

In Organizational Encounters with Risk, Bridget Hutter and Michael Power, eds. New York and Cambridge: Cambridge University Press, 2004.

## ORGANIZATIONAL RITUALS OF RISK AND ERROR

Diane Vaughan

### Abstract

In this chapter, I compare organizational rituals of risk and error in two U.S. government agencies: The Federal Aviation Administration's National Air Traffic System (NATS) and the National Aeronautics and Space Administration (NASA). I take the position that all organizations, even those categorized as High Reliability Organizations, are subject to routine nonconformity: they regularly produce mistakes and errors. For both agencies, I describe their technologies of control, which are the rules and procedures, work practices, and surveillance technologies for regulating risk. Then I examine definitional processes: how anomalies are identified, tracked, and converted into formal organizational categories. Because of differences in the certainty of space shuttles and airplanes, the technologies of control employed as these two agencies identify, define, and control risk and error vary. This variation has important consequences. I show that the variation in these agencies' technologies of control produces different cultural understandings about risk and error. These cultural understandings have social psychological consequences, affecting how technical workers interpret signals of potential danger. Finally, I consider the implications for this comparison for organizational encounters with risk.

Organizational encounters with risk and error are not restricted to the sensational cases that draw media coverage when mistakes, near-misses, and accidents become public. They are, instead, a routine and systematic part of daily organizational life that only occasionally become visible to outsiders. Merton was the first to observe that any system of action can generate unexpected consequences that are in contradiction to its goals and objectives (1936, 1940, 1968a). Recent research affirms his observation: unanticipated events that deviate from organizational expectations are so typical that they are "routine non-conformity" - a regular by-product of the characteristics of the system itself (Vaughan 1999). The public only learns about the most egregious of these. Because routine non-conformity is a regular system consequence, complex organizations that use or

produce risky technologies may have encounters with risk daily.

In this paper, I compare daily encounters with risk for two organizations for which mistakes result in public failures and have high costs: the Federal Aviation Administration's National Air Transportation System (NATS) and the National Aeronautics and Space Administration's (NASA) National Space Transportation System (NSTS), otherwise known as the Space Shuttle Program. My comparison of these two agencies is grounded in two related strands of research. Barry Turner investigated the causes of 85 different "man-made disasters." Turner found an alarming pattern: after a disaster, investigators typically found a history of early warning signs that were misinterpreted or ignored. A problem that seemed well-structured in retrospect was ill-structured at the time decisions were being made (Turner 1978; Turner and Pidgeon, 1997). Turner did not have micro-level decision making data to explain how and why this could happen. Instead, he explained it skillfully by drawing on organization theories about information flows. However, no one had tested this theories with data about how people decisions and why they make the choices they did. The second strand is Star and Gerson's (1986) study of anomalies in scientific work. Star and Gerson point out that every anomaly has a trajectory during which it is subject to processes of definition, negotiation, and control that are similar to the response to anomalies in other types of work. They found that how anomalies are defined and negotiated depends upon the occupational context and the evaluation systems that have been developed to meet unexpected deviations in the work flow. Star and Gerson concluded that a mistake or anomaly is never defined in isolation, but always is relative to the local and institutional context of work (1986: 148-50).

These two strands of research led me to the following conclusions. First, Turner's discovery of warnings of hazards over long incubation periods preceding accidents suggested the importance of research on daily encounters with risk. Second, Star and Gerson's work indicated that an appropriate focus was the trajectory of anomalies - how they are identified, measured, and assessed for risk. I had a chance to investigate both these issues in my research on NASA's Challenger accident (Vaughan 1996). In the years preceding the accident, I found a history of early warning signs - signals of potential danger - about the technology of the flawed Solid Rocket Boosters that caused the disaster. These early warning signs were misinterpreted as decisions were being made, contributing finally to the disastrous outcome. Their seriousness only became clear in retrospect, after the tragedy. In contrast to the accidents Turner investigated, I had micro-level data on NASA

decision making and also macro-level data on the NASA organization and its relations in its political/economic environment. These data made it possible for me to explain how and why signals of potential danger were misinterpreted. The answer was "the normalization of deviance": a social psychological product of institutional and organizational forces. The consequence was that after engineering analysis, technical anomalies that deviated from design performance expectations were not interpreted as warning signs but became acceptable, routine, and taken-for-granted aspects of shuttle performance to managers and engineers. The trajectory of anomalies, as they were identified and their risk measured and assessed showed the importance of both the local organizational and institutional contexts of work. Air Traffic Control was the next logical research site for subsequent study because I wanted to locate myself in an organization where people successfully were trained to recognize early warning signs so that small mistakes did not become accidents. This research strategy contrasted with the retrospective analysis necessary in the Challenger analysis. The theoretical framework I use for the analysis of both cases is "situated action," which examines the link between institutions, organizations, and individual decisions and actions (Vaughan 1998; cf. Suchman 1987).

In the following pages, I examine the process of identification and definition of anomalies and the technologies of control in use at these two agencies. By technologies of control, I mean 1) rules and procedures that guide decisions about risk, 2) work practices, and 3) surveillance technologies consisting of both procedures and hardware designed to identify anomalies that are potential signals of danger.<sup>1</sup> I locate their encounters with risk within the institutional and organizational context of work. The comparison that follows is necessarily brief: I must omit much of the institutional and organizational context, as well as much of the rich ethnographic and interview data on which these studies are based. The NASA analysis is from a completed project (Vaughan 1996). The air traffic control analysis is one part of a project, still in progress. The research is based on ethnographic and interview data I collected during 11 months with air traffic controllers in four air traffic control facilities in the New England region during 2000-2001. For both agencies, their technologies of control are so extensive that even concentrating on one of these organization in a paper of this length necessarily would omit important aspects, so I will focus only on some of the more salient factors.

I begin by laying out some basic similarities and differences between the two kinds of space vehicles that the two agencies control. My purpose is to show how differences in the technologies of shuttles and airplanes affect the social organization of work and therefore, encounters with risk. Then, I examine the trajectory of anomalies in both organizations showing 1) the variation in the technologies of control employed by each agency as it attempts to identify, define, and control risk and error, 2) the symbolic use and meaning of their technologies of control, as they come to function in both agencies as organizational rituals of risk and error, and 3) how these ritualistic practices have social psychological consequences with significant affects on how technical workers interpret signals of potential danger. Finally, I conclude with some thoughts on the implications of this comparison.

#### LAUNCHING SHUTTLES AND AIRPLANES: THE SOCIAL ORGANIZATION OF RISK ENCOUNTERERS

Both the FAA's National Air Transportation System (NATS) and NASA's National Space Transportation System (NSTS) do space work: launching vehicles into space, monitoring them while in flight, and returning them to earth. Both have responsibility for high-risk technologies in which routine nonconformity may result in loss of life. Both organizations are government bureaucracies and large scale technical systems in which rules, procedures, and routines govern all facets of organizational life. Both are dependent upon Congress and the White House for funding; both are dealing with getting vehicles into space on a pre-established, schedule-driven, timetable over which the people doing the hands-on work have little control. During the periods of my research, both agencies were experiencing production pressure generated by the external environment. NASA was under pressure to launch a predefined number of shuttles every year in order to subsidize inadequate Congressional funding for the program with funding from payloads; air traffic control was under pressure to reduce delays, a situation arising from congested skies due to deregulation of the airlines, a thriving economy that had more people able to fly, and inability to sufficiently increase the number of airports and/or runways of existing airports in order to accommodate the increased demand. Any accident, under these circumstances, has negative consequences in addition to the tragic and irrevocable loss of life for both agencies.

For air traffic control, the threat of privatization hangs over all negative performance outcomes. For NASA, an accident or mistake could result in decreased funding or program cut-backs. Because of the extensive harmful consequences of mistake, both agencies invest considerable resources into technologies of control: their apparatus for encountering risk are well developed and central to everyday work, perhaps more so than many organizations where other types of risky work are done or where failures are less public or less costly.

Regulating risk tends to take two forms: a compliance strategy or a deterrent strategy of control (Reiss 1984; Hawkins 1984), although most organizations use some combination of compliance and deterrence strategies of control. A compliance strategy is forward-looking and preventive. It aims for early identification and measurement of deviations from the expected so an appropriate response can bring the technology back into compliance with the established parameters for system safety. A deterrent strategy is backward-looking, usually implemented after a mistake or accident, designed to prevent or reduce the probability of a similar event. Deterrent strategies are punitive, invoking sanctions that are directed toward the individuals or organizations responsible for an error or accident. A compliance strategy tends to be the strategy of choice when accidents are so costly that society (and the organization responsible) cannot afford them; the harmful consequences are so egregious that an after-the-fact deterrent strategy, on its own, simply will not do. For both the Space Shuttle Program and air traffic control, encounters with risk are a routine aspect of daily work that prioritize a compliance approach designed to prevent mistakes and accidents. Prevention is the goal: both use technologies of control designed to ferret out anomalies and correct them before small problems can turn into big ones.

Although both agencies do space work and have a lot in common, variation in the certainty of the two vehicles' technology translates into differences in the social organization of encounters with risk. NASA's Space Shuttle Program still is an experimental system, the shuttle itself an experimental technology. No two vehicles are ever exactly alike because technical changes are made to each vehicle after every flight. Air transportation, in contrast, is an operational system, the technology of aircraft highly standardized. As a consequence, time and division of labor are markedly different in encounters with risk at both places. The differences in certainty shape the time available for risk assessment, providing different windows of opportunity for identifying potential

signals of danger and correcting mistakes. The pre-launch process for both organizations (controllers in towers, like engineers at NASA, also talk about "launching" airplanes) reflect these differences. A shuttle flight has a long prelude. Planning, preparation, and assessment of the readiness of a shuttle for flight begin from 18 months to 24 months ahead. The final countdown, launch, and mission of a shuttle are a very small portion of the pre-launch process that has to be monitored for anomalies. Many anomalies are discovered and corrected during this long pre-launch assessment period. Moreover, shuttles, like airplanes, are reusable, so the post-flight engineering analysis, repair of damage from a flight, and technical adjustments to improve performance on the next shuttle flight are part of the pre-launch flight readiness assessment, which can take months.

The launch and flight of an airplane does not have this long prelude, though it may seem that way to passengers. The comparatively leisurely process of risk assessment in the Shuttle Program is in stark contrast to the condensed window of opportunity in air traffic control. The turn-around time of a flight is short, and any technical anomalies on the aircraft prior to takeoff are the responsibility of the airlines. Instead, the pre-launch process for an aircraft begins in an air traffic control tower when a computer-generated Flight Progress Strip, identifying the aircraft, its equipment, its point of origin and the proposed route to destination arrives at a controller's position, indicating the plane is coming into his or her domain of responsibility. Tower controllers monitor pre-flight plans, check planned routes, control all ground movement and take-off. The gate. In this very short pre-launch phase, controllers are making decisions about aircraft with people on board, whereas in the shuttle's nearly two-year pre-launch phase, lives are not yet at risk.

Differences in division of labor at the two agencies are as opposite as the allotment of time for identifying and correcting anomalies. All decisions about anomalies and risk acceptability in the Shuttle Program are a collective effort. In the pre-launch phase, each component part of the Shuttle is the responsibility of a Project work group, to which NASA managers and engineers are permanently assigned. The experimental technology requires specialized skills, so engineers on the same work group have a division of labor dependent upon whether they are thermal engineers, structural engineers, or some other specialty. However, the Project members, in work groups assigned to technical components of the shuttle, collectively identify anomalies, assess

risk, and determine whether or not their component part is ready to fly. Post-flight analysis - assessing the performance of a few similar objects/vehicles after each flight - is an important part of the task for shuttle engineers that has no equivalent in air traffic control.

In air traffic control, division of labor has individual controllers making risk assessments incrementally as flight progresses. The work is divided sequentially: the controller responsible for a plane at a given moment is the decision maker assessing risk. Responsibility is passed from controller to controller as the aircraft is handed off to the next controller as the plane taxis along the ground and through sectors of airspace. Once an airplane starts moving, a controller's opportunity to identify anomalies, assess risk, and take corrective action can be a split second affair. Each controller may have control of a plane for from five to twenty minutes, depending upon the geographic size of the controller's air space, the flight plan, the speed and destination of the plane, the amount of traffic, and the weather. Although in some circumstances that controller may have input from a supervisor or other controllers nearby, the typical case is one controller, aided by whatever technology is available in his or her facility, assessing risk, making decisions, and communicating those decisions to pilots.

In both types of space work, more than one set of human eyes is monitoring a flight, but the number of human eyes monitoring a given flight, either directly or indirectly via some technology, is huge. For NASA, a shuttle launch is a relatively unique event: only one shuttle is in orbit at a time; in a "good" year, eight shuttles may be launched. NASA personnel are monitoring progress and performance not only at the launch and landing sites, but at numerous NASA facilities around the globe. Each vehicle is monitored by thousands of eyes. In contrast, launching airplanes is a routine event. At the busiest times of each day, major terminals are launching 2 planes a minute; typically, over 5,000 planes are in the air over the United States. Also in contrast with NASA, every controller has many more than one airplane to control at a time. The path of each plane is watched by a small number of people. Controllers "own" sectors of airspace; a plane in a sector is controlled by that controller, but others in the room (at a Center [high altitude], the radar controllers' radar assistant watches all activity on the same scope; at a tower or a TRACON (intermediate altitude) airspace and architectural space differences make it possible for other controllers and supervisors in the room to see and hear what's happening

for every controller, so provide extra sets of eyes.

For NASA and NATS, the variation in certainty of their respective vehicles affects the time and division of labor for encounters with risk. These differences, in turn, lead to distinctive technologies of control.

#### THE TRAJECTORY OF ANOMALIES:

##### TECHNOLOGIES OF CONTROL AND THE DETERMINATION OF RISK

Technologies of control can be thought of as manufactured sets of "extra eyes": Rules and procedures are designed to point human eyes and behaviors in designated directions; surveillance technologies are designed to supplement the human eye and cognitive capacity by revealing in detail what the human eye cannot see. Both have a significant impact on work practices, the third technology of control. I turn now to examine technologies of control as they affect the identification and trajectory of anomalies in these agencies.

#### Space Transportation System

At the outset of the Space Shuttle Program, NASA acknowledged the uncertainty of shuttle technology. Its experimental character was reflected in a formal system of organizational and institutional rules and procedures that left wide discretion to engineers in work groups in all their attempts to identify, measure, and assess risk. The formal rules set forth decision making procedures that frame these expert technical decisions, setting standards for what happens to the information that the work groups produce after work groups produce it. The assumption from the start of the program was that this was a risky technical system because it was without precedent. Because the forces of nature experienced in flight could not be replicated in the laboratory and in field tests, the expectation was that always a shuttle would return from a mission showing physical damage and deviations from predicted performance. Part of the organizational context of risk assessment at NASA was the institutionalized belief that a) a vehicle would never be risk-free and b) anomalies were expected throughout the Space Shuttle vehicle after every mission. These cultural understandings were encoded in a written protocol governing all engineering decisions. Prior to the first shuttle launch in 1981, NASA created an overarching universalistic guideline for all risk assessments titled "The Acceptable Risk Process" (Hammack and Raines 1981). The Acceptable Risk Process was the basis for all technical decision

making at NASA, from daily decision making to the final formal decision process immediately prior to a launch known as Flight Readiness Review.

The document stated that hazards that could not be eliminated or controlled would be subject to risk assessment determining the "credibility and probability of the hazardous risk occurring" and required an engineering rationale for retaining the existing design or procedure. It explained that a hazard would be classified as an Acceptable Risk only after all risk avoidance measures were implemented and documented and after upper management, upon consideration, decided to accept risk on the basis of the documented engineering rationale. It concluded by acknowledging that after all feasible corrective actions have been taken, some residual risks would remain. Consequently, all risks needed to be monitored "to insure that the aggregate risk (of the shuttle) remains acceptable (Hammack and Raines, 1981: 10)." NASA had no coordinated or intensive in-house job training in risk assessment for engineers assigned to a project. Instead, the agency trusted in professional training and previous experience. Engineers are trained to work for corporations and government agencies that will assign them to specialized tasks, which they will learn on the job. In every position, they carried with them occupational rules: universalistic engineering principles and skills learned in schools and other jobs. At NASA, the identification of any anomaly was to be followed by an engineering investigation into its seriousness and risk acceptability that was to include the measurement of risk: how likely is it to recur and how risky is it?

Identification and Measurement. The experimental character of the space shuttle made the entire process of identifying and assessing the seriousness of anomalies murky, but action had to be taken. Any problems would need to be addressed and corrected prior to the next launch. Because of the unprecedented character of the shuttle technology, engineers assigned to a Project had no particularistic rules to guide their technical decisions. They had to create their own criteria for their component part's performance. These standards were developed inductively from experience, first determined by predictions from engineering principles, then continuously adjusted in response to experience in bench tests, lab tests, and finally, shuttle missions. This situation was not unique to NASA, but typical of engineering as an occupation. Wynne (1988)

writes about the ad hoc character of engineering rules, noting the absence of appropriate rules to guide fundamental engineering decisions. They tend to develop informally out of evolving engineering experience with a technology. Formal rules established prior to experience with a complex technical system are designed to apply to diverse circumstances, but are not sufficiently specific to guide decisions in specific situations.

They were, in short, learning incrementally, by mistake. Their long time line between launches allowed them in many cases to repair the technical cause of the problem, and in others, to find out if it was tolerable in the shuttle: measure the amount of deviation from expected performance (in engineering-speak, calculate the margin of safety) to ascertain if the anomaly was an Acceptable Risk. The latter determination was a matter of producing evidence that the damage might continue but would be contained, so that it did not threaten the safety of a mission. Consequently, the trajectory that Star and Gerson identified in scientific work was replicated in the technical work of NASA work groups: Incidents that engineers first identified as anomalies that were potential signal of danger often, on the basis of post-flight investigations, were redefined as normal and acceptable in subsequent flights.

Wynne (1988) observed further that as engineers begin to develop experiential knowledge, much of what they understand is intuitive, based on tacit knowledge that is difficult to express. Yet, especially at NASA, I found, engineers are challenged to convert their tacit knowledge into visible, discrete, units that make communication and decision making possible. To do so, they have to rely on the methods of scientific positivism that they learned in engineering school in order to quantify, rank, classify, standardize performance and assess risk. For NASA engineers, a compliance strategy of control was one that had them identifying anomalies that had to be brought into compliance with continuously shifting rules that were created ad hoc, rather than rules that provided standardized measures that were constant and consistent. But guidelines for measurement were also murky. Because of the unprecedented technology, engineers had to create their own measurement devices, in the form of models and field or bench tests. Not only did they invent the rules as they went along, but they also had to invent the technologies of control they used for surveillance of anomalies.

Research in science and technology studies shows that for complex technologies, experiments, tests,

and other attempts to measure outcomes are subject to "interpretive flexibility" (Pinch and Bijker 1984). Engineers in the Solid Rocket Booster Project found that various tests of the same phenomenon produced differing and sometime contradictory results. Even the results of a single test were open to more than one interpretation. Attempts to settle controversies by having additional tests run by outside research institutes or contractors not involved in the shuttle project, often led to yet another set of findings, which Collins (1985) research found to be common in attempts at replication of experiments. Consequently, controversies were part of the daily routine. Disagreements about what was risky and how risky it was were ongoing and continually negotiated in conversations and formal work group meetings. Further, as they learned more about both the performance of the boosters and the measurement devices, they developed increasing sophisticated tests and models, so data comparison at, say, time 1 and time 6 was impossible. For all these reasons, judgments were always made under conditions of imperfect knowledge (Marcus 1988).

As shuttle engineers performed these acts of identifying, measuring, and determining risk, they were engaged in work that was more like a craft than a science. They were experimenting tinkering, adjusting, inductively recreating the shuttle as they went along. They eliminated the anomalies that they could with targeted adjustments that fixed the problem. For those that remained, criteria for how risky an anomaly was and how much deviation could be safely tolerated were adjusted based on post-flight analysis after each shuttle mission. The retrospective nature of their interpretive work and standard setting was more a project in the invention of accuracy than developing standardized technical guidelines for assessing risk. For those doing the assessments, however, their shifting standards were not questionable actions. Indeed, within professional engineering culture, they were practicing "normal science" (Kuhn 1962).

The end point for the trajectory of an anomaly was a formal status of classification in which the nebulous became an organizational fact. Measurement of the amount of risk and probability of failure, however nebulous and controversial, had to be converted into information that was conveyed to the rest of the system. Two NASA requirements reinforced the trajectory from anomaly to Acceptable Risk. Their unintended consequence was to aid in the conversion of uncertainty into certainty: the ongoing engineering controversies,

negotiations, changing risk estimates and toleration levels for deviations were transformed into stable, formal, institutional facts that were binding. One procedure was a classification system called a Critical Items List that attempted to rank the relative risk of each shuttle component. The Criticality categories were formal labels assigned to each part, identifying the "failure consequences" should the component fail. These designations (Criticality 1, Criticality 2, Criticality 3) were conservative because their assignment was based on failure analysis under worst case conditions.

A work group composed the entry for its component part of the shuttle, giving the engineering rationale explaining their decision for Acceptable Risk. Since all items were risky, the entries described the data and actions taken that the engineers believed precluded catastrophic failure and stated why, in spite of the risk, the design should be retained. As the document represented an official work group position, engineering differences of opinion had to be negotiated to get the document written and approved. Not everyone agreed about the level of risk or the engineering rationale and evidence included. But the document required a consensus, so one was always produced; decisions had to be made about which evidence or test had more weight. The documentary representation of reality was one of stable knowledge confirmed by scientific evidence. Producing these documents classified Acceptable Risks by sorting them into Criticality Categories that affected the definitions held by management and others in the Space Shuttle Program. Ambiguity was erased in the process

The second NASA procedure that forced fact making in the process of defining risk was Flight Readiness Review (FRR). FRR is NASA's pre-launch decision making structure that is the final, formal review in an intricate process of launch preparation and decision making that involves thousands of people and engineering hours and takes 18 to 24 months. The goal is to collectively determine that the shuttle is ready to fly, and fly safely. FRR brings all parts of NASA and relevant contractor organizations together to assess engineering determinations of Acceptable Risk for each part and thus the aggregate risk of the shuttle. The review is a tiered process with four levels. To proceed with a launch, the decision to go must be based on consensus of all parties. It begins about two weeks before a launch at Level IV, with Project work groups and their contractor

counterparts bringing forward risk assessments and documented engineering evidence, based on the most recent data from tests and post-flight analysis, that their component is safe to fly. To prepare for this stage, the people assigned to the Project must negotiate their differences and agree on the meaning of tests and other evidence about performance.

The entire FRR process is rigorously adversarial. As one engineer said, "I've seen grown men cry." The initial documentation and statement of risk acceptability are presented before knowledgeable others, who will challenge the work group's analysis. Hiding evidence, sweeping problems under the carpet is not possible, because every stage has a critical informed audience, some of whom have a vested interest in making a rival Project look bad by finding problems that will hold up the process, giving their own Project more time to work out their analysis. If the scrutiny produces unanswered questions, inconsistencies, or inadequate evidence, the work group will be sent back to the drawing board to do more tests, analysis, or whatever it took to produce convincing evidence that the component was an Acceptable Risk. When the evidence withstands scrutiny and all challenges are met, Project Managers sign a document indicating the groups' decision that the component is safe to fly. As the engineering analysis for each component works its way through the 4 levels of the FRR hierarchy, the adversarial audience gets larger, and the analysis gets tighter and tighter, as more and more data are brought to bear affirming the original position, stabilizing and formalizing the definition of risk. At each Level, senior management participants are required to sign a document stating their approval that the design is an Acceptable Risk and therefore safe to fly. These signed documents not only indicate consensus, but also affirm management's oversight responsibility in launch decisions. Officially and formally, responsibility for error, mistake, and accident belongs to all participating parties.

After FRR, the final phase of the pre-launch preparation switched the responsibility for risk and error to a Mission Management Team that supervised a standardized system in which computers and other hardware surveillance technologies dominated. Governed by two volumes of items that had to be checked one-at-a-time, most decision criteria were quantified and precise, with predetermined criteria for stop or go (unless an unprecedented situation arose, as it did on the eve of the Challenger launch, when wind and cold produced 18"

icles and no one knew if the hurtling icicles would damage the vehicle at lift off.) The final countdown was computer surveillance of each operating component that automatically stops the countdown by shutting down the system when an anomaly is detected. NASA engineers don't expect major glitches in this phase of the process. The purpose of FRR *is* flight readiness. As one engineer said, "By the time it goes to the Mission Management Team, the hardware is all cleaned and ready to go."

To sum: I have singled out for attention several technologies of control for the Space Shuttle Program that grew out of encounters with risk on an uncertain technology. Notably absent was any in-house training as a technology of control; in response to uncertain technology, the agency placed trust in professional engineering training and experience, allocating deference to professional expertise. Technologies of control in the form of rules and procedures were of three types: an overarching guideline, "The Acceptable Risk Process", that created and reinforced an institutionalized belief system about risk that governed all NASA decisions about flight safety; a flexible schema of engineering rules, created bottom-up by engineers, that were the basis of specific engineering decisions evaluating the performance of the shuttle and resulting risk assessments; and top-down imposed, standardized rules and procedures, institutionalized across the NASA organization designed to coordinate decision making across the system. The work practices that were employed as a result had the technical experts doing the hands-on work and encountering risk on a daily basis creating rules from experience with the technology in an ad hoc fashion. These functioned as standards that were used to identify anomalies and convert anomalies to Acceptable Risks. These rules were consistent with the interpretation of evidence and determination of what was or was not a fact at the moment of their creation, but as new learning occurred with the shuttle technology, the facts changed. The rules were constantly adjusted.

Engineers also crafted their own hardware surveillance technologies: measuring devices, such as instruments, bench and lab tests. The results of measurement had interpretive flexibility that made controversy part of the daily routine. The NASA organization required consensus, however. Engineers' decision making was framed and shaped at crucial points by institutionalized beliefs and organization rules for decision making and documentation of risk that were bureaucratic. These formal rules coordinated and regulated encounters with

risk between shuttle project work groups and the rest of the organization. Because of the complexity of the shuttle as a technical system and the complexity of the NASA/contractor system, there were myriad daily documentation and coordination requirements. I focused on two because they play significant roles in the trajectory of converting anomalies to Acceptable Risks: the Critical Items List and Flight Readiness Review. These standardized rules and procedures are in contrast to the experiential and flexible rules of engineers. Although the strategy of having rules and procedures frame engineer decisions by defining what should happen to data and risk assessments after engineers have produced them would appear to have no effect on engineering data and determinations of risk, at NASA the formal rules and procedures have a significant effect on the technical decisions. They converted uncertainty to certainty, forcing contradictory evidence, ambiguous results, and controversies into consensus and organizational facts. These processes were reinforced by structural conditions: structural secrecy within the organization, and the culture of production, which developed from political and economic relations in the environment. Neglected in this telling, they were equally influential in the normalization of deviance at NASA (see Vaughan, 1996, Chapters 6-7).

### Air Traffic Control

The certainty of airplane technology and the standardization of technologies of control in the air traffic system invest the tasks of air traffic controllers with a clarity unknown by the technical experts in shuttle work groups. The National Air Transportation System is part of a global system of regulation with both universalistic and particularistic rules that apply across national and cultural boundaries. However, particularistic rules dominate the system. English is the international language; the sky is divided into regions of air space that are "owned" by specific air traffic facilities. Each defined air space has identifiable boundaries, and each is crossed with named highways and intersections (e.g., Weir, Bronc) that indicate the routes of airplanes, like a city street map guides the path of automobiles. Commercial flights fly these standard routes at scheduled intervals that are known and predictable. Rules and procedures are designed both to control airplanes and coordinate the activities of controllers. How traffic is handled by controllers is set forth in the "bible" of air traffic controllers, known as 7110.65. It is two inches thick and codifies all rules and procedures for communicating with pilots

and controlling aircraft. Also, a published volume titled "Letters of Agreement" is specific to the needs of each facility. The Letters of Agreement contains all the procedures for coordinating activity with other controllers in other air traffic facilities within and adjacent to its airspace.

In contrast to the shuttle program, in which engineers educate themselves in shuttle technology and invent many of their own rules, for ATCs these technologies of control all are imposed top-down. The goal is to produce an air traffic controller whose behavior is as standardized and predictable as the system in which they are to work. Thus, education is a significant technology of control. The first line of defense against anomalies is controlling the decisions of the individual controller. Education aims directly at the decision making process itself. The goal is to have controllers master the rules so well that they can make decisions without thinking (Vaughan 2002). An extremely important part of their job is the quick identification and correction of errors and mistakes in the movement, direction, speed, altitude or airplanes in motion, carrying people and cargo. The timing of decisions is so rapid that having to think or calculate everything is a threat to safety - a process that stands in contrast to the prolonged group processing of information and debate style decision making in the STS program. Therefore, the making of a controller is an intensive, arduous process. The training of controllers is in stages, its duration varying depending upon the type of facility to which each is assigned. The average training time across types of facilities is 3 years. Moreover, even after formal training is successfully completed and capped by the formal certification of a controller, retraining goes on continually, integrated into the work routine in all facilities, because the existing system keeps being refined to meet changing traffic demands and new types of aircraft.

Controllers are trained to achieve the FAA goal of "the safe, orderly, and efficient coordination of air traffic." Unlike engineers, who come to the workplace already trained in engineering principles, ATCs bring no universalistic principles from professional training to the workplace. Their background work experiences and training are not in air traffic control. There are some exceptions: some were pilots or military air traffic controllers, but they still need to learn and be able to apply the technologies of control that are peculiar to the FAA system. Most controllers are either recruited off the street or from technical schools that specialize in

aeronautics and give degrees for air traffic control specialist. Controller training requires mastery and integration of three layers of knowledge: institutional, organizational, and tacit knowledge. The institutionalized rules of the national and international system are taught at the FAA training facility in Oklahoma City, where all candidates must spend time training. First they must pass a "screen": tests on a computer-based instruction system that simulate skills believed important for air traffic controllers. Following success, candidates take classes in which they learn air space maps with thousands of bits of information, phraseology, technique, how to use the technology. They memorize and are tested daily. On simulators, they work traffic situations so realistic that the adrenaline is pumping. It is survival of the fittest, and many controllers told me it was the most stressful experience of their lives.

Those who survive are assigned to a facility, where they acquire local knowledge: the facility's air space, traffic patterns, types of aircraft, hardware technologies, and Letters of Agreement. New controllers begin by memorizing maps and the rules for their airspace. Prior to working "live" airplanes, they spend time on traffic problems on simulators, building up skill at handling complex situations and more and more traffic. The next phase is working traffic in an apprentice system. "Developmentals" are assigned a primary and secondary trainer. After lengthy work assignments that allow them to observe experienced controllers at work, developmentals begin working traffic with a trainer who corrects them on the spot and takes over when necessary. This aspect of training is stressful for both the developmental and the trainer. Like NASA engineers, they learn by mistake, but working live traffic is a qualitatively different experience than the pre-launch or post-launch weighing of evidence and learning by mistake by lab tests and shuttle parts on the ground at NASA. What controllers also begin to acquire during facility training is tacit knowledge: the essential subtleties of air traffic control - which types of aircraft can do what, or predicting competency of a pilot by the tone of voice, or sensing when the wind is changing, or how to coordinate with the varying skills and temperaments of colleagues. An extremely important skill they learn in training is "room awareness": controllers are taught to pay attention not only to their own traffic but to what everybody else in the room is doing. At the same time they are talking to pilots, watching the scope or looking out the window, writing on Flight Progress Strips, and coordinating activities

with other controllers through the computer or ground phone lines, they are acutely aware of what everyone else is saying, even to the private whispered joke of a coworker on the other side of a room. No privacy is possible. All decisions immediately are public, and so are mistakes.

The outcome of the top down imposition of technologies of control on controllers is that institutionalized rules, local knowledge from the facility, and tacit knowledge become embodied: controllers can control traffic without thinking. Some bottom up invention is possible in air traffic control: controllers identify air space problems and help design realignments of the boundaries in the sky in order to improve the traffic flow; they also participate in the development of new technologies. However, particularistic rules and standardization govern every aspect of decision making by controllers - what they do, what they say, and how they say it. They do not have universalistic guidelines, like the Acceptable Risk Process and Flight Readiness Review as frames for wide discretion and bottom-up invention of technical standards for what is or is not an anomaly. Controllers' search for anomalies - the identification, measurement, risk assessment, and correction of any deviations they spot - is done within detailed and narrow constraints. This layered system of institutional and organization rules and procedures leave them a limited range of discretion. Although limited, their discretion is as critically important to their work as it is at NASA, and so is the intuition that comes from experience and tacit knowledge.

Identification and Measurement. Both shuttle project work groups and air traffic controllers are engaged in prediction and the search for deviation from expectations. However, identification of anomalies for an air traffic controller is far from the murky process of the Space Shuttle Program work groups. In a standardized system that is ingrained in controller cognition through extensive, intensive training, any deviations from expected behavior instantly stand out. In air traffic control, aircraft behavior either fits the standard pattern or it does not. Any deviation attracts the attention of controllers, who immediately act to bring the situation back into compliance with standards. They are trained to automatically spot and correct, for example, a pilot's deviation from assigned route; a routing error printed on the computer-generated Flight Progress Strip; a pilot error in repeating a controller instruction ("climb and maintain heading one-two-zero" instead of

"one-three-zero"). These sensitivities are part of tacit knowledge, a result of formal training and experience, learning by mistake being a part of both. Controllers daily encounters with risk are equally reliant on hardware surveillance technology that aid in monitoring and identifying deviations - computer systems, radar, navigational devices, radios, and radio frequencies. All are important aids to memory as well as tracking, communication, and identification of anomalies. In contrast to shuttle work groups, whose hardware technologies for anomaly surveillance and assessment of risk are invented and changed as knowledge and shuttle technology changes, in air traffic control hardware surveillance technologies are standardized and slow to change because of cost of replacing them throughout the system.<sup>ii</sup>

Remarkably, within the rigid standards of this system and the plethora of hardware technologies at their disposal, much of controllers' ability to identify and correct anomalies nonetheless rests upon experientially-based human assessment and intuition. Controllers even are able to be prescient, taking preventive action before anomalies occur. Their tacit knowledge gives them almost a sixth-sense about pilots, allowing them to predict and prevent a mistake waiting to happen. Their attention is attracted to differences between pilots' voices that are indicators of a need for what controllers call "special handling." Foreign pilots, for example. Many pilots with foreign accents, controllers know from experience, may have mastered English well enough to be a qualified pilot, but may misunderstand a directive or not have sufficient language to comprehend anything beyond the most routine instruction, so cannot handle an emergency in English. For foreign pilots, controllers speak more slowly, more clearly, and break down instructions into separate parts. Controllers recognize the difference between the experienced pilot's command of phraseology and that of the infrequent flier or pilot in training, and pay extra attention accordingly. Controllers know the performance capability of the types of aircraft common in their airspace and recognize when one is behaving inappropriately for that type. Any signals of potential danger get passed on to the next controller ("Watch this guy, he's...".)

In contrast to NASA. in air traffic control anomalies do not wait to be leisurely considered. The quick fix ability of air traffic controllers renders most anomalies momentary, so there is no protracted organizational trajectory in which every anomaly is routinely and systematically considered by a team of people, then formally

and officially converted to some equivalent of the shuttle concept of "Acceptable Risk." However, there is a trajectory of conversion from anomaly to error that has a formal classification as end point. Measurement with hardware surveillance technologies is central to the distinction. As in the shuttle program, scientific positivism and the assumption that science produces accuracy prevail. Controllers' primary responsibility is to avoid collision above all else. To that end, they must abide by the "Rules of Separation." These rules define the amount of space that must exist between aircraft in the air and on the runway. The spacing requirements vary by type of facility, its assigned airspace, and type of aircraft, and are specified in 7110.65. For example, for high altitude centers the rules of separation require spacing of 1,000 feet above or below an airplane and 5 miles between aircraft. If the rules of separation are violated, the violation may be charged to the controller, so the anomaly is officially defined as an "Operational Error." The anomaly is transformed into an Operational Error (an OE) when an investigation shows that the controller failed to act to avert the incident or acted erroneously or too late.

But there are other possible outcomes. The example of the high altitude center, which uses radar to follow the route of aircraft flying at altitudes of 14,000 ft. and above, shows the role hardware surveillance technology and interpretive flexibility play in the trajectory of an anomaly. In all facilities, radar is present, but it is exclusively relied on for controlling traffic at 3,500 feet and above. A radar scope contains not airplanes but representations of airplanes in the form of data blocks giving call signs, destination, type of equipment, speed and altitude. These data blocks are shown on controllers' radar scopes attached to targets that indicate the location of an aircraft in the sky. The blocks move as the aircraft move, allowing controllers to track the paths of the aircraft. Controllers can predict the future location and potential points of conflict of planes in their airspace by using the information in the data block, so are able to change headings or altitudes to avoid possible conflicts. Also, the computer can calculate the distance between two planes for them, based on the radar tracking system: figures appear printed on the scope.

Finally, the computer keypad offers a function key that allows them to put a 6-mile diameter circle (formally known as a J-ring, or informally to controllers, a "sissy ring") around an airplane in a position of potential spacing conflict with other aircraft in a congested traffic situation. Two consequences follow, one

human and one technological: First, the 6-mile diameter J-ring gives a visual boundary that increases the margin of safety one mile beyond that dictated in 7110.65 so controllers have ample time to give corrective headings; second, the computer always is measuring the distance and altitude spacing between aircraft, based on the radar indication of a plane's location, whether the controller asks for it to be displayed or not. Consequently, when the rules of separation are violated (i.e., the separation between planes is 4.9 miles instead of 5 miles), a J-ring automatically appears around the two airplanes in conflict. The ring flashes off and on, also flashing the letters CA-CA, which indicates a Conflict Alert. This technological indicator of a Conflict Alert is a signal of potential danger that attracts controllers' attention so they can - and do - give pilots instructions that provide the necessary separation.

It is measurement that sorts these otherwise momentary anomalies into a formal organizational classification system and converts them into organizational fact. An official investigation follows every computer Conflict Alert. The flashing computerized warning on the scope of a Center controller is matched by an alarm that is triggered on a computer (known to controllers as "the snitch machine") located at the Watch Desk of the Operations Manager. The Ops. Manager then calls the supervisor in charge to ask about the incident. The supervisor removes and replaces the controller (who has already fixed the conflict), putting in a new controller to work traffic while the Conflict Alert is investigated. Having an Operational Error is an emotional experience, from the moment the of Conflict Alert through the investigation, the controller's decertification, and reinstatement. For investigation, the controller writes a description of the incident, the supervisor retrieves a print-out of the two aircrafts' flight paths from the computer, and also makes a duplicate tape of the recorded conversation between the controller and the pilots in his or her airspace at the time. The incident is investigated by a Quality Assurance officer and the Operations Manager; participants include the controller, a union representative to assure the process is fair for the controller, the supervisor, and the controller's assistant if one was present because in that situation the responsibility is shared between the two controllers.

The radar track printout first is examined to decide whether the 5 miles, 1,000 feet rules of separation

were violated. There are several possible outcomes, and here we see where measurement of anomalies is typified by the same interpretive flexibility and the invention of accuracy as in the Space Shuttle Program. The computer may have made an error. By the time the computer sets off the Conflict Alert, one of the airplanes may have turned to a different heading that averts a conflict. Representations stand in relation to but are not identical to the object that they represent (Latour 1987). All radar representations are inaccurate because the radar is not in constant surveillance of all airspace. but sweeps it at intervals. At the high altitude centers, the sweep is in intervals of 12 seconds. The controller sees where the plane was 12 seconds ago, taking that into account when calculating the location of an airplane. Pilot responses to controller instructions are already underway before a controller sees them on the scope. The 12 second interval has the deleterious effect of creating distrust among controllers for the Conflict Alert as a signal of potential danger. They report that often it seems like "the boy who cried wolf." If, on the other hand, the investigating team finds the rules of separation were violated, the next step is to find out what happened, identify the cause, and fix the problem so that it won't happen again.

However, interpretive flexibility of the technology for measurement again confounds the outcome. Interpretive flexibility increases at Towers and TRACONS. First, a computer Conflict Alert is not an OE when the controller gave the planes' pilots visual control of their descent or ascent and advised to watch for possible conflicts, a condition in which responsibility for separation shifts from the controller to the pilot. Second, it is not an OE when the rules of separation are violated, but the two aircraft were on diverging headings, so collision was not possible. Finally, at all facilities, the investigators may attribute the error to the pilot, thus the incident is classified as a "Pilot Deviation." An example would be when the controller gives the pilot a heading or altitude change and the pilot verbally confirms the instruction but fails to make it or does not make it quickly enough, thus violating the rules of separation.

If the controller is charged with an OE, what has so far remained an informal conversion of anomaly to Operational Error now becomes a formal, public, organizational fact. The controller, who was pulled off of his or her position at the time of the computerized Conflict Alert, is not allowed to return to working air

traffic until he or she has completed some retraining, usually consisting of computer-based instruction designed to remedy the problem. Depending on the type of error, the controller may be retraining for a number of days. The outcome of the investigation becomes known through gossip channels. Word of the decision, who is to blame, and the traffic situation that led to the Operational Error spread informally and are discussed among other controllers in that facility, who are uniformly and obsessively interested in the details of OEs and accidents, trying to find out what went wrong and how it might have been avoided. The incident is formalized, written up and distributed to supervisors, who place it in the Read-and-Initial files at his or her desk. More formal attempts to learn from mistake follow at high altitude centers. Annually, all OEs for the year are reviewed in the facility by the Operational Error Review Board, comprised of union members and quality control representatives, who examine what happened in each case and try to establish patterns. However, the degree to which it is public goes beyond the local. Every facility records its Operational Errors, this count tallied daily, monthly, quarterly, and annually and conveyed to regional and national headquarters. Part of the national data system of the FAA, these totals are distributed to all facilities. Facilities are ranked by the number of OEs they have in a period.

Although the official response to OEs is designed primarily to learn from mistake and correct the problem so that the same type of mistake is not repeated, it is a deterrent strategy of control that has no parallel in the shuttle system. Having Operational Errors affects not only the status of the facility, but also the status of the controller who made it. Controllers refer to having an Operational Error as "having a deal" because it is a Big Deal, organizationally and personally: three deals in a year, you're out of a job. Moreover, the controllers' crew also experiences it as punishment to the group because OEs are attributed to the crew as well as the controller. In some facilities, rewards, in the form of time off, are given to crews that have long records with no OEs; when one happens, they lose this "good time." Although almost all controllers have deals - they are impossible to avoid in some air spaces and some traffic situations - controllers experience it as punishment. The moment of the deal is emotionally traumatizing in its own right; being taken off of position, being subject of an official inquiry, and then to remedial education is stigmatizing.

To sum: Encounters with risk in air traffic control leave little to guess-work. As one controller said, "Structure and routine, structure and routine. If we have to improvise, we improvise from this base." The division of labor in the system gives the responsibility for risk assessment to individual controllers with individually owned airspace, not a team. The certainty of the airplane technology controllers regulate has resulted in a high degree of standardization in the technologies of control employed in the air traffic control system. The primary focus is controlling the controller, which makes education and training significant as a technology of control. Institutional and organizational rules and procedures do not stand apart from decision making by technical experts, framing them, as in the shuttle program, but are the very essence of decision making, becoming cognitively embedded, as controllers integrate institutional, organizational, and tacit knowledge (Vaughan 2002). Controllers quickly identify anomalies, which stand out from standardized, predictable patterns. Hardware surveillance technologies of control, also standardized across the system, are not only tools for communication and coordination, but also aids in the identification and correction of anomalies.

The trajectory from anomaly to categorization and classification that Star and Gerson describe takes a different path in air traffic control than the shuttle program. Anomalies are converted by hardware technologies of control into Pilot Deviations or Operational Errors, the latter adding a deterrent strategy of control to air traffic's primarily compliance system. Significantly for our consideration of risk and error is that controllers *are not trained to avoid collision per se; they are trained to avoid a violation of the rules of separation*, a practice that builds in safety by timing their corrective practices sufficiently prior to a possible collision that accidents are avoided. Despite the standardization of the system and tight constraints on controller decision making, preventing errors, and identification of anomalies by controllers relies to a great extent on individual experience and intuition. In a system where more and more sophisticated hardware technology is being introduced to relieve congested traffic and reduce the reliance on humans, the experience and sensitivity of human interpretation and cognition is still essential. It is critical in air traffic control due to the sometime inadequacies of representation and measurement by hardware surveillance technologies. Ironically, however, exactitude and scientific principles govern the system, emphasizing standardization, technology, and scientific measurement,

while underemphasizing the contribution of individual intuition and thought.

#### ORGANIZATIONAL RITUALS OF RISK AND ERROR

This comparison joins other recent work that is designed to bridge the so-called Great Divide: recently, scholars studying the risk of complex technologies have been categorized as either followers of High Reliability Theory or Normal Accident Theory (see, e.g., Sagan, 1993; La Porte, 1994). Following Perrow (1984), Normal Accident Theorists study failures, emphasize structures, and argue that complex systems will inevitably fail. Normal Accident theorists tend to study harmful incidents after the fact, concentrating on public failures with high costs (for an exception, see Sagan (1993) on near-misses). High Reliability Theorists, on the other hand, study safe systems, emphasize processes, and argue for effective prevention. They study processes by linking them to successes. Rather than after the fact analysis, their research is done by locating themselves in an organization to interview and watch work practices and how risk is managed. The distinctions made between the two approaches have been very useful in pointing out the advantages and disadvantages of each. However, more recent work blurs the genres created by the dichotomy: for example, research on a risky industry (Carroll and Perin 1995; Bourrier 1999), research comparing organizations with failures to those without (Weick 1990, 1993, Roberts and Libuser 1993, LaPorte 1994), studying the connection between structure, processes, and cognition (Clarke 1992, 1993, Vaughan 1996, 1999, 2002), and variations in structure, process, performance, and accidents across organizational settings (Schulman 1993, Marcus 1995).

In this paper, I bridge the so-called Great Divide by arguing first that all formal and complex organizations are subject to routine nonconformity: all are systems that regularly produce incidents deviating from expectations (Vaughan 1999). Then, I examined two organizations, both of which have been defined as High Reliability Organizations. Both NASA and the FAA's air traffic control system are organizations that routinely produce mistakes/deviations/anomalies that are not visible to the public (for NASA, failures on the rocket test range at a rate of 1 out of 25; Operational Errors in air traffic control). I have treated anomalies as signals of potential danger, examining the organizational response. Defining technologies of control as 1) rules and procedures that guide decisions about risk, 2) work practices, and 3) surveillance technologies consisting

of both procedures and hardware designed to identify anomalies that are potential signals of danger, I examined how anomalies are analyzed and transformed into formal organizational categories. In the comparison, I focused on the trajectory of anomalies in both agencies: how anomalies get identified, measured, and classified as risks and errors at each workplace.

I found that both agencies are continuously engaged in "clean-up work" (Vaughan 1999), applying technologies of control so that many deviations get corrected early, preventing small mistakes from turning into costly, harmful public failures. In response to differences in the uncertainty of airplane and shuttle technologies, the technologies of control for each agency differ because the tasks differ. So, for example, I found the agencies have very different kinds of rules and procedures, NASA allowing for more discretion in many circumstances and NATS much less. Variation in technological uncertainty also shaped fundamental differences in the time and division of labor for risk assessment and clean-up work. What are we to conclude from this comparison? In spite of the different technologies of control in use, they had common consequences. Because they were repeated, regularly and routinely integrated into the daily work experience, the technologies of control developed ritualistic properties with symbolic meanings at both agencies. They created taken-for-granted cultural understandings that had an effect at the social psychological level, affecting the interpretation of signals of potential danger. As the technologies of control varied, so did their social psychological impact. The result was significant differences in technical experts' sensibilities about anomalies and risk at the two agencies.

At NASA, the cultural understanding that encompassed all risk assessment was that problems and anomalies were normal and tolerated (Vaughan 1996: Chapters 3-5). Due to the uncertain, experimental nature of the shuttle, technical problems were expected to occur on every mission. Ironically, procedures, documentary requirements, and rules designed to separate unacceptable risks from acceptable risks also had the consequence of declaring many anomalies "acceptable", and therefore tolerable. The formal "Acceptable Risk Process", the official category of "Acceptable Risk", and bureaucratic language formalizing determinations of Acceptable Risks permeated every day discussions of technical anomalies. Words like "mistake" and "error" and "accident" were not part of this language. "Failure" and "catastrophe" were, however; in fact, they were so

much a part of standard engineering language, that these words lost their attention-getting quality. Engineers routinely did analysis of the effects of possible failures, but the ability of these words themselves to serve as warning signs ("We're going to have a catastrophe if we do X") was mitigated by routine use on bureaucratic forms calling for engineers to describe the consequences of failure of every shuttle item (e.g., if a wing fell off, because the shuttle had no back up for its wings, the failure outcome was "catastrophe.") Further, these Acceptable Risk procedures helped to create a culture where there was always a second chance. Thus, things going wrong was not a signal of potential danger but a routine and taken-for-granted aspect of the daily routine.

The Space Shuttle Program's technologies of control desensitized technical experts to the possibility of error, mistake, and failure. My research on the Challenger accident showed how these formal rules, categories and procedures, like Flight Readiness Review and the Critical Items List, had the effect of nullifying the sense of risk at the social psychological level, contributing to "the normalization of deviance." On the Solid Rocket Booster, once the technical deviation was identified and labeled an Acceptable Risk (due to some engineering fix, or engineering calculations predicting the damage incurred would be within worst case limits), NASA engineers and managers continued to build on this original classification, so that the increased frequency and seriousness of the same anomaly also was acceptable. Only in retrospect, after the Challenger's tragic demise, did participants who made the risk assessments for the Solid Rocket Boosters redefine their actions as mistakes in judgment. At the time the decisions were made, however, accepting anomaly after anomaly in booster performance seemed normal and acceptable within the organizational and institutional context of work that existed at the time.

In air traffic control, in contrast, the technology of airplanes is standardized and certain, the mechanical problem being the exception, not the rule. Technologies of control are similarly standardized. However, the cultural understanding is that error and mistake are not acceptable and will not be tolerated. The language of error and mistake permeates both formal categories and informal conversations: hear-back read-back error, violation of air space, pilot deviation, violation of rules of separation, accident. Anomalies that cannot be fixed result in a violation of the rules of separation and are classified as Operational Errors. This formal designation

is real in its effects, because the procedures attached to a violation have social psychological consequence: the organizational response to these incidents leaves the responsible controller feeling stigmatized and punished. Individuals who have "deals" are assigned remedial exercises and not allowed to work traffic. Even though the organizational response was designed to instruct and rehearse skills associated with the error to insure that it did not happen again, the social psychological consequence was the experience of having failed to live up to professional standards for the controller.

Controllers learned not to fear accident, but to fear mistakes that violated the rules of separation. Recall that controllers are not trained to avoid collision per se; they are trained to avoid a violation of the rules of separation, a practice that builds in safety by structuring their corrective practices while airplanes are still at a sufficient distance that collision can be avoided. The training process that produces controllers makes abiding by the rules of separation a priority. Moreover, the very language and surveillance technology in use assures a constant sensitivity to error and mistake. This message is continually reinforced. A technique started in 2002 uses a new technology that visually reconstructs real incidents of good controlling, Operational Errors, and accidents that include the pilot-controller conversations as an incident unfolds. All controllers are required to watch, which they do in small groups, then discuss how the controller handled the situation and possible alternatives.

As the monthly tally of Operational Errors and Runway Incursions is circulated to all facilities, showing year-to-date figures and how they rank in number of errors, controllers are reminded of the status implications of errors for their facility. Another more detailed description circulated across facilities lists errors by facility with a sentence or two about the circumstances. One controller showed me the monthly list in the Read-and-Initial file. She said, "Look at this. Isn't this ridiculous? I can't learn anything from this because they don't give enough detail about what happened. Why should I bother to read it?" She was correct about the insufficient detail, but the message that got through whether controllers read it or not was that errors matter in this business: don't make them.

## IMPLICATIONS

I have focused on two compliance-oriented organizations and their preventive strategies. I have noted

the differences in their technologies of control and the unintended consequences of those differences, as they generate vastly different cultural understandings about risk. We need to bear in mind that these two case studies undoubtedly have distinguishing characteristics. Although all organizations systematically produce incidents deviating from expectations, not all organizations try to proactively reduce risks, and not all try to measure and assess risk. How organizations respond to anomalies will depend on both the organizational and institutional contexts of work and important properties of the anomalies themselves: their "seriousness, frequency, number and types of persons and activities involved, cumulative impact, and rectifiability" (Star and Gerson 1986: 150). Uncorrected anomalies at NASA and the FAA's National Air Traffic System can result in public failures with high costs. Thus, both agencies devote enormous resources - time, energy, personnel, money - to regulating risk. Both have compliance systems for regulating risk, although the compliance system of the FAA is alone in also having a deterrent strategy when mistakes result in Operational Errors. Although their investigation and subsequent retraining of a controller charged with a violation was designed to bring the controller into compliance with professional standards, the social psychological consequence is one of punishment and stigma for the responsible individual. Both agencies systematically employ the principles of science to identify, measure, and classify anