun, who co-edited a book about the *Columbia* accident, catastrophic failures "dramatize how things can go wrong, particularly in large, complex social systems, and so they afford opportunities for reflection, learning, and improvement." (Starbuck and Farjoun, 2005: 3). Moreover, as John Carroll has argued, these events prove to be fertile ground for the development of new organizational theory, particularly since the subsequent investigations provide rich descriptions of the inner workings – and failings – of individuals, groups, and organizations. These thick, rich descriptions of organizational systems and behaviors provide the data required for provocative new theory-building.

The purpose of this chapter is to examine the major streams of research about catastrophic failures, describing what we have learned about why these failures occur as well as how they can be prevented. The chapter begins by describing the most prominent sociological school of thought with regard to catastrophic failures, namely normal accident theory. That body

4

of thought examines the structure of organizational systems that are most susceptible to catastrophic failures. Then, we turn to several behavioral perspectives on catastrophic failures, assessing a stream of research that has attempted to understand the cognitive, group and organizational processes that develop and unfold over time, leading ultimately to a catastrophic failure. For an understanding of how to prevent such failures, we then assess the literature on high reliability organizations (HRO).   These scholars have examined why some complex organizations operating in extremely hazardous conditions manage to remain nearly error free. The chapter closes by assessing how scholars are trying to extend the HRO literature to develop more extensive prescriptions for managers trying to avoid catastrophic failures.

A Structural Perspective

Research on catastrophic failures traces its roots to a groundbreaking study of the Three Mile Island nuclear power plant accident and the development of normal accident theory. In his 1984 book, Normal Accidents, Charles Perrow examined the structural characteristics of organizational systems that involve high-risk technologies such as nuclear power. Perrow's conceptual framework classifies all high-risk systems along two dimensions: interactive complexity and coupling. Interactions within a system may be simple/linear or complex/nonlinear. Coupling may be either loose or tight. Perrow argues that systems with high levels of interactive complexity and tight coupling are especially vulnerable to catastrophic failures. In fact, he argues that accidents are inevitable in these situations; certain failures constitute "normal accidents." Perrow concludes that, "Normal accidents emerge from the characteristics of the systems themselves. They cannot be prevented." (Perrow, 1981: 17)

Interactive complexity refers to the extent to which different elements of a system interact in ways that are unexpected and difficult to perceive or comprehend. Often, these interactions among elements of the system are not entirely visible to the people working in the organization. Simple, linear interactions characterize systems such as a basic manufacturing assembly line. In that instance, the failure of a particular piece of equipment typically has a direct, visible impact on the next station along the line. The operations of a nuclear power plant do not follow a simple linear process; instead, they are characterized by complex and nonlinear interactions among various subsystems. The failure of one component can have multiple, unanticipated effects on various subsystems, making it difficult for an operator to diagnose the symptoms of a developing catastrophe.

Tight coupling exists if different elements of an organizational system are highly interdependent and closely linked to one another, such that a change in one area quickly triggers changes in other aspects of the system. Tightly coupled systems have four attributes: time-dependent processes, a fairly rigid sequence of activities, one dominant path to achieving the goal, and very little slack. When such rigidity exists within an organization, with few buffers or slack among the various parts, small problems can cascade quickly throughout the system leading to catastrophe. Loose coupling exists when subsystems are not as tightly integrated, such that small errors in one area can be isolated or absorbed without affecting other subsystems.

Engineers naturally try to account for the fact that components or subsystems might fail, or human error might occur, in a high-risk technological system. They design back-up systems to deal with those contingencies. However, engineers encounter difficulty when a strong possibility exists for multiple failures to occur in a nearly simultaneous fashion. When interactive complexity and tight coupling exist, then the possibility exists for a series of

unanticipated, interconnected breakdowns that can quickly build upon one another. In short, normal accidents often do not have a single cause; they involve a chain of multiple failures. That chain often proves difficult to detect and break before catastrophe strikes.

Psychologist James Reason (1997) also argues that catastrophic accidents typically involve a chain of failures rather than a single cause. In his famous "Swiss cheese" analogy, an organization's layers of defense or protection against accidents are described as slices of cheese, with the holes in the block of cheese representing the weaknesses in those defenses. In most instances, the holes in a block of Swiss cheese do not line up perfectly, such that one could look through a hole on one side and see through to the other side. In other words, a small error may occur, but one of the layers of defense catches it before it cascades throughout the system. However, in some cases, the holes become completely aligned, such that an error can traverse the block, i.e. cascade quickly through the organizational system.

