The limitation of Safety
- Organizations, accidents and nuclear weapons.

Scott D. Sagan

Princeton University Press, Princeton, New Jersey. 1993

Chapter 1. The origins of accidents

Accidental nuclear war is a very difficult subject to study. The first reason for this difficulty is the most obvious, the most important, and the most fortunate one: there has never been a single accidental nuclear detonation, much less an accidental nuclear war. The traditional comparative methodology used by social scientists to explain complex political phenomena – comparing and contrasting cases in which the phenomena occurred against cases in which it did not – cannot be used here. There are, of course, many other difficulties involved in a thorough investigation of this subject. Many important pieces of evidence about past nuclear weapons incidents remain classified; some critical documents about sensitive military operations have been destroyed; faulty historical records have sometimes been created by military units; and inadequate social science theories exist to help us understand both the causes of war and the origins of accidents. But these difficulties pale next to the basic conceptual dilemma posed by the problem of accidental nuclear war. How does one even begin to study something that has never occurred?

One possibility is to assume that this central fact proves that the danger of nuclear weapons accidents and accidental nuclear war is minimal. Such an optimistic assessment is not usual, in fact, since it can be argued that history has demonstrated that nuclear weapons can be maintained and operated in a safe and secure manner. This assessment, however, is inadequate for at least three reasons. First, things have never happened before happen all the time in history. There must be a first time for every type of historical event that has occurred in the past, and the lack of earlier nuclear accidents is therefore insufficient evidence for making such a strong statement about future possibilities. Second, nuclear weapons have existed for less than fifty years and have been in the possession of only a small number of nations. This is a very limited pool of experience on which to base confident assessments of long-term nuclear weapons safety, especially under what could be quite different conditions during the next century. Third, an assessment of the risk of accidental nuclear war should examine close calls to catastrophe, and not be satisfied with the simple fact that accidental nuclear war has never occurred. For if we have had numerous “near-accidents” with nuclear weapons – incidents that could have resulted in an accidental nuclear war had they occurred under other plausible circumstances - even an apparently perfect final safety record may not inspire extreme confidence.

A more thorough assessment therefore requires a deeper investigation into the hidden history of nuclear weapons. What has been the complete safety record, of accidents and near-accidents, with these weapons and their command and control systems? How have the military organiza-

tions that have custody over nuclear weapons been able to manage their complex operations with such apparent success? Have these organizations done something extremely intelligent to avoid accidents or have they merely been extremely lucky?

A useful place to start is to examine the causes of accidents and safety problems in other similar sociotechnical systems. For although there has never been an accidental nuclear weapons detonation or war, there have been numerous serious accidents in recent decades in other complex high-technology systems such as nuclear reactors, commercial and military aircraft, space programs, international shipping, and large petrochemical plants. Fortunately, a rich scholarly literature studying the causes of reliability and safety in these industries exists. What has caused serious accidents with these hazardous technologies? What organizational designs and strategies have been used to prevent accidents and enhance safety? A number of scholars have sought to explain successes and failures in organizational safety, and their ideas, if used very carefully, can help us understand the risks of serious accidents with nuclear weapons.

This chapter will examine the two most important schools of thought within the organization theory literature concerning the issue of safety and reliability in complex technological systems. Subsequent chapters will then apply these theories to the military organization that control U. S. nuclear weapons and test the theories competitively against one another by probing into the historical record of nuclear weapons accidents and "close calls" to accidental nuclear war. The goal is to provide a clearer understanding of the origins of accidents and the causes of safety.

ORGANIZATION THEORY AND ACCIDENTS

Even a brief glance at the recent history of high-technology industries cautions against complacency. Why have such tragedies as Chernobyl, the Exxon Valdez, and Bhopal occurred? Are such accidents preventable? Or are they the inevitable consequence of the widespread use of hazardous technologies in the modern world?

The scholarly literature about complex organizations is large and diverse, but two general competing schools of thought on this specific issue exist. The first is the optimistic view of what I will call "high reliability theory," whose proponents argue that extremely safe operations are possible, even with extremely hazardous technologies, if appropriate organization design and management techniques are followed. The second school, what I will call "normal accidents theory," presents a much more pessimistic prediction: serious accidents with complex high technology systems are inevitable.

The term school of thought was used deliberately, since it is in many ways a better description of what exists in this literature on hazardous technologies than the term theories. The scholarship I will be analyzed is based on mixtures of abstract deductive logic and inductive empirical observation, and the authors within each school by no means agree on all details concerning organizational safety. Specific terms that appear often in this literature are not always used in a consistent manner. And perhaps most importantly the predictions of both schools are often im-

precise. Nevertheless, proponents of each school do focus attention on a specific set of factors that they believe contributes to or decreases safety, and each school develops a set of general hypotheses that is meant to hold true in a variety of organizations across space and time. These ideas can therefore be viewed as nascent social science theories and can usefully be tested against one another.

These two schools of thought have intellectual roots in different traditions within the organization theory literature; they have different basic understandings of how organizations work and hold different views on how best to analyze complex organizations. The theories offer competing general explanations for the causes of accidents with hazardous technological systems and offer alternative prescriptions for improving safety in the future. At the broadest level, they have conflicting visions about what could be called the degree of perfectibility that is possible in complex organizations. Finally, and importantly for this book's purposes, high reliability theory and normal accidents theory lead to very different predictions about the causes and the likelihood of serious nuclear weapons and command and control accidents. Each will therefore be analyzed in some detail.

HIGH RELIABILITY ORGANIZATION THEORY

How safe are nuclear power plants, commercial aircraft, oil tankers, petrochemical factories, and other potentially dangerous high-technology systems? Is it possible to design and manage such complex organizations so well that, even though they use inherently hazardous technologies, they are unlikely to produce serious accidents? One group of organization theory scholars – the high reliability theories – are in essential agreement with the professional risk analysts and engineers who build these systems: serious accidents with hazardous technologies can be prevented through intelligent organizational design and management. Scholars in this school have studied a variety of high risk organizations and have reached quite optimistic conclusions about the prospects for safely managing hazardous technologies in modern society.

Three major scholarly efforts to understand safety problems in such hazardous high-technology organizations can be best represent the high reliability school. First, Joseph Marone and Edward Woodhouse's *Averting Catastrophe: Strategies for Regulating Risky Technologies* is an innovative study of the management of toxic chemicals, nuclear power, recombinant DNA research, ozone layer depletion, and global warming problems in the United States. The authors maintain that "given the challenge posed by modern technologies, the record is surprisingly good: despite dire warnings, no catastrophes have occurred in the United States." This positive historical record of safety was, according to Marone and Woodhouse, "a systematic product of human actions – the result of a deliberate process by which risks are monitored, evaluated, and reduced," and much of their research therefore focuses on identifying the specific organizational processes and strategies which produced this result. Although they acknowledge that the strategies they discovered were not always fully developed or perfectly implemented, Marone and Woodhouse nevertheless believe that "taken together, the strategies we found in use suggest the

elements of *a complete system for averting catastrophe*," Continued use of such wise management practices can therefore maintain and improve this safety record well into the future: "There is a good chance in areas of civilian technology that catastrophes will be prevented – even in new problem areas where society is not presently expecting trouble."

The second example of high reliability theory is the work of a multidisciplinary group of scholars based at the University of California at Berkeley that has studied the "design and management of hazardous organizations that achieve extremely high levels of reliable and safe operations," Their empirical research has focused on three hazardous organizations that they argue have achieved "nearly error free operations": the Federal Aviation Administration's (FAA) air-traffic control system, the Pacific Gas and Electric Company's electric power system (which includes the Diablo Canyon nuclear power plant), and the peacetime flight operations of two U. S. Navy aircraft carriers. Their extensive field research into the daily operations of these specific organizations has identified a number of strategies and processes that re believed to have produced such impressive safety records. Although the Berkeley authors emphasize that further comparative research is needed to provide confident prescriptions for other hazardous organizations, they nonetheless maintain that "we have begun to discover the degree and character of effort necessary to overcome the inherent limitations to securing consistent, failure free operations in complex social organizations." These successful organizations are therefore viewed as providing important lessons for the management of other risky high-technology systems: "Most of the characteristics identified here should operate in most organizations that require advanced technologies and in which the cost of error is so great that it needs to be avoided altogether."

The third work that will be used to represent the high reliability school is Aaron Wildavsky's *Searching for Safety*, a more deductive effort to develop "a theory accounting for the considerable degree of safety achieved in contemporary society," The avowed purpose of the book is to alert readers to "the increases in safety due to entrepreneurial activity" in a variety of complex systems: "If this essay is persuasion achieves its purpose, the focus of the risk debate will shift from the passive prevention of harm to a more active search for safety," Wildavsky's major focus is on examining the cost and benefits of what he calls two "universal strategies" for improving safety: "anticipation" (efforts to predict and prevent potential dangers from arising before they have ever occurred) and "resilience" (efforts to cope with dangers once they have ever occurred). The book then illustrates its central ideas by presenting evidence on the degrees to which anticipation and resilience have improved safety in as diverse a set of systems as nuclear power plants, the human body's' immune system, and the Food and Drug Administration's (FDA) drug approval process.

How can serious accidents be prevented, and how can safety be improved in organizations managing highly hazardous technology? These three groups of high reliability theorists have not focused entirely upon then same set of hazardous organizations, and they differ on a number of specific details of explanation and prescription. They share, however, one essential assumption about how such organizations function, and their analyses point to a common set of four major factors that are seen as contributing to high degrees of safety.

The common assumption of the high reliability theorists is not a naïve belief in the ability of human beings to behave with perfect rationality; it is the much more plausible belief that organizations, properly designed and managed, can compensate for well-known human frailties and can therefore be significantly more rational and effective than can individuals. The high reliability theory can be best seen therefore as fitting into the tradition that W. Richard Scott has called the "closed rational systems" approach in organization theory. High reliability hazardous organizations are seen as "rational" in the sense that they have highly formalized structures and are oriented toward the achievement of clear and consistent goals (in this case, extremely reliable and safe operations). They are relatively "closed systems" in the sense that they go to great efforts to minimize the effects that actors and the environment outside the organization have on the achievement of such objectives. As Todd La Porte and Paula Consolini have argues, high reliability organizations "are characterized by very clear, well-agreed-upon operational goals":

> Those in the organizations carry on intensive efforts to know the physical and dynamic properties of their production technologies, and they go to considerable pains to buffer the effects of environmental surprises. In most regards, the organizations come close to meeting the conditions of closed rational systems, i.e., a well-buffered, well-understood technical core requiring consistency and stability for effective, failure-free operations.

The research of the different high reliability theorists has produced similar explanations for positive safety records within a wide variety of organizations. Four critical causal factors have been identified: the prioritization of safety and reliability as a goal by political elites and the organization's leadership; high levels of redundancy in personnel and technical safety measures; the development of a "high reliability culture" in decentralized and continually practiced operations; and sophisticated forms of trial and error organizational learning. These four factors continue, according to this school of thought, a route to extremely reliable operations even with highly hazardous technologies.

Leadership Safety Objectives

The first and most obvious requirement for high reliability organizations is that extreme reliability and safety must be held as a priority objective by both political leaders and by the heads of the organization. Such organizations maintain "the goals of avoiding altogether serious operational failures," according to La Porte and Consolini: "This has nurtured an organizations perspective in which short-term efficiency has taken a second seat to very high-reliability operations." Consider the case, for example, of the U. S. air traffic control system for commercial airlines. La Porte argues that "the public (and especially its Congressional leaders) demands a system" in the United States that:

- is always safe
- carries anyone, anywhere, anytime (and is always safe)
- enables private carriers to make a reasonable profit (while always being safe)

Although political tensions and disagreements continue to exist over who should pay the costs of airline safety improvements, according to La Porte, "the twin pressures from the travelling public and elites for extraordinarily reliable and safe performance resulted in a system … (in which) the goal of failure-free performance is a central objective of everyone in the system.

The literature postulates two central reasons why political and organizational leaders must place very high reliability is to be achieved. First, high reliability organizations require both significant levels of redundancy and constant operational training and both of these factors cost a great deal of money. It should therefore be no surprise that, in Wildavsky's succinct phrase, "richer is safer," since wealth increases the "level of general resources upon which our safety mostly depends." If political authorities and organizational leaders are not willing to devote considerable resources to safety, accidents will therefore become more likely. La Porte's study of airline safety, for example, emphasizes the fact that Congress has never reduced the budget requested by the FAA in support of air traffic control. Similarly, in her review of safety in U. S. Navy aircraft carrier operations, Karlene Roberts argues that the navy has "gotten Congress to recognize the virtual impossibility of doing this job without enormous amounts of redundancy (in jobs, communication structures, parts, etc.) and training, both terrifically expensive. When hazardous organizations cut corners on either of these issues disaster is likely to occur.

Second, if high reliability organizations require "very clear and well-agreed-upon operational goals," then organizational leaders must place high priority on safety in order to communicate this objective clearly and consistently to the rest of the organization. Over time, the organization will develop a "culture of reliability" as members are socialized into accepting the organizations' operational goals. Roberts notes, to give one example, that the commanding officer of the aircraft carrier she observed was constantly "laying down the culture of the organization" by briefing newly arriving crewmen on the importance of such safety procedures as the "buddy system" and training them never to break the ship's rules "unless safety is at stake." Such clear communication of the importance the leadership places on reliability and safety is considered necessary in order to "assure that there is agreement about organizational mission" by all members of the organization.

The Need of Redundancy

Numerous psychological studies have rigorously demonstrated that we all have suspected from our daily lives: human beings are not perfectly rational machines, but rather operate with limited and fallible cognitive capabilities. Organization theorists, of course, have long been aware of the limits of human rationality and have therefore spent a great deal of effort seeking to solve a basic puzzle. It is possible, in John von Neumann's phrase, to build "reliable systems from unreliable parts"?

The answer, according to the high reliability organization theorists, is a resounding yes, and the key design feature of such organizations is redundancy. Multiple and independent channels of such organization, decision-making, and implementation can produce, in theory, a highly reli-

able overall system, even if each component of the organization is subject to error. Jonathan Bendor's sophisticated study of the effects of redundancy in U. S. government agencies provides a simple safety-related analogy to illustrate why "a system's reliability is not necessarily limited by its components' fallibility":

> Suppose an automobile had dual breaking circuits: each circuit can stop the car, and the circuits operate independently so that it one malfunctions it does not impair the other. If the probability of either one failing is 1/10, the probability if both failing simultaneously is $(1/10)^2$, or 1/100. Add a third independent circuit and the probability of the catastrophic failure of no brakes at all drops to $(1/10)^3$, or 1/1000.

Although many politicians continue to call for eliminating government waste and bureaucratic overlap for the sake of efficiency, these theorist point to the important contribution made by such redundancy to a larger system's reliability. First, duplication (two different units that perform the same function) often exists. The Berkeley group's studies of U. S. aircraft carrier operations, for example, have stressed the critical importance of having both technical redundancy (such as backup computers, antennas, and safety devices specifically kept on board to take the place of any components that fail) and personnel redundancy (such different personnel given the same job of checking the carrier's arresting gear settings before each landing). The use of reserve units, to be brought into action only if the main unit is unavailable, is another obvious example of duplication. Second, overlap (two units with some functional areas in common) also often exists. Different communications systems are used for related purposes during day-to-day operations and therefore each can pick up the slack for another if necessary. Different officers are also given overlapping responsibilities, although their primary duties may differ, to ensure that they cross-check each other's work. Each time a plane lands on the carrier, for example, a continuous loop of orders and verifications is broadcast simultaneously over multiple channels: "This constant flow of information about each safety-critical activity, monitored by many different listeners on several different communications nets, is designed specifically to assure that any critical element that is out of place will be discussed or noticed by someone before it causes problems.

Scholars in this school have found redundancy to be critical to the success of virtually all the successful high-reliability organizations they have studied. La Porte notes, for example, that U. S. air traffic controllers continued to use a voice radio system as a redundant backup device to map the locations of aircraft, even after more accurate radar technology became available, and that the functional responsibilities of individual controllers often overlap in an effort to ensure that more than one set of eyes are monitoring each safety-critical event. Marone and Woodhouse similarly emphasize the central role that redundancy plays in nuclear reactor "defense in depth" safety programs: for example, there must be at least two independent outside power sources for the plant, redundant instruments for measuring the reactor's operating parameters, and several coolant loops, so that if one system failed, others would be able to take over. Roberts has noted the importance of both technical and personnel redundancy in Pacific Gas and Electric's distribution grid, in which reserve power sources compensate for any losses of regular sources and

where multiple control centers monitor the output. Finally, Dennis Coyle and Aaron Wildavsky's analysis of the human immune system also places considerable importance on the degree to which "the defense systems of the body are highly redundant": for example, several antibodies are capable of binding to a single antigen; two kidneys exist to eliminate harmful substances; and bone marrow can take over the job of producing red blood cells if one's spleen is removed. The lessons of theory and evidence are clear for the high reliability school of organization theorists. Redundancy is absolutely essential if one is to produce safety and reliability inside complex and flawed organizations. In Bendor's terms, "duplication is a substitute for perfect parts."

Decentralization, Culture, and Continuity

Although high levels of redundancy in an organization's structure and operations can greatly enhance reliability, it is still beneficial to reduce the number and severity of individual component failures in order to avoid stressing the redundant systems beyond their capacity. The high reliability theorists have therefore focused considerable attention on operations and management strategies that can reduce the burden placed on redundancy. Three related characteristics of operations and management have been identified as contributing to organizational reliability and safety in this literature.

First, it is argues that considerable decentralization of decision-making authority must exist concerning safety issues in high reliability organizations in order to permit rapid and appropriate responses to dangers by the individuals closest to the problems at hand. The need for "decentralized anticipation" runs throughout Wildavsky's theoretical analysis emphasizing the superiority of entrepreneurial efforts to improve safety over centralized and restrictive government regulatory policies. He applies this idea at a more microlevel, for example, by explicitly calling for more discretion to be given to nuclear power plant operators concerning how best to run the plant in a safe manner. Decentralized decision-making has also been found in other high-technology systems studied by these scholars. Although the U. S. Navy aircraft carriers and the U. S. air traffic control system may appear at first glance to be very hierarchical in their decision-making structure, for example, close observation has suggested that surprisingly collegial processes are at work and that considerable operational authority in fact rests at very low levels of the organization. During high-tempo carrier operations, higher ranking personnel, and even the lowest ranking individual on the deck of the ship has the authority (and the obligation) to suspend immediately any takeoff or landing that he believes would result in an accident. Similar patterns of nonhierarchical and decentralized decision-making have been observed in air traffic control centers, where supervisors and controllers may switch responsibilities when necessary and where informal teams are often formed to trade advice and manage dangerous operations at the radar screens. Finally, it is considered important that commercial airline captains have the capability to delegate the task of flying the aircraft to subordinates during onboard crises; otherwise, the crisis can take everyone's attention away from the details of flying, resulting in a crash.

This leads directly to the second operations management factor contributing to safety: the creation of a "culture of reliability" within the organization. Common organizational practices such as the promulgation of formal rules and standard operational procedures can contribute to reliability if the outside environment is stable, that is if the decisions required of operators all fall within a predictable set of contingencies. The central problem for organizations operating with hazardous technologies, however, is that this is quite often not the case. They must cope with unexpected and unique environmental dangers in a very rapid fashion, which is precisely why a significant degree of decentralized authority is deemed necessary. Yet how can such an organization ensure that lower-level personnel will identify situations properly, behave responsibly, and take appropriate actions in crises?

The answer, according to the high reliability theorists, is to recruit, socialize, and train personnel to maintain a strong organizational culture emphasizing safety and reliability. This organizational culture will enable lower-level personnel, even when acting independently, to behave similarly and to make operational decisions that meet the approval of higher authorities. Karl Weick has described the concept extremely well:

> Before you can decentralize, you first have to centralize so that people are socialized to use similar decision premises and assumption so that when they operate their own units, those decentralized operations are equivalent and coordinated. This is precisely what culture does. It creates a homogeneous set of assumptions and decision premises which, when they are invoked on a local and decentralized basis, preserve coordination and centralization.

Concerned efforts to socialize personnel into a "culture of reliability" have been observed within a variety of hazardous organizations. U. S. Navy aircraft carrier commanders use rituals, exercise, and punishments to train officers and enlisted men to follow established rules when it is appropriate, to improvise on their own authority when necessary, and to know which is which. According to Roberts, the navy has unusual advantages in this effort, even when compared to other high reliability organizations, since aircraft carriers are, borrowing Erving Goffman's phrase, a "total institution" in which members are isolated from a wider society and can therefore be more intensely socialized and trained. Air traffic controllers and airplane crews must be also socialized into a set of shared missions and beliefs so that, when important authority is delegated to lower-level personnel, no one questions the action and everyone is confident that the correct operational decisions will be made on the spot. Such manifestations of a "culture of reliability" should promote reliability without hierarchy or extreme centralization. As La Porte and Consolini argue: "Such organizations invest a great deal in recruiting, socialization, and incentives to assure that there is agreement about organizational mission. At the operating levels, there is rarely and question at all. Consensus is unequivocal."

The final elements of intelligent operations management that has been identified in successful hazardous organization is the maintenance of continuous operations and training. "One of the great enemies of high reliability," according to the Berkeley group, "is the usual 'civilian' com-

bination of stability, routinization, and lack of challenge and variety that predispose and organization to relax vigilance and sink into a dangerous complacency that can lead to carelessness and error." A constant process of on-the-job training improvements, frequent and realistic simulations of emergencies, and challenging operational work loads are therefore believed to contribute greatly to reduced error rates. Air traffic control accidents, for example, are reportedly more likely under light traffic conditions, when vigilance is low, than under heavy traffic conditions. The Berkeley team similarly found that the constant flight training mode of aircraft carriers at sea has been critical to their ability to operate reliably: "It is in the real-world environment of workups and deployment, through the continual training and retraining of officers and crew, that the information needed for safe and efficient operation is developed, transmitted, and maintained. Without that continuity, and without sufficient operational time at sea, both effectiveness and safety would suffer."

These three factors re-enforce one another. Constant training, strong cultural norms, and decentralized decision authority can produce, according to Karlene Roberrs, Denise Rousseau, and Todd La Porte, "a self regulating work unit where operators are empowered to directly address risks and uncertainties." In this sense, high reliability organizations are seen as being similar to modern industrial factories that have been designed according to "sociotechnical principles" in which multiskilled workers are organized into loosely supervised and semiautonomous production teams. According to Larry Hirschhorn, "the normal accident will become increasingly abnormal" in such settings since "workers are vigilant and committed to production quality and safety because of their desire to learn, their understanding of plant dynamics and policies, and their close relationships with teammates." To the degree that organizations utilizing hazardous technologies can mimic such industrial practices, the high reliability theory would predict that their operations would become increasingly safe.

Organizational Learning

The final factor necessary for high reliability in hazardous organizations, according to this school of thought, is a strong capability to learn. Two common modes of such organizational learning have been discussed in the theoretical and empirical literature. First, and most importantly, a high reliability organization, like other effective organizations, must adjust its procedures and routines over time, learning through a process of trial and error which activities promote safety and which do not. Such an incremental learning process need not be centrally controlled and will produce evolutionary progress as long as successful designs and standard operating procedures are maintained while unsuccessful ones are eliminated.

A belief in the effectiveness of trial-and-error organizational learning lies at the heart of the high reliability theory. It can be most clearly witnessed in Wildadvsky's argument against excessive government regulations on potentially hazardous technologies: " Without trials there can be no new errors; but without these errors, there is also less new learning." According to Wildavsky:

> Trial and error us a device for courting small dangers in order to avoid or lessen the damage from big ones… Because it is a discovery process that discloses latent errors so we can learn how to deal with them, trial and error also lowers risk by reducing the scope of unforeseen dangers. Trial and error samples the world of as yet unknown risks; by learning to cope with risks that become evident as the result of small-scale trial and error, we develop skills for dealing with whatever may come our way from the world of unknown risks.

This perspective is echoed in the Marone and Woodhouse case studies of nuclear power safety (which note the changes in operational training, emergency procedures, and the design of reactor control panels after the incident at Three Mile Island) and toxic substance controls (which note how the government learned to conduct routine audits of testing laboratories after a series of fraudulent test results were exposed). The Berkeley group's study of aircraft carrier operations similarly maintains that many of the safety-related innovations on U. S. carriers were implemented after aircraft incidents and deck fires in the past and therefore predicts that it will take many years and "some loss of lives in the learning process" before aircraft carriers in other nation's navies can equal the U. S. Navy's reliability and safety rates.

Because the organizations and social costs of accidents with such hazardous technologies are so high, however, a second organizational learning strategy – improving procedures through simulations and imagination of trials and errors – is often used. Marone and Woodhouse have argued that such processes are an important part of a "sophisticated trial and error" strategy. In the most dramatic case, the nuclear power industry deliberately boiled away the coolant under controlled circumstances in small experimental reactors, in an effort to gain understanding of the consequences of such incidents in regular operating nuclear power plants. Less dramatic examples abound. In the mid-1970s, a set of biotechnology experiments were run in order to determine whether a genetically engineered organism could survive outside the laboratory. The FDA's elaborate drug testing process, using laboratory experiments and limited human tests as a simulation of what would happen if the drug was put on the marketplace, is another obvious example. Airplane pilots, air traffic controllers, aircraft carrier crews, and nuclear power plant operators also undergo rigorous training in simulated crises in order to gain some of the experience of trails without large-scale errors. Finally, engineers and other consultants are hired by many hazardous organizations to provide risk analyses: to imagine potential operation and design errors, to draw fault tree diagrams to discover any hidden failure modes, and to identify technical solutions to the problems. In such cases, as Marone and Woodhouse put it, "Rather than wait to learn from experience whether their precautions were adequate, researchers and regulators chose to speed up the learning process."

High Reliability Summary

High reliability theorists believe that hazardous technologies can be safely controlled by complex

organizations if wise design and management techniques are followed. This optimistic conclusion is based on the argument that effective organizations can meet the following four specific conditions, which are necessary to create and maintain adequate safety: (1) political elites and organization leaders place a high priority on safety and reliability; (2) significant levels of redundancy exist, permitting backup or overlapping units to compensate for failures; (3) error rates are reduced through decentralization of authority, strong organizational culture, and continuous operations and training; and (4) organizational learning takes place through a trial-and-error process, supplemented by anticipation and simulation. These conditions have been witnessed in an number of high reliability organizations, and if these conditions exist in other organizations, then the theory would predict that serious accidents and catastrophes can be prevented. As Marone and Woodhouse put it, "While the exact mix of strategies appropriate in a given case obviously depends on the nature of the particular problem, the catastrophe-aversion strategy outlined above should be applicable to virtually any risky technology." The Berkeley group holds similar views: "We feel that most of the characteristics identified here should operate in most organizations that require advanced technologies and in which the cost of error is so great that it needs to be avoided altogether."

Thus, while the high reliability theorists do not state what precise amounts and mixtures of these factors are necessary for operational success with hazardous technologies, their overall optimism is clear. Properly designed and well-managed organizations can safely operate even the most hazardous technologies. Although this literature does not explicitly address the issue of nuclear weapons and command and control safety, the logic of the theory would lead to a similar optimistic prediction: if nuclear weapons and command and control systems are designed and managed according to the factors outlined above, a high degree of safety and reliability will be maintained.

NORMAL ACCIDENTS THEORY

A very different approach to understanding how complex organizations work, however, leads to a much more pessimistic conclusion about the risks of using hazardous technologies in modern society. A second group of scholars, the normal accidents theorists, have examined some of the same industries from a different theoretical perspective and have concluded that although such complex organizations may work hard to maintain safety and reliability, serious accidents are nonetheless a "normal" result or an integral characteristic of the system. Serious accidents in organizations managing hazardous technologies may be rare, but they are inevitable over time. The belief that intelligent design and management will result in complex organizations that are capable of safely operating hazardous technology is an illusion according to this perspective. What are the central assumptions and theoretical underpinnings of the normal accidents school that lead to this conclusion?

The high reliability theorists view successful hazardous organizations as reasonably rational

actors: they have consistent and clear goals and can therefore learn how to maximize those objectives over time. If one relaxes those assumptions, however, a far more complicated and conflictual vision of organizational behavior emerges. James March has highlighted the profound changes in perspective that result from starting with a different set if assumptions:

> As long as we assume that organizations have goals and those goals have some classical properties of stability, precision and consistency, we can treat an organization as some kind of rational actor. But organizations do not have simple, consistent preference functions. They exhibit internal conflict over preferences. Once such conflict us noted, it is natural to shift from a metaphor of problem solving to a more political vision."

Organized Anarchies and the Garbage Can Model

The resulting approach to understanding complex organizations fits more closely into the "natural open systems" tradition. Organizations are seen as "natural" in that they (as is natural with all social groups) actively pursue goals of narrow self-interest, such as their own security and survival, and not just their official goals, such as profit, production, or reliability. Organizations are also seen as "open" in that they are constantly interacting with the outside environment, bit influencing and being influenced by broader social and political forces.

March and his colleagues, Michael Cohen and Johan Olsen, developed an extremely important perspective on organizational behavior, with the infelicitous title of the "garbage can model," which seeks to explain how complex organizations make decisions under conditions that differ radically from those that reign under rational models. According to the model, such "organized anarchies" exhibit three general properties. First, instead of having clear and consistent objectives, "the organization operated on the basis of a variety of inconsistent and ill-defined preferences." Different individuals at different levels of the organization may hold conflicting goals; the same individuals may hold different and incompatible goals at different times; organizations may not even know their preferences until after choices are made. Second, such organizations use extremely "unclear technology" in their operations: "Although the organization manages to survive and even produce, its own processes are not understood by its members." The organization's left hand does not know what the right hand is doing; what happened in the past and why it happened is not clear; and the connections between the organization's actions and the consequences of its actions are obscure. Third, there is extremely "fluid participation" in the organization's decision-making process. Participants come and go; some pay attention, while others do not; key meetings may be dominated by biased, uninformed, or even uninterested personnel.

This is a deliberately provocative theory of organizational behavior and Cohen, March, and Olsen acknowledge that such extreme properties will not be found at all times in all organizations. But they insist that the theory explains many of the actions of some organizations and some of the actions al almost any organization. This approach replaces the concept of rational organizations pursuing clear and consistent goals, with a more political vision in which "solu-

tions" are actively looking for problems to attach themselves to, "problems" are ill-defined and often unrecognized, and "participants" have limited attention, shifting allegiances, and uncertain intentions. This complex mixture is often haphazardly dumped together at "choice opportunities" – such as budget conferences or hiring committee lunches, and board of directors meetings (such choice opportunities are the garbage cans in the title) – during which the "organization is expected to produce behavior that can be called a decision." The result is a very different way of conceptualizing organizational behavior:

> Such a view of organizational choice focuses attention on the way the meaning of choice changes over time. It calls attention to the strategic effects of timing, through the introduction of choices and problems, the time pattern of available energy, and the impact of organizational structure… Such a theory of organizational decision-making must concern itself with a relatively complicated interplay among the generation of problems in an organization, then deployment of personnel, the production of solutions, and the opportunities for choice."

The garbage can theory has become very influential since it was developed in the early 1970s. By not treating organizations as rational, problem-solving actors, the theory has illuminated the hidden and, more capricious aspects of organizational life. It has encouraged scholars to examine how conflicting goals emerge and coexists, how organizational actors use power and information to promote their favored solutions to problems, how the rigid structure of an agenda can dominate the uncertain intent of a decision-maker, and how organizational leaders reinterpret the haphazard events of history to fit their preconceptions that the results must have been what they intended all along. Ideas borrowed from this somewhat "deconstructionist" approach to organization theory have been utilized to illuminate complex decision-making behavior in organizations as diverse as the Physics Department of the University of Oslo, the National Institute of Education; secretive Pentagon procurement agencies, the U. S. Congress, and a variety of New York state disaster relief agencies. It has also strongly influenced the most important work of the normal accidents school studying how modern organizations manage and mismanage hazardous technologies.

Structure, Politics, and Accidents

In an important 1977 essay, Charles Perrow argues that academic theorists often construct models of organizations whose behavior is far more rational and effective than that displayed by complex organizations in the real world. He therefore urged scholars to be more skeptical and conduct "revelatory" studies of "gross malfunctioning" organizations. Such studies, Perrow maintained, would be "more likely to reveal what most managers know but social scientists cannot afford to acknowledge, namely that complex social systems are greatly influenced by sheer chance, accident, and luck; that most decisions are very ambiguous, preference orderings are incoherent and unstable, efforts at communication and understanding are often ineffective, subsys-

tems are very loosely connected, and most attempts at social control are clumsy and unpredictable. In his influential 1984 book, *Normal Accidents:; Living with High-Risk Technologies*, Perrow followed his own advice and used and modified many of the ideas found in the garbage can model in an effort to understand the safety risks in such hazardous systems as commercial airlines, nuclear power plants, international shipping, the petrochemical industry, and (in a very brief section) nuclear weapons. His pessimistic conclusion – that "serious accidents are inevitable, no matter how hard we try to avoid them" – sharply contrasts against the optimism displayed by the high reliability theorists.

What are the assumptions and arguments that lead to this conclusion? Compared to the high reliability approach, the normal accidents theory is both more structural and more political. It is more structural because Perrow identifies two specific structural characteristics of many organizations operating dangerous technologies – "interactive complexity" and "tight-coupling" – which make them highly accident prone regardless of the intent of their leaders or operators. The theory is also more political because it focuses attention on the interaction of conflicting interests both within these organizations and between the organizations and the broader political community. Such conflicting interests can exert a strong influence on the frequency of catastrophic accidents, on their interpretation and therefore who receives the blame for failures, and, finally, on the degree to which the organizational structures that make normal accidents inevitable are modified and abandoned.

Complex and Linear Interactions

What does Perrow mean when he writes that an organization or technological system displays high degrees of interactive complexity and tight coupling? Interactive complexity is a measure, not of a system's overall size of the number of subunits that exist in it, but rather of the way in which parts are connected and interact. According to Perrow, "complex interactions are those of unfamiliar sequences, unplanned and unexpected sequences, and either not visible or not immediately comprehensible." The opposite is a system with linear interactions, "those in expected and familiar production of maintenance sequence, and those that are quite visible even if unplanned." An automobile assembly line is the prototypical linear system: it may be very large, but equipment is spread out; production steps are segregated and take place in a planned and familiar sequence; the assembly process is relatively well understood and if a broken conveyer belt stops the line, the problem is visible to operators and is, relatively speaking, easily comprehensible; and, finally, feedback and information on the what is happening on the plant floor is usually direct and simply verified.

None of these characteristics are found, however, in a nuclear power plant, the prototypical system with high interactive complexity. Critical components are kept, by necessity, in close proximity within a containment building, increasing the possibility of unplanned interactions. The nuclear energy production process is not a set of largely independent and serial steps, but rather requires many coordinated actions by numerous mechanical components and operators.

Despite years of operation, not all aspects of nuclear physics are completely understood. (Experts disagreed, for example, about whether the zirconium and water outside the fuel rods could interact under extreme heat and produce dangerous hydrogen bubbles until the accident at Three Mile Island proved that this was possible.) Finally, power plant operators cannot directly observe all the components involved in the production process. Many critical components and safety valves are inside the containment building or along the maze of pipes used for cooling purpose; control room operators must therefore rely on numerous (and fallible) warning devices, control panel lights that indicate whether components are functioning properly, and redundant monitoring systems to manage operations inside the plant.

Organizations and system with high degrees of interactive complexity – which include universities, biotech firms,, and NASA space launch missions as well as nuclear power plants – will share a number of problems, according to Perrow. They are likely to experience unexpected and often baffling interactions among components, which designers did not anticipated and operators cannot recognize. They are highly vulnerable to common-mode failure: the sometimes deliberate, but usually inadvertent condition where critical components share a common feature, the failure of which causes them all to break down. Designers and operators in hazardous organizations may work hard to anticipate and fix all likely potential problems, but Perrow's case studies repeatedly discover that it is the unlikely problem, even the bizarre and often banal failure, that initiates a normal accident. A commercial airplane's coffee machine causes a fire that shorts out the wires controlling both the aircraft's warning lights and its landing flaps. A maintenance worker changing a light bulb at a nuclear power station accidentally drops it onto some sensors and controls, causing a short circuit and the automatic scramming of the reactor. At Three Mile Island, a critical warning light properly indicated that the emergency feedwater valves were not open during the emergency, but the light could not be seen by control room operators because it was partially covered by a maintenance tag. Such unanticipated and even freakish incidents are inevitable in organizations with high interactive complexity. As Perrrow puts it:

> The argument is basically very simple. We start with a plant, airplane, ship, biology laboratory, or other setting with a lot of components (parts, procedures, operators). Then we need two or more failures among components that interact in some unexpected way. No one dreamed that when X failed, Y would also be out of order and the two failures would interact so as to both start a fire and silence the fire alarm. Further, no one can figure out the interactions at the time and thus know what to do. The problem is something that just never occurred to the designers.

Tight and Loose Coupling

Although interactive complexity will increase the likelihood of such bizarre and dangerous incidents, the second structural condition of tight coupling is necessary to produce escalation to a full-blown normal accident. Whether a system is tightly coupled or loosely coupled affects its ability to recover from small-scale failures before they cascade into larger problems. Perrow has

identified a number of related characteristics that distinguish tightly from loosely coupled systems. First, tightly coupled systems have more time-dependent processes: planned and unplanned interactions occur quickly; items must move continuously through the production process; delays, extensions, and storage of incomplete products are not possible. In contrast, in loosely couple systems, production moves more slowly, it is possible to put the system on a stand-by mode, and delays are possible because unfinished products can sit for a while or can be stored without damage. Second, in tightly coupled systems the sequences and coordinated activities need to produce the products are invariant: there is only one way to make the item and each step in the process must be taken in a sequence. In loosely coupled systems, the products can be produced in a number of ways, items can be rerouted and schedules changed, and parts can be added later if delays or shortage occur. Third, tightly coupled systems have little slack: quantities used in production must be precise and the process must be done right the first time or not at all. Such precision is not required for success in loosely coupled systems, and one part of the production process can always be repeated if necessary. Fourth, safety devices, redundancies, and buffers between parts of the production process in tightly coupled systems are largely limited to those that have been planned and designed into the system. Because of the time-dependent process, invariant production sequences, and lack of slack in these systems, there is little opportunity to improvise when things go wrong, and fortuitous recovery aids rarely have time to emerge. Adequate mechanisms for safety and recovery must therefore be deliberately built into the system. In loosely coupled systems, one finds much more opportunity for unplanned, but nonetheless successful, responses to individual failures.

Organizations that display high interactive complexity can be either loosely or tightly coupled. An example of each combination can provide a better understanding of the general characteristics that have just been listed. A large university, on the on hand, shows all the signs of high interactive complexity, but is also very loosely coupled system. Complex and unexpected interactions abound in the daily activities of students, staff, and faculty members living and working together on a campus. As college presidents undoubtedly know, no one can observe every activity that takes place on campus, much less plan or control them all, and the existence of bizarre, small-scale failures in the education process are commonplace. (It should be no surprise that universities have provided some of the best examples of the garbage can model.) No one thought that the professor who went to Los Angeles to give a talk to an alumni group would have her return flight canceled and therefore not show up for the semester's first lecture. No one anticipated that the secretary would forget to send a philosopher's course announcement sheet to either the department chairman or the course catalog office, leading to the course not being taught during the spring semester. Who would have guessed that the undergraduate entertainment committee would schedule a Grateful Dead concert on the afternoon of the Music Appreciation 101 midterm exam?

Such incidents do not lead to an accidental end of a student's education, however, because universities are loosely coupled systems. The learning process is not terribly time dependent, and the professor's lecture can be rescheduled for the next week without much damage to the

class. The steps involved in greeting a degree are usually very flexible, and it does not matter very much whether a student takes the philosophy course this semester or the next. There is a lot of slack in the system: the midterm could be offered the next morning; a good student might skip it altogether and rely on getting an A on the final in order to pass; a poor student could always take the class again next year. Finally, improvisation or fortuitous events could save the day: a smart undergraduate could run to the classroom and take the exam during the drum solo, or the whole concert might be canceled because of rain. In short, small failures happen all the time in universities, but they rarely escalate to the point where a student's education is irreparably harmed.

Contrast this situation, on the other hand, with a nuclear power plant, which has high degrees of both interactive complexity and tight coupling. If something goes seriously wrong due to a bizarre interaction—imagine that a forgetful maintenance worker accidentally leaves a coolant pipe valve closed and that he also forgets to fix the broken warning light indicator for that vale in the control room—the tightly coupled system will not so easily recover. Nuclear energy production is a highly time-dependent and very precise process: there has to be continual and sufficient amounts of coolant moving through the reactor to avoid dangerous overheating. The process is invariant and operators cannot substitute air circulation for the water coolant or wait until tomorrow to extract the excess heat. The opportunity for successful improvisation during an emergency is also highly limited. This is precisely why there must be redundant coolant pipes, backup power sources, and multiple warning indicators carefully designed into nuclear power plants.

It is this particular mixture that produces Perrow's pessimism. If a system has many complex interactions, unanticipated and common-mode failures are inevitable; and if the system is also tightly coupled, it will be very difficult to prevent such failures from escalating into a major accident. Such accidents may not happen often; indeed, they may be rare. But this is of little consolation if one is talking about highly hazardous technologies such as nuclear reactors, toxic chemical plants, or recombinant DNA (gene-splicing) research. If significant degrees of interactive complexity and tight coupling are combined, according to the theory, the organization's structure will eventually lead to a serious accident.

Four Factor Revisited

The high reliability theorists identified four factors that should produce extremely safe operations, even in organizations operating hazardous technologies. How are these specific safety factors viewed from the perspective of a normal accidents theorist? Can such conditions exist, and would they significantly reduce or even eliminate the dangers inherent in hazardous technologies?

In each case, the normal accidents approach provides reasons to suspect that significant improvements in safety will be much less forthcoming than the high reliability theorists suggest. Perrow's *Normal Accidents* directly address some of these factors, but in other cases I have been forced to deduce additional propositions, which are not explicitly raised by Perrow, from the

general causal concepts present in the theory. In addition, research by other scholars who have been influenced by Perrow's ideas has deepened and broadened the normal accidents theory by applying it to further case studies and new substantive areas. This literature will therefore also be utilized to develop propositions and illustrate concepts.

Conflicting Objectives

The first condition identified by the high reliability theorists was the safety is considered to be the priority objective by leaders of the organization. Normal accidents theorists do not disagree that I can be beneficial to have political elites and organizational leaders place higher priority on the objectives of safety and reliability. Error rates might be reduced if more money is spent in the right places and if clear operational safety goals are set. In fact, Perrow places considerable emphasis on the unnecessary risks caused by the leaders' and elites' lack of interest in improving safety in a number of high-risk technologies, arguing that "the nature of the victims in contact with the system should have some effect on the safety of that system":

> Elites fly on airplanes all the time, and airline pilots are in high demand, well paid, and in a position to have some, though not great, influence on the operation of the system. Captains and Admirals cannot escape naval vessels and a serious aircraft accident aboard a carrier will endanger their lives. But nobody of any great importance works directly in a nuclear power plant, travels aboard tankers and freighters loaded with explosives and toxic cargos, sits in the potato fields sprayed by genetically engineered microbes, or gets very close to a huge chemical plant.

Yet, even though Perrow argues that increased leadership interests in safety can have some beneficial effects, the approach also leads one to expert that the goals of elites and organizations' leaders will only have a quite limited effect on the behavior of the entire organization. The structural factors of interactive complexity and tight coupling are seen as significantly increasing the likelihood of accidents in such organizations, regardless of the intent of the leaders. In addition, because complex organizations are conceived as having poorly defined and often inconsistent preferences, this perspective suggests that the pursuit of other objectives will continue, and conflict over organizational goals will remain, even if improved safety is recognized as the priority official goal by leaders.

This continued conflict over goals within high-risk organizations is significant in three major ways. First, significant pressure to maintain high production rates exist in most hazardous industries and may be only slightly moderated by increased interests in safety. The signs of such production pressures – hasty decision-making, violations of safety rules, and jerry-rigged procedures – are found during many normal accidents. Subtle and sometimes not so subtle reminders of the "need" to keep the ship on schedule, for example, often lead to captains and crews "cutting corners" and hence to accidents at sea. Diane Vaughan's study of the space shuttle *Challenger* incident similarly emphasizes the degree to which, despite leadership's great concern about safety,

NASA continued to provide strong pressures to its contractors and officials for immediate delivery of parts and strict maintenance of launch schedules. Second, even if political leaders desire increased safety, differences in prioritization of goals between an organization and its ostensible political authority can continue. Organizations will seek to maintain their autonomy against outside pressures and numerous studies have noted the degree to which tensions exist between an industry and a regulatory agency created by the political authorities to monitor its behavior. Under such conditions, the normal accident approach focuses attention on the likelihood of government watchdogs being "captured" by the forces they are supposed to regulate. Third, even if political elites and organizational leaders have consistent objectives favoring safety, they may be misinformed about the nature or frequency of dangerous operation by lower-level operators, whose interests include keeping their jobs and therefore not getting caught when rules are violated. Michael Tamuz's innovative research on the reporting of near-accidents in the U. S. airline industry, for example, discovered that commercial pilots consistently underreported the number of safety violations they have witnessed. Raymond Levitt's study of the California construction industry similarly showed how the incentives to withhold information about small accidents and near misses increase when concerned company officials offer rewards to employees with the best safety records.

In short, although the normal accidents theory accepts that there can be "organizations with an avowed goal of safe operations," it treats the step from avowed goals to real organizational objectives as very problematic. The theory points to a number of reasons why safety and reliability may not increase even if political and organizational leaders claim to have placed increased priority on such goals. Even the serious intent by an organization's leadership to accord safety and reliability higher priority than other objectives is not seen as automatically leading to the acceptance of such goals by others within the organization. Inconsistent goals and conflicting interests remain and can increase the risks of accidents in important ways.

The perils of Redundancy

Given that individual parts of a complex technological system can always fail, the high reliability theorists place great emphasis on the need for redundancy – duplication and overlap in critical components and personnel – to produce improved safety. Accidents, of course, can still occur. It is always possible that more than one component can fail independently at the same time: during the 1984 Bhopal plant chemical disaster, for example, the deadly gases escaped because three separate safety devices – a flare tower to burn off the gas, a scrubber to clean air emissions, and a water sprinkler system to neutralize remaining fumes – all failed simultaneously. Such an event, however, does not directly contradict the high reliability theory. Indeed, it supports the theory, since simultaneous and independent failure of all critical components should be made less likely when more redundant backups are added to the system.

The normal accidents perspective, in contrast, focuses attention on the potential negative consequences of redundancy and, indeed, maintains that adding redundant parts to a complex tech-

nological system can increase the likelihood of accidents for three central reasons. First, while these authors do not disagree that redundant safety systems, if truly independent, can enhance system reliability in theory, they do suggest that, in reality, redundant systems are often less independent than their designers believe. Although redundant backups are supposed to be independent of one another, they also usually increase the interactive complexity of high technology organizations and hence lead to unanticipated common-mode failures. A number of nuclear power plant accidents, for example, have been caused by recently added safety devices that interacted with other components in bizarre and unexpected ways to produce critical failures. The 1966 near meltdown of a power plant near Monroe, Michigan, was produced by a broken safety device that blocked the flow of coolant to the reactor, and the 1986 Chernobyl accident was caused by a test of backup safety power sources. Second, adding redundancy often makes the system more opaque: individual component or human failures will often be less visible, since they have been compensated for by overlap or backup devices. If this occurs, such individual errors may remain unfixed, and latent problems will therefore accumulate over time, increasing the likelihood of both interactive common-mode failures and simultaneous failures of independent and redundant components. Third, when redundancy makes the system appear more safe, operators often take advantage of such improvements to move to higher and more dangerous production levels. "Fixes, including safety devices," Perrow argues, "often merely allow those in charge to run the system faster, or in worse weather, or with bigger explosives."

Decentralization, Culture, and Continuity Revisited

From the perspective of a normal accidents theorist, doubt can also be raised about the three operational characteristics that high reliability theorists identified as being essential for safety. Decentralization of decision-making authority to individuals close to the operations may well be necessary to produce appropriate low-level responses to unanticipated events in systems that display high degrees of interactive complexity. Yet, Perrow also argues that the characteristics of highly centralized decision-making -- strict standard operating procedures, unquestioned individual obedience and immediate responses – are equally necessary in tightly coupled systems. Because tightly coupled systems have invariant and time-dependent production processes with little built-in slack, there is no time for improvisation by operators in crises and little scope for recovery from faulty field-level decisions. When a system is both interactively complex and tightly coupled, therefore, the requirements for reliability are simply incompatible. Such system are "organizational Pushmepullyus, straight out of the Dr. Doolittle stories, trying to go in opposite direction at once." Neither strict centralization nor decentralization can therefore ensure safety under such conditions.

High reliability theorists argue that a solution to this dilemma is a strong organization culture, which creates a common set of decision-making premises and assumptions so that lower-level personnel, responding to surprises according to their own best judgment, nevertheless react in a uniform and appropriate manner. A normal accidents perspective, however, is accordance with

garbage can mode assumptions, would lead to doubts about whether organizational leaders know enough about their operations and technology to determine how low-level authorities should respond in all contingencies. Uniformity of decentralized decision may be helpful if the leadership has accurately identified how to prevent accidents; if it has not done so, however, then uniformity of low-level responses will decrease safety and reliability. In addition, the intense organizational culture, which is viewed as necessary by the high reliability theorists, is seen as being simply impossible to achieve in an individualistic, democratic society. As Perrow argues:

> It calls for a wartime, military model in a peacetime civilian operation. A military model reflects strict discipline, unquestioning obedience, intense socialization, and isolation from normal civilian life styles… Efforts to extend this model to industry in the nineteenth and early twentieth centuries failed; it was too incompatible with American social values and culture.

In short, citizens in democratic societies are not willing to manage all hazardous organizations as if they were "total institutions," completely isolating their members from broader social life, socializing them into conformity, and controlling every aspect of their behavior.

Continuous training and operational experience helps develop and maintain the information necessary for maximum safety, according to the high reliability school. Analysts with a normal accidents perspective might agree in principle, but would also point out that an organization simply cannot develop operational experience with reactions to all the unanticipated and undesirable failure modes produced by high interactive complexity in hazardous technologies. Potential accident scenarios that have not been imagined obviously will not be practiced. Some dangerous operations cannot be practiced often, precisely because they are so dangerous. Finally, in the politicized environment of a conflicted organization, some highly unpalatable accident scenarios, especially those that clearly place blame on certain individuals or subunits, will not be addressed because to do so would require that those responsible implicitly acknowledge that such events are possible.

Restrictions on Learning

In contrast to the belief that high-risk organizations will improve safety and reliability through trial-and-error learning and simulation, the normal accidents school suggest a number of reasons why effective learning will be very difficult, if not impossible. In the first place, the causes of accidents and near-accidents are often unclear, and in such situations even well meaning organizational leaders are likely to reconstruct history to conform to their preconceptions, to attribute success to their own actions, and to develop lesion to fit into their sense of mission. Unless the causes of an incident are relatively clear and cannot be ignored, biased interpretations will therefore reign. Under severe conditions of ambiguity, instead of learning more effective behavior, March and Olsen have argues, "modern organizations develop myths, fictions, legends, folklore, and illusions."

Second, analyses of real or potential accidents often take place in extremely politicized environments in which blame for failures and credit for successes must be assigned to someone within the organization. If this is the case, such analyses are likely to be designed to protect the parochial interests of the most powerful actors instead of promoting more objective learning. These powerful interests help explain why, for example, internal investigations of industrial disasters almost always find that they were caused by "operator errors" and were rarely the result of mistaken design of faulty designs by the senior management. In such situations, organizational leaning, in the sense of changing procedures to avoid repeating errors, would be expected to be highly constrained and severely limited in the types of lessons that can be accepted. Third, normal accidents theory highlights the possibility that faulty reporting will make it extremely difficult for organizations to assess their performance accurately. Unfortunately, the incentives to cover up serious errors and to fabricate positive records can be strong. Especially if filed-level operators of lower-level officials know that they will be blamed for the occurrence of any accidents or near-accidents. It therefore should not be surprising that faulty reporting has been witnessed in a variety of organizations in which it was not in the interests of lower-level individuals to acknowledge the truth. When the Federal Aviation Administration installed automatic flight path recording devices in air traffic control centers, the number of "near miss" collisions reported by pilots increased enormously. When U. S. Air Force pilots over North Vietnam bombed a Soviet naval vessel in error, the evidence was quickly covered up. After the December 1989 invasion of Panama, the secretary of defense inaccurately told the press that the new F-117A Stealth fighter had delivered its weapons with "pinpointing accuracy," because the commander of the Tactical Air Force command failed to report that the bombs has accidently been dropped away from the intended targets.

In addition, organizational leaders who are informed of improper activities by subordinate often have an interest to cover up the mistakes so as to minimize outside criticisms of the organization's management. "Shielding members from external perception of their deviance also shields the organization from embarrassment," argues Jack Katz: when an official is fired, it is often represented as voluntary resignation; commanding military officers have been found to list AWOL soldiers as being on official leave; and corporation officials have been known to resist prosecution of discharged embezzlers in order to protect their image with stockholders. While such shielding may enhance the organization's reputation with outsiders, it also decreases the incentives, which often stem from external pressures, to improve procedures after errors are discovered.

Finally, a fourth constraint on organizational learning is secrecy inside complex organizations and between organizations. Although compartmentalization of knowledge within an organization may be necessary to guard secrets from outsiders, it can also limit the overall organization's ability to learn and implement lessons from historical experience. It is not uncommon in such circumstances for one part of the organization to be unaware of what other subunits have learned from trials and errors. Secrecy between organizations is also a common practice since it protects production methods or procedural innovations. This also has the negative effect, however, of

limiting vicarious learning, that is, learning from the mistakes of others.

Normal Accidents Summary

Normal accidents theorists take a natural open systems perspective in which organizations and members of organizations are self-interested actors with potentially conflicting interests, and in which organizations are strongly influenced by broader political and social forces in the environment. These theorists utilize a modified form of the garbage can model: they view organizations as often having inconsistent preferences, unclear technologies, and fluid participation. The theory predicts that serious accidents are inevitable if the organizations that control hazardous technologies display both high interactive complexity (which produces bizarre and unanticipated failures) and tight coupling (which causes the failures to escalate rapidly out of control). Each of the four factors previously identified as contributing to high reliability are seen from the normal accidents perspective as being ineffective, unlikely to be implemented, or even counterproductive. Even if leaders place a very high priority on safety and reliability, which is by no means a given, competing organizational and individual objectives will remain: the continuing desires to maximize production, maintain autonomy, and protect personal reputations, however, can severely impair efforts to improve safety. Adding redundancy does not necessarily enhance reliability because it also increases interactive complexity, encourages operators to run more risks, and makes the overall system more opaque. Decentralized decision-making does not necessarily enhance safety because tightly coupled systems require rapid responses and strict adherence to standard operating procedures. Intense socialization and a strong organization culture are unlikely to be very productive in hazardous organizations both because their leaders cannot know how operators should respond in all contingencies and because democratic societies are unwilling to isolate and control all aspects of the lives of such organizations' members. Constant training and practice will not address accident scenarios that are unanticipated, excessively dangerous, or politically unpalatable. Finally, a number of factors will severely constrain the process of trial-and-error learning: uncertainty about the causes of accidents, the political interests and biases of organizational leaders and low-level operators, and compartmentalization within the organization and secrecy between organizations.

Although there has previously been no thorough study of nuclear weapons and command control systems from the normal accidents perspective, the theory would not lead to optimistic predictions about safety with nuclear weapon systems. A normal accidents theorist would look with alarm at the global command and control system that manages U. S. nuclear weapons operation, sine it appears to be highly complex, by necessity, and appears to be tightly coupled, by design, so as to perform prompt retaliation in a war. Indeed, from a normal accidents perspective, the fact that there has never been an accidental nuclear weapons detonation or an accidental war is surprising.

These two theories explaining the origins of accidents in complex organizations are clearly competitive in a number of ways. At the most general level, they represent very different perspectives on how modern organizations function. May of the specific conditions that the high reliability theorists argue will promote safety will actually reduce safety according to the normal accidents theorists. Other strategies that the first school argues are both necessary for safety and possible to implement in complex organizations are seen as simply impossible to achieve by the second school. Finally, the central conclusions reaches by the two groups of scholars are very different from one another. One theory leads to the highly optimistic prediction hat near-perfect reliability and safety are possible in complex organizations using hazardous technologies. The other theory leads the more pessimistic prediction that serious accidents are inevitable. Table 1.1 outlines the contradictory assumptions, prepositions, and conclusions drawn from both schools of thought.

Table 1.1
Competing Perspectives on Safety with Hazardous Technologies

High Reliability Theory	Normal Accidents Theory
Accidents can be prevented through good organizational design and management	Accidents are inevitable in complex and tightly coupled systems
Safety is the priority organizational objective	Safety is one of a number of competing objectives
Redundancy enhances safety: duplication and overlap can make "a reliable system out of unreliable parts"	Redundancy often causes accidents: it increases interactive complexity and opaqueness and encourages risk-taking
Decentralized decision-making is needed to permit prompt and flexible field-level responses to surprise	Organizational contradiction: decentralization is needed for complexity, but centralization is needed for tightly coupled systems
A "culture of reliability" will enhance safety by encouraging uniform and appropriate responses by field-level operators	A military model of intense discipline, socialization, and isolation is incompatible with democratic value
Continuous operations, training, and simulations can create and maintain high reliability operations	Organizations cannot train for unimagined, highly dangerous, or politically unpalatable operations
Trial and error learning from accidents can be effective, and can be supplemented by anticipation and simulations	Denial of responsibility, faulty reporting, and reconstruction of history cripples learning efforts.

What lies at the heart of the disagreement between these two perspectives on safety in hazardous organizations? The opposing schools do not dispute the basic logic of Perrow's argument that the structural conditions of interactive complexity and tight coupling should, in theory, lead to accident-prone organizations. They do, however, hold very different visions of the degree to which the basic forces of human "agency" – the elements of culture, design, management, and

choice – can counter or compensate for these dangerous structural pressures within hazardous organizations. Perrow's basic pessimism on this issue is unmistakable: "No matter how hard we try, no matter how much training, how many safety devices, planning, redundancies, buffers, alarms, bells and whistles we build into our systems, those that are complexly interactive will find an occasion where the unexpected interaction of two or more failures defeats the training, the planning, and the design of safety devices." In direct contrast, the high reliability theorists argue that certain complex, tightly coupled "organizations have developed strategies for avoiding the negative effects of these (structural) characteristics."

This position is more clearly laid out in the Berkeley group's studies of U. S. aircraft carriers. Weick and Roberts, for example, maintain that U. S. carriers "tend to be tightly coupled and interactively complex, a combination which normally increases the tendency toward consequential failures but in these organizations does not." Rigorous exercises, continual training, and realistic simulations combat the negative effects of interactive complexity since such efforts are "directed to understanding the complexities of the technologies aboard the ships…in such a way that baffling interactions do not occur." Redundant communication systems also play a key role because "indirect sources of information that can lead to baffling perceptions are meditated against by many direct information sources"; indeed, "in one important decision making place aboard ship, 20 different phone are used by one person." Personnel redundancy is similarly found to counter the negative effects of tight coupling. Roberts, for example, argues that having many officers responsible for checking the aircraft's landing gear can "decompose the tight time frame that are part of tight coupling" since "in a short time three pairs of eyes should be able to spot a problem that may take one pair of eyes longer to detect." Finally, it is argues that continuous training and a strong "culture of responsibility" aboard U. S. aircraft carriers enables "simultaneous centralization and decentralization" to exist, which overcomes the contradictory organizational pressures caused by complexity and coupling. When such intense socialization occurs, according Weick, "there are previously shared values concerning reliability which then allow for coordinated, decentralized action." For the high reliability theorists, if these conditions are met, a highly optimistic prediction of near-perfect safety is warranted: "Parts of these systems do fail", Roberts concludes in a critique of Perrow's book, but "it is really not clear that all high-risk technologies will fail."

Evaluating the Theories

Which perspective on hazardous organizations is more accurate? It is not an easy task to evaluate these two theories. Both have relatively plausible assumptions and the prepositions and conclusions of both theories appear to flow logically from those assumptions. Moreover, both the high reliability and the normal accidents theorists provide numerous hypothetical and empirical examples to illustrate and support their arguments. So how can one assess the general strength of the two theoretical approaches?

The most serious difficulties should be acknowledged from the start. Because neither theory

provides a precise estimate of the likelihood of serious accidents with hazardous technologies, it is impossible to determine the precise number of accidents, which, if discovered over time, would support or weaken the theories. Members of the high reliability school are usually prudent enough to avoid the extreme claim that organizational perfection is possible. Marone and Woodhouse therefore argue that "there is a good chance… that catastrophes will be prevented"; Roberts only states that there are "hazardous organizations that engage in nearly error free operations"; and La Porte and Consolini only claim that such organizations have "a very low error rate and an almost total absence of catastrophic failure." Similarly, Perrow only claims that disastrous accidents will occur eventually: "Accidents are inevitable and happen all the time, serious ones are inevitable though infrequent; catastrophes are inevitable but extremely rare." Such imprecise language suggests that the two theoretical schools have a common estimate about the probability of dangerous accidents despite the strong difference in the tone of their conclusions: Perrow may look at a glass of safety and find it 1 percent empty; high reliability theorists may see the same glass of safety as 99 percent full.

Yet, when one considers that causal mechanisms involved in the theories – the specific factors outlined in table 1.1 that each theory claims will lead to or prevent organizational accidents – the contradictions between them become more clear. This suggests that a more valuable test of the theories would involve detailed specific historical case studies in which a range of these factors existed, focusing on whether these factors had the predicted influence on safety and reliability. For example, what effect did leaders placing a high priority on safety have on the beliefs and behavior of the rest of the complex organization? What was the impact of adding redundant safety devices? Did effective organizational learning take place, fixing the cause of safety problems, after serious incidents? Or did faulty reporting, denial of responsibility, and the reconstruction of history occur?

The following chapters attempt to evaluate these two theoretical approaches by applying them in a competitive manner to a new subject area. In these chapters I will deduce what each theory should predict about specific efforts to prevent the ultimate safety system failure – an accidental nuclear war – and will then compare these predictions to the historical evidence of U. S. nuclear weapons command and control. Which theory provides better predictions of what happened and more compelling explanations of why it happened? Which theory leads to the discovery of more novel facts and new insights? Which one is therefore is a better guide to understanding?

This is bound to be a very imperfect exercise in theory testing for a number of unfortunate, but unavoidable, reasons: The imprecision of each theory's predictions, the limited information available about the history of U. S. nuclear weapons safety and military operations, and the inability to control for the effect of perturbing variables all conspire to make this a less than fully satisfying test. I believe, however, that despite such imperfections, these theoretical perspectives will be made more credible to the degree that they lead to expectations that are found to be historically accurate. To the degree that empirical findings do not fit such predictions, the claims of the theory in question will be weakened.

A Tough Test for Normal Accidents Theory

It is critical that social science theories be tested competitively against one another: since no theory is perfect, its relative strengths and weaknesses can only be measured in comparison to the best available alternative. It is also important, however, to subject theories to tough tests, "least-likely" cases in which the theory's applicability is not obvious or well-accepted. Under such conditions, as Kenneth Waltz notes, "if we observe outcomes that the theory leads us to expect even though strong forces works against them, the theory will begin to command belief."

In this respect, a study of the command and control of U. S. nuclear weapons and the prevention of accidental nuclear war is a tough test fro the normal accidents theory. At first glance, one would not expect that the normal accidents theory would provide a strong explanation for nuclear weapons safety. Indeed, there are four important reasons to suspect, at his initial point, that high reliability theory will prove a better tool for studying this subject.

The first reason should be obvious: there has never been an accidental nuclear war, nor even a single accidental detonation of a nuclear weapon. This impressive safety record with nuclear weapons therefore appears to conform more closely to the optimistic predictions of the high reliability theorists. Indeed, even Perrow is optimistic about the possibility of high reliability with respect to at least one important aspect of nuclear weapons safety, the risk of an accidental war caused by false warnings. "There is much to fear from accidents with nuclear weapons such as dropping them or an accidental launch, but with regard to firing them after a false warning we reach a surprising conclusion, one I was not prepared for: because of the safety systems involved in a launch-on-warning scenario, it is virtually impossible for well-intentioned actions to bring about an accidental attack." If this research discovers evidence that false warning problems have been more serious than Perrow recognized, his theory will thereby be given even more support than he provided for it.

The second reason why this subject is a tough test for normal accidents theory is that the prevention of an accidental nuclear war is undoubtedly a high priority objective of the senior leadership of the United States. Many hazardous systems, such as oil tankers, are far less safe than they could be simply because the individual s with the power to influence change are not potential victims of accidents and therefore do not place great importance on safety improvements. Nuclear weapons are not, however, discriminatory: political authorities, elites, and organizational leaders would suffer just as much as would local operators if there was an accidental nuclear war. As Perrow puts it, on the one hand, "elites do not sail on Liberian tankers"; on the other hand, "the nuclear weapons system is finally democratic." Because the president would be a victim (indeed, potentially one of the first victims) in a nuclear war, even a normal accident theorist would expect that nuclear weapons safety would be a high priority of all American presidents.

Third, there is one aspect of the problem of safety about which Perrow and the high reliability theorists appear to agree: isolation away from society, intense socialization, and strict discipline of organization members, as in the ideal military model, can enhance reliability and safety. Where Perrow disagreed with the high reliability theorists is over the degree to which this mili-

tary model can or should be applied to members of civilian hazardous organizations in a democracy. The military organizations that control U. S. nuclear weapons, however, have been designed to come as close as possible to that idealized military model. To the degree that the historical evidence suggests that such isolation, discipline, and socialization did not lead to enhanced safety, it will provide additional support for a pessimistic appraisal of the prospects for safety in other organizations.

Fourth, there is an unavoidable restriction in my selection of case studies to United States command and control problems because the secrecy surrounding this subject makes it virtually impossible to do the same kind of detailed historical study about the command and control of nuclear weapons in other countries. This practical necessity is a virtue of theory testing, however, since safety efforts cost a great deal of money, and one would therefore expect that the wealthiest nuclear power would also be the safest nuclear power. "Richer is safer," as Wildavsky states, all other things being equal. One would anticipate, for example, that just as U. S. nuclear power plants and commercial airlines are more reliable than their counterparts in the former Soviet Union, there should be better nuclear command and control safety in the United States. If serious problems with the nuclear weapons system are found in the U. S., one would anticipate finding them to an even greater degree in Russia and in other nuclear powers.

In short, this study is a tough test for the normal accidents theory. It makes the theory compete on the home turf of high reliability theory. If the normal accidents theory does well here, it will be shown to be evne stronger than previously estimated.

The Catch 22 of Close Calls

The methodology used in this study us quite simple in conception, although it is often difficult in execution. To determine which theory provides a more useful tool for understanding nuclear weapons safety, I identify historical situations in which the theories would provide contradictory predictions. For the high reliability theory, I identify the extent to which the factors that they argue lead to safety were in fact present and then examine whether those factors, in fact, had the predicted effect. For the normal accidents theory, I deduce where and when accidents should have been likely if the theory is correct, and then check the historical record for evidence of safety problems occurring in those specific areas.

It is important to emphasize that the mere existence of near-accidents is inadequate evidence either to diminish the credibility of high reliability theory or to increase the credibility of normal accidents theory. Indeed, there is an irony here that could be called the catch 22 if close calls: the more near-accidents I discover, the more it could be argues that the system worked, since the incidents did not in the end lead to an accidental nuclear war. The fact that individual errors will occur is, after all, the basic reason why redundancy and other safety measures are built into complex technological systems.

This is yet another reason why it is essential, in order to weight the relative explanatory power of the two theories, to focus on their specific predictions that are in conflict. What were the

causes of the incidents and the reasons why they did not escalate? For example, did adding more than one safety device prevent a serious accident, as would be suggested by high reliability theorists? Or was redundancy the cause of the problem, as predicted by normal accidents theory? And if redundancy caused the incident, what prevented it from escalating?