Critique of Normal Accident Theory

Normal accident theory has certain limitations. First, many scholars and practitioners find the theory frustrating, in that it does not move us toward an understanding of how to prevent catastrophic accidents. It appears to have little prescriptive value. They point out that we should not simply resign ourselves to the inevitability of failure in certain situations. Instead, we should try to develop an understanding of how we might reduce the probability of normal accidents. To begin, scholars point out that we can design safer systems by addressing the very attributes defined by Perrow. By trying to reduce interactive complexity and tight coupling, engineers and managers can begin to construct safer systems. We will discuss other remedies for preventing catastrophic failures when we address the HRO literature in a subsequent section. For now, it

should simply be pointed out that the HRO scholars have tried to identify the behaviors, norms, and processes required to raise people's awareness of interdependencies and interaction effects, as well as to catch small errors before they cascade throughout a system.

A second major criticism refers to the problems inherent in the classification scheme itself. The two system dimensions articulated by Perrow are useful in helping us understand the vulnerability of organizations. However, one cannot easily classify organizations in his 2x2 matrix. For instance, many scholars have characterized commercial aviation as a complex system prone to normal accidents. Marais, Dulac, and Leveson (2004) point out the flaws in this argument. They argue that the U.S. air traffic control infrastructure has been deliberately designed as loosely coupled system. Various measures exist to ensure that problems in one sector of airspace or on one particular flight are unlikely to have an impact on the safety of other flights. Similarly, Snook (2000) examines the 1994 accidental shootdown of two U.S. Blackhawk helicopters by U.S. fighter jets in northern Iraq. Based on Perrow's earlier discussion of military and aviation systems, he cannot discern whether the incident should be characterized as a tightly or loosely coupled system. Snook goes on to argue that complex organizational systems are fundamentally dynamic entities, and coupling may not be a static property of such systems. Drawing on Weick's (1976) earlier work, he makes the case that the organization responsible for maintaining a no-fly zone over northern Iraq shifted between loose and tight coupling over time. He argues that coupling may be situational, rather than an enduring property of organizations. A system might be loosely coupled most of the time, but then in certain relatively rare situations, it becomes rigidly interdependent and interconnected. In those instances, the likelihood a catastrophic accident escalates dramatically. The critical question, then, is not whether a system exhibits tight or loose coupling, but instead how and when does a

system migrate from loose to tight coupling, and what can be done to prevent that migration from occurring.

Behavioral Perspectives

In contrast to normal accident theory, another stream of researchers has studied the cognitive, group and organizational processes that lead to the failure to detect errors and/or the failure to address errors before they lead to a catastrophic failure. These researchers have focused much more attention on behavior, rather than solely focusing on the structural dimensions of organizational systems. They also have focused much more on the historical evolution of catastrophic accidents. They do not only examine the momentous decision that might have immediately preceded a tragedy (such as the critical eve-of-launch meeting that took place prior to the *Challenger* accident), nor do they focus exclusively on the immediate chain of events that led to a failure (such as the breakdowns in communication that took place on the morning of April 14, 1994 in northern Iraq). Instead, these scholars examine the gradual development of norms, beliefs, and attitudes that contribute to unsafe action. They try to understand the history of decisions that were made in an organization, as well as the danger signs that may have been downplayed or misinterpreted over time. This approach to studying failures can be traced back to Barry Turner's seminal book, Man-Made Disasters, published in 1978. He argued that many catastrophic failures were characterized by incubation periods that stretched over many years, not days or hours. He made a strong case for how catastrophic accidents were *processes, not events*. They did not simply happen at a point in time; they unfolded in a gradual accumulation of actions, decisions, and interpretations by many actors within an organization. As such, catastrophic failures required intensive longitudinal study. Since Turner's work, many

behavioral studies have been done, but a few major developments in the field should be highlighted here, not simply because they help us understand better why catastrophic failures occur, but also because they provide important guidance as to how we should study failures moving forward.

From Proximate Choices to Patterns of Decisions

After the *Challenger* accident in 1986, many scholars focused on the decision to launch in cold weather on a January morning, and in particular, on the critical meeting that took place on the eve of the launch. During that infamous midnight teleconference, some engineers argued that the unprecedented cold temperatures expected for the following morning could lead to O-ring failure during launch. They pushed for a postponement of the mission. Scholars studied this meeting quite carefully, examining the communication failures that took place among the group members. Many people described the *Challenger* launch decision as a classic case of *groupthink* – a term that social psychologist Irving Janis (1982) coined to describe the pressures for conformity, ardent striving for unanimity, and premature convergence that take place within many cohesive groups (Moorhead, Ference, and Neck, 1991; Esser and Lindoerfer, 1989). A well-known producer of management training materials even created a best-selling film about groupthink which focused, in part, on the Challenger launch decision. Executives around the world participated in seminars, using this video, in which they learned how to avoid the groupthink pathology that characterized the eve-of-launch meeting at NASA in January 1986.

Ten years after the tragedy, sociologist Diane Vaughan (1996) offered a sharp rebuke to the conventional wisdom regarding the *Challenger* accident. In her widely acclaimed book, Vaughan argued that the many antecedent conditions of groupthink, as described by Janis, did

not exist. For instance, the group that assembled on the midnight teleconference hardly could be described as a small cohesive in-group. It consisted of 34 individuals, some of whom did not know one another. More importantly, however, Vaughan explained that one could not understand the causes of the catastrophe simply by studying only the specific decision, made at the eve-of launch meeting, to launch the Challenger on January 28, 1996 despite the unusually cold temperatures. She argued that one had to understand the *pattern of decisions* over time that caused NASA officials to view O-ring erosion as an acceptable risk.

Vaughan described the *normalization of deviance* which took place over many years. In that gradual, evolutionary process, engineers and managers moved down a dangerous slippery slope. At first, they did not expect or predict O-ring erosion on shuttle flights. However, when a small amount of erosion was discovered on an early, successful shuttle mission, engineers considered it an anomaly.   Then, it happened again. Gradually, the unexpected became the expected. O-ring erosion began to occur regularly. Engineers rationalized that sufficient redundancy existed to ensure no safety-of-flight risk.   Small deviations became taken-for-granted. Over time, however, deviations from the original specification/expectation grew. Engineers and managers expanded their view of what constituted acceptable risk.   As the years unfolded, Vaughan explained that the "unexpected became the expected became the accepted" (Roberto, Bohmer, and Edmondson, 2006). The launch decision, therefore, could only be understood in the context of this long pattern of decisions, during which a gradual normalization of deviance took place. In short, history matters a great deal. Decisions and catastrophic failures cannot be understood without examining their historical context. When the *Columbia* accident occurred, the investigative board worked closely with Dr. Vaughan, drawing on her expertise to examine the historical origins of that disaster. They once again found a disturbing pattern of

normalization of deviance. Moreover, they studied NASA's history to try to understand why the organization had not learned the lessons of *Challenger*.

From Decision-Making to Sense-Making

Karl Weick offered another important admonition to scholars of catastrophic failure, calling attention to whether or not decisions represented the appropriate unit of analysis for studies in this area. In his classic 1993 paper on the Mann Gulch wildfire disaster, Weick argued that a sensemaking perspective, rather than a decision-making lens, offered more insight as to why thirteen firefighters died in that tragic incident.

According to Weick, scholars need to re-think the classic linear view that people decide and then act when confronted with a particular situation. He made the case that individuals often act, and then try to make sense of that action. They engage in an effort to make meaning out of the reality around them. That rationalization process compels them to make certain kinds of choices. In short, individuals often act, and then decide, rather than the other way around.

At Mann Gulch, the firefighters initially viewed the situation as what they routinely termed "ten o'clock fires," meaning that they expected to contain the relatively small fire by ten o'clock the next morning. Some indications arose fairly quickly that this blaze would be more serious than the typical "ten o'clock fire." However, early actions on the part of the crew leader, Wag Dodge, as well as other crew members, seemed to confirm the initial assessment of the situation. For instance, Dodge ordered the crew to eat dinner when they landed, and then he later left the crew for a short period of time to eat his own supper. As members tried to make sense of those actions, they rationalized that the blaze must not be too serious – perhaps it was only a ten o'clock fire. Later, a prominent crewmember stopped to take photos of the fire, while the

men were digging the fire line. Again, how does one make sense of that action, despite some signals that the fire might be more dangerous than originally thought? Gradually, as the situation worsened, Weick argued that the sensemaking processes of the crewmembers collapsed, as it became increasingly implausible to rationalize what was taking place in a manner consistent with their original conceptualization of the fire. The firefighters became immobilized, cognitively if not physically, by their inner desire to cling to their initial categorization of the fire. Weick explained, then, that the key question is not the quality of decisions made by the firefighters, but their inability to socially construct a new reality as the fire progressed:

> "People in Mann Gulch did not face questions like where should we go, when do we take a stand, or what should our strategy be? Instead, they faced the more basic, the more frightening feeling that their old labels were no longer working. They were outstripping their past experience and were not sure either what was up or who they were. Until they develop some sense of issues like this, there is nothing to decide." (Weick, 1993: 636)

Weick's argument remains incredibly powerful, yet one finds it hard to distinguish or disentangle decision-making from sensemaking processes in many situations. To understand the actions and behaviors that lead to catastrophic failures, one has to examine the iterative cycle of actions, sensemaking, and decisions that takes place over time. Gary Klein's (1998) work attempts to accomplish this synthesis through his model of recognition-primed decision-making. In that research, based on observations and interviews of firefighters, nurses, and naval commanders, Klein examines how people – particularly experts – engage in pattern recognition based on past experiences. In his model, a situation generates certain cues, which enable individuals to identify patterns based on past experiences. That pattern recognition process triggers a set of action scripts, which individuals assess by mentally simulating the sequence of events that will take place if they act in accordance with those scripts. What Klein describes is

essentially a *simulated sensemaking process* that takes place within the mind as one rapidly assesses an action script. Sensemaking and decision-making come together inside of this mental simulation exercise in that split-second moment of intuitive decision on the part of an expert. Of course, pattern recognition processes can go astray, even for experts. After all, pattern recognition is based to a large extent on analogical reasoning. Scholars have shown that even the most intelligent and capable experts can make mistakes when trying to reason by analogy, often because individuals focus too intently on the similarities between a current situation and a past experience, while downplaying or ignoring the differences (Neustadt and May, 1986; Rivkin and Gavetti, 2005).

From Competing Theories to Multiple, Mutually Reinforcing Lenses

Individuals typically recognize that equifinality characterizes many situations. In other words, they understand that there are many ways to arrive at the same outcome. Nevertheless, we have a natural tendency to want to play the "blame game" when catastrophic failures occur. For instance, when we witness a medical accident, we cry malpractice; we seek to blame someone for the error, rather than attributing the poor performance to external and contextual factors. Scholars have avoided this trap to a large extent, by focusing on the systemic, cultural, and contextual conditions that make organizations susceptible to catastrophic failures. However, scholars too fail to embrace fully the notion of equifinality. Rather than bringing multiple theoretical perspectives to bear on a particular situation, we often enjoy pitting competing theories against one another. We try to argue that our theoretical perspective outperforms other explanations of a particular failure. Many scholars in this field have gone to great lengths to argue that one body of theory is best-suited to explain a specific catastrophic failure. For

instance, Vaughan (1996) spends considerable time explaining why various theories – including groupthink – do not accurately explain the causes of the *Challenger* accident. She argues quite persuasively that her theory of the normalization of deviance provides a more compelling explanation based on the facts of the case. Similarly, scholars often focus on a particular level of analysis – individual, group, or organizational – rather than looking at how activities at each level may interact with one another.

Some scholars, however, have shown that applying multiple conceptual lenses, as well as employing multiple levels of analysis, can be a powerful way to explore a complex organizational phenomenon. In his 1971 book, Essence of Decision, Graham Allison provided a riveting explanation of the Cuban Missile Crisis, drawing on three quite distinct intellectual traditions. Yet, scholars of organizational failure often did not employ such a methodology, until two recent studies that examined complex failures both from multiple levels of analysis and different conceptual lenses. Snook's (2000) study of the 1994 friendly fire incident in Iraq and my study of the 1996 Mount Everest tragedy (Roberto, 2002) each tried to examine a failure at the individual, group, and organizational levels of analysis.

The Everest study employed behavioral decision theory at the individual level, the theory of psychological safety at the group level (Edmondson, 1999), and normal accident theory at the organizational system level to explain why an unprecedented tragedy took place in May 1996. Snook employed a wider range of theories to explain the friendly fire accident, drawing on quite diverse conceptual lenses such as the theories of social impact, team design, and sensemaking. Both studies have one thing in common. Unlike Allison, they did not present these various theories as alternative explanations for the failures. Instead, they sought to examine the linkages among the psychological and sociological forces involved at the individual, group, and

15

organizational system level. These two studies demonstrated that multiple conceptual lenses, applied at three levels of analysis, could serve as complementary and mutually reinforcing explanations.

The Everest study provided important lessons about how the problems of cognitive bias, team psychological safety, and complex systems relate to one another and together enhance the risk of serious organizational failures. Often, scholars viewed these individual, group, and system-level explanations as distinct ways to explain flawed organizational strategies and outcomes. The Everest study demonstrated that an absence of team psychological safety makes it more difficult to avoid cognitive bias, because individuals do not question one another, test assumptions, and express minority views. Systematic biases in judgment become especially problematic in complex systems, because one mistake can trigger a series of other breakdowns in the system. Finally, an absence of psychological safety enhances the risk inherent in complex systems, because candid discussions do not occur about the sources of failure and the interconnections among different components of the system.

Snook's study took a slightly different, but highly fruitful, approach to multi-level analysis. After examining the tragedy at three distinct levels of analysis, he concluded that he had not yet created a complete explanation of the failure. Thus, he sought to formulate a cross-levels account of the tragedy. From that effort, he posited a new theory of organizational behavior, which he coined "practical drift." According to this theory, organizations establish rules and procedures, but individuals and units within the organization constantly seek better ways of doing things. They engage in practical actions that appear to be locally efficient. These locally efficient procedures become accepted practice and perhaps even "taken-for-granted" by many people. Gradually, actual practice drifts from official procedure. The drift is not a

problem most of the time; but in certain unstable situations, it gets us into big trouble. Many

instances of practical drift occurred in the friendly fire incident examined by Snook. For

instance, official rules and procedures called for helicopter pilots to switch radio frequencies

when flying from Turkey into the No Fly Zone in northern Iraq. However, pilots had decided

that it was more efficient, and perhaps even safer, to remain on the same frequency when flying

short missions just over the border. Gradually, this commonplace technique became accepted

practice and taken-for-granted. As pilots rotated out of the Iraqi task force and new pilots

replaced them, this accepted practice, rather than the official procedure, was passed on from one

group to another. Remaining on the same frequency did not cause a tragedy for over one

thousand days of the operation, but on that fateful day in April 1994, the other layers of

organizational defense fell by the wayside, and the failure to change radio frequencies served as

one of the causes of the catastrophe. For Snook, multiple levels of analysis proved crucial to

the development of a compelling new theory, as well as a comprehensive explanation of a

tragedy. Snook concluded: "I am more convinced than ever that we cannot fully capture the

richness of such complex incidents by limiting ourselves to any one or even a series of isolated,

within-level accounts… [we must] capture the dynamic, integrated nature of organizational

reality." (Snook, 2000: 179)

From Failure to Success: High Reliability Organizations

As researchers began to gain a better understanding of why catastrophic failures

occurred, a group of scholars started to work concurrently to study complex organizations that

have operated with very, very few major safety incidents for many years. Scholars coined the

term "high reliability organizations" (HRO) to describe these entities. Prominent scholars in the field include Karlene Roberts, Karl Weick, Todd La Porter, Paula Consolini, Gene Rochlin, and Kathleen Sutcliffe. They have studied organizations such as aircraft carriers and air traffic control centers. The error rates for these organizations are remarkably low, given the hazardous conditions in which they operate. For instance, Roberts (1990) reported that the aviation safety record for pilots operating on naval aircraft carriers were amazingly low – slightly less than three fatalities per 100,000 hours of flight time. Roberts argued that HROs have managed to keep error rates very low by developing mechanisms to cope with interactive complexity and tight coupling. For instance, she found that HROs employ technical and social redundancy to reduce the likelihood of a catastrophic failure. Technical redundancy refers to components and equipment that serve as back-ups in case of failure by the front-line systems. Social redundancy refers to the processes by which individuals back up one another and check on the work of others. Besides redundancy, Roberts discovered that HROs push decision-making down in the organization, to ensure that the appropriate experts apply their knowledge in critical situations. They also employ training and communication mechanisms to ensure that all individuals understand the interconnections among various subsystems and maintain updated, accurate situational awareness throughout a hazardous procedure such as the landing of a jet on an aircraft carrier.

In 2001, after roughly a decade of HRO research, Weick and Sutcliffe wrote a book, titled <u>Managing the Unexpected</u>, in which they tried to synthesize and integrate what they and others such as Karlene Roberts had learned about these complex organizations that performed very reliably in hazardous conditions. They coined the term *mindfulness* to describe the simultaneous existence of five key characteristics of HROs. First, HROs appeared to be

preoccupied with failure of all sizes and shapes. They did not dismiss small deviations, or settle

on narrow, localized explanations of these problems. Instead, they treated each small failure as a

potential indication of a much larger problem. David Breashears, the highly accomplished

mountaineer, has described the preoccupation with failure that characterizes successful

expeditions. He has argued that many climbers focus almost exclusively on the notion of success

(with success being defined as reaching the summit of one of the world's tallest mountains).

These climbers think that they will enhance their odds of reaching the summit if they wall

themselves off to any possible image or conception of failure. Breashears argues that this is

precisely the wrong approach. Great climbers obsess with the notion of failure. They simulate

failure scenarios in their head, long before they head to a mountain. Moreover, they pay close

attention to every detail, and they painstakingly assess whether a small deviation from plan

might compromise the entire expedition. Organizations such as Toyota similarly have been

found to obsess with failure, building organizational processes for highlighting and assessing the

smallest of deviations that occur during the automobile production process.

Second, Weick and Sutcliffe argued that HROs exhibit a reluctance to simplify

interpretations. In short, we all seek to simplify the complex world around us. We create

categories and labels to help us cope with complexity. For the most part, this simplification

process serves us well. Without it, we sometimes would find ourselves paralyzed in the face of

ambiguous data and stimuli. However, we get into trouble when we fail to shed old categories

and labels as contextual conditions change. In some cases, we oversimplify a situation, by trying

to slap a label or category on a situation for which it doesn't quite fit. HROs try to maintain a

healthy diversity of perspectives within the organization, and they constantly test their simplified

models of reality. They build in requisite variety, particularly with regard to the diverse composition of teams, so as to reduce the likelihood of over-simplification or mis-categorization.

Third, HROs demonstrate sensitivity to operations. They do not allow the emphasis on the big picture – strategic plans, vision statements, etc. – to minimize the importance of front-line operations, where the real work gets done. HROs encourage the people on the front lines to identify latent errors – those lapses in organizational defenses that do not usually cause serious negative consequences, but which might suggest system deficiencies that should be corrected, lest a large-scale failure eventually occur (Reason, 1997).

Fourth, HROs exhibit a commitment to resilience. They recognize that no hazardous and complex system will be error free. However, they develop mechanisms for catching and recovery from small failures before they cascade through multiple subsystems of an organization. HROs have the ability to improvise when small errors occur, and they move quickly to learn from failure.

Finally, HROs ensure that expertise is tapped into at all levels of the organization. They push decision-making authority down, and they migrate decisions in real-time to the location in the organization where the most relevant expertise lies. They do not let formal hierarchy, status, or power dictate decision-making during times of stress. Instead, they try to ensure that people at lower levels can apply their expertise to solve thorny and urgent problems without fear of retribution from their superiors.

Taken together, these five characteristics comprise what Weick and Sutcliffe call "mindfulness." When they use the term mindful, they mean that HROs are highly attuned to the unexpected, to the ambiguous threats that emerge at all levels. They remain intellectually curious at all times, and they operate in a spirit of inquiry when faced with tough issues. In other

words, they seek to open up new avenues of discussion when the unexpected occurs, rather than trying to press quickly toward a simple explanation or solution.

Critique of HRO Theory

The HRO literature has several important limitations and weaknesses. First, some scholars question the very definition of HROs. They point to Roberts' definition, in which she states that if an organization could have encountered catastrophic failures thousands of times, but did not, then it can be described as highly reliable. These researchers suggest that a wide array of organizations fit these criteria, rendering the classification rather meaningless (Marais, Dulac, and Leveson, 2004). Scholars also argue that many of the organizations studied by HRO researchers do not actually exhibit the high level of interactive complexity and the tight coupling described in Perrow's model of normal accidents (Marais, Dulac, and Leveson, 2004).

A second major line of critique centers on the concept of redundancy, which is at the heart of many HRO studies. Here, scholars can point back to Perrow's original theory, in which he stressed that redundancy often can *enhance* the likelihood of catastrophic failure in high-risk systems, because redundant features can either increase interactive complexity or tight coupling. That negative side effect holds true not only for more redundant technical systems, but also social ones. More behavioral rules and procedures can create more rigidity, i.e. tight coupling. Redundant systems also can make it even more difficult for individuals to diagnose what is wrong with a complex technology. Moreover, people can become complacent, indeed overly reliant on redundant technological systems at times, as the military officers might have been in the friendly fire incident in northern Iraq. Snook also pointed out that social redundancy can be problematic, because individuals are not always acting independently – a key precondition for

effective redundancy. Snook found that efforts to create social redundancy perversely led to a diffusion of responsibility and ambiguous lines of authority. As a result, the human checks and balances did not function properly.   The fighter pilots' over-reliance on that flawed social redundancy may have contributed to the tragedy.

Finally, the HRO literature has been criticized because it appears to be arguing that increased vigilance lies at the heart of all efforts to increase safety.   In short, are we simply asking organizations to be more careful?   HRO scholars have identified some key attributes of organizations that are highly attuned to safety considerations, but more work could be done to identify the concrete mechanisms and processes that organizations can put into practice to increase safety.   How precisely do I become more careful?   How do we move from individual care to organizational vigilance? Second, and perhaps more importantly, some scholars have questioned whether HRO scholars have given due consideration to the tradeoff between safety and other organizational goals and objectives. Marais, Dulac, and Leveson (2004) have pointed out, for instance, that aircraft carriers operating during peacetime have safety as a primary goal, given that they are simply running training missions. At the end of the day, the objective is to return each pilot safety to the deck of the ship, not to hit enemy targets, etc. Few thorny tradeoffs between safety and other objectives exist during peacetime operations, which is when HRO scholars studied aircraft carrier operations.

Future Directions

While this chapter clearly illustrates that we have learned a great deal about the causes of catastrophic failures, more work needs to be done to help us understand how such accidents can be prevented. The HRO literature has made substantial strides in this area, but it has some

limitations. The question remains: How does an organization become "preoccupied with failure" and without completely sacrificing its other goals and objectives?   If organizations become attuned to all the ambiguous signals or threats that arise, will they ever get any of their work done?   Can a company serve customers and make reasonable profits if it is constantly sounding alarms about subtle and weak signals of potential safety or quality issues?

Recent studies have begun to explore this question. One promising line of inquiry involves the study of Rapid Response Teams in health care organizations (Buist, Moore, et al., 2002; Park, 2006; Roberto, Bohmer, and Edmondson, 2006). Many hospitals have begun to implement a process for amplifying and exploring weak signals.   When "code blues" occurred in hospitals in the past (i.e. crisis teams responding to a patient experiencing cardiac arrest), staff members often noted that the patient's condition had been deteriorating for several hours prior to the code.   However, front-line staff members – typically nurses – often did not voice their concerns to those with more status and formal authority in the organization (such as surgeons). Researchers found that several early warning signs tended to precede unexpected cardiac arrests. In some cases, these pieces of data were quantifiable, i.e. acute changes in respiratory rate or oxygen saturation. In other instances, nurses might simply notice a shift in the patient's appearance, cognitive ability, or demeanor.

Hospitals have now created lists of these early warning signs, and they have empowered nurses to call in a "rapid response team" – if they see one of these warning signals – so as to help them assess the significance of these ambiguous threats.   Hospitals often describe this process as "calling a trigger" – meaning that front-line staffers are identifying a small problem that may trigger a serious incident in the near future. These rapid response teams are cross-disciplinary groups that are highly skilled at quickly assessing whether a warning sign merits further action.

This rapid response process enables and empowers inexperienced nurses to speak up when something doesn't look right. Many hospitals have reported substantial decreases in the number of "code blues" after implementing rapid response teams. They also report that many other improvement ideas have emerged from this process, even in instances when the threats did not prove to be "real." Perhaps most importantly, the benefits of these new rapid response teams easily outweigh the costs, and they do not appear to interfere unnecessarily with the care of other healthier patients.

Why do these rapid response teams increase preoccupation with failure and enhance patient safety without unduly increasing costs or sacrificing other organizational goals and objectives (such as profit maximization)?   Here, we must go back to the concept of shared cognition which is at the heart of the HRO literature. Individuals and groups working in HROs build a strong sense of shared situational awareness, as well as a strong sense of system awareness. They understand how situations are unfolding, how their work relates to that done by others, and how the human and technical systems function. Bigley and Roberts (2001) point out that HROs constantly update these cognitive representations of a complex reality, processing real-time information very quickly to maintain a current shared picture of a hazardous situation.

To understand how rapid response teams enhance mindfulness without unduly burdening the organization with excess costs, we have to bring together two streams of work: the HRO scholars' work on mindfulness and Klein's (1998) work on pattern recognition. Recall that Klein has shown how expert decision-makers discern patterns very quickly based on past experience, and this pattern recognition triggers action scripts. One can argue that rapid response teams build individual as well as shared pattern recognition capabilities within the organization.

They are a concrete mechanism for not only surfacing weak signals, but also rapidly and inexpensively evaluating those signals and refining the diagnosis of those signals over time.

As experts interact with front-line employees, and constantly react to weak signals, they may become better and better at recognizing patterns – as individuals, as teams, and as an entire organization. First, people may become better at identifying the subtle, non-quantifiable signals of future harm. In particular, as experts interact with novices, they help develop the pattern recognition capabilities of newer, inexperienced front-line employees. Consequently, novices get better at spotting weak signals, and perhaps spot them earlier as well. Novices get better, in particular, when the warning signs are non-quantifiable, e.g., when the patient's appearance has simply changed without alterations in his or her vital signs. Second, people may begin to develop a stronger capability to discern which signals are real, and which are merely "false alarms." In the long run, that will reduce the number of triggers that are called for non-serious situations. Third, people may learn how to more quickly assess the validity of a weak signal – thereby reducing the time associated with each deployment of a rapid response team. Finally, people begin to see patterns that cut across the organization, i.e. these types of patients tend to experience triggers for this specific set of reasons – we need to be proactively attentive to those situations. The organization can then move from simply being reactive – responding to weak signals after they arise – toward a more proactive orientation in which they seek out the situations in which signals are more likely to arise and accidents are more likely to occur. That proactive approach can not only improve patient safety, but it also may reduce costs, as it becomes more expensive to treat a patient if their condition deteriorates.

More work certainly needs to be done to develop our understanding of the concrete mechanisms by which organizations become more mindful without sacrificing other goals and

objectives. The rapid response team work represents one promising line of inquiry; many others surely exist. For instance, in a 1999 Harvard Business Review article, Steven Spear and Kent Bowen examined the inner workings of the Toyota Production System. Toyota represents an organization which is extremely preoccupied with small failures. Without question, they are an auto industry leader in terms of quality. However, they also have one of the highest levels of productivity in the industry. They appear to have developed a system whereby reliability is enhanced without sacrificing cost efficiency.

As scholars continue to study catastrophic failure, they should take heart from the recent work that has been done in the health care field as well as on the *Columbia* shuttle accident. Hospitals around the world are trying to learn from the literature discussed in this chapter. Similarly, the *Columbia* accident investigators reached out to scholars such as Diane Vaughan and Karl Weick to understand what went wrong at NASA. In both instances, we see practitioners seeking out scholars to help them understand the organizational causes of accidents and the potential prescriptions for reducing the risk of catastrophic failures. Research appears to be impacting practice, and rapid advances in practice are informing new theory development. That synergy between scholarship and practice promises to be much fruit in the years ahead.

Bibliography

Bigley, G. and K. Roberts (2001). "The incident command system: High reliability organizing for complex and volatile task environments." *Academy of Management Journal*. 44(6): 1281-1299.

Buist, M., G. Moore, S. Bernard, B. Waxman, J.Anderson, and T. Nguyen. (2002) "Effects of a Medical Emergency Team on Reduction of Incidence of and Mortality from Unexpected Cardiac Arrests in Hospital: Preliminary Study." *British Medical Journal.* 324:387-390.

Deal, D. W. (2004) "Beyond the Widget: *Columbia* Accident Lessons Affirmed" *Air and Space Power Journal,* Summer.

Edmondson, A. (1999) "Psychological Safety and Learning Behavior in Work Teams." *Administrative Science Quarterly* 44(4): 350-383.

Esser, J. and J. Lindoerfer (1989) "Groupthink and the Space Shuttle Challenger accident: Toward a quantitative case analysis." *Journal of Behavioral Decision Making*. 2: 167-77.

Gavetti, G., D. Levinthal, and J. Rivkin (2005). "Strategy-making in Novel and Complex Worlds: The Power of Analogy," *Strategic Management Journal* (26): 691-712.

Janis, I. L. (1982). Groupthink: Psychological studies of policy decisions and fiascos. 2nd Edition. Boston, Houghton Mifflin.

Klein, Gary A. (1998) Sources of power: How people make decisions. Cambridge, MA, MIT Press.

Marais, K., N. Dulac, and N. Leveson (2004). "Beyond normal accidents and high reliability organizations: The need for an alternative approach to safety in complex systems." *MIT Engineering Systems Division Symposium.*

Moorhead, G., R. Ference, and C. Neck (1991). "Group decision fiascos continue: Space shuttle Challenger a revised Groupthink framework." *Human Relations*. 44(6): 539-550.

Neustadt, R.E & May, E.R. (1986). Thinking in time : the uses of history for decision-makers. New York: Free Press.

Park. J. (2006). "Making rapid response real: Change management and organizational learning in critical patient care." Unpublished thesis. Harvard University.

Perrow, C. (1981). "Normal accident at Three Mile Island." *Society*. 18: 17-26.

Reason, J.T. (1997). Managing the Risks of Organizational Accidents. Aldershof, England, Ashgate

Roberto, M.A. (2002) "Lessons from Everest: The Interaction of Cognitive Bias, Psychological Safety, and System Complexity." *California Management Review* 45(1): 136-158.

Roberto, M., R. Bohmer, and A. Edmondson (2006). "Facing Ambiguous Threats." Harvard Business Review. 84(11): 106-113.

Roberts, K. (1990). "Managing high reliability organizations." California Management Review. 32(4): 101-113.

Snook, S. A. (2000). Friendly Fire: The Accidental Shootdown of U.S. Black Hawks Over Northern Iraq. Princeton, NJ: Princeton University Press.

Starbuck, W. and M. Farjoun (2005). Organization at the Limit: Lessons from the *Columbia* Disaster. London: Blackwell Publishing.

Vaughn, D. (1996) The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA Chicago, University of Chicago Press.

Weick, K. (1976). "Educational organizations as loosely coupled systems." *Administrative Science Quarterly*. 21: 1-19.

Weick, K. (1993). "The collapse of sensemaking in organizations: The Mann Gulch Disaster," *Administrative Science Quarterly*.

Author Name

Michael A. Roberto, Trustee Professor of Management Bryant University

Suggested Readings

Allison, G.T. (1971). The essence of decision: explaining the Cuban missile crisis. Boston: Little Brown.

Columbia Accident Investigation Board (CAIB) Report, August 26, 2003, Washington, D.C.

Edmondson, A. (1996). "Learning from mistakes is easier said than done: Group and organizational influences on the detection and correction of human error." *Journal of Applied Behavioral Sciences* 32(1): 5-32.

Perrow, C. (1984). Normal accidents. New York, Basic Books.

Roberto, M. (2005). Why great leaders don't take yes for an answer: managing for conflict and consensus. Upper Saddle River, NJ: Wharton School Publishing.

Russo, J. E., and P. Schoemaker (1989) Decision Traps: Ten Barriers to Brilliant Decision-Making and How to Overcome Them. New York: Doubleday.

Spear, S. and K. Bowen (1999). "Decoding the Toyota Production System." *Harvard Business Review*. Sept-Oct: 96-106.

Turner, B. (1978). Manmade Disasters. London: Wykeham.

Useem, M. (1998) <u>The Leadership Moment: Nine True Stories of Triumph and Disaster and Their Lessons for Us All</u>. New York, Random House.

Weick, K. and K. Sutcliffe (2001). <u>Managing the Unexpected</u>. San Francisco: Jossey Bass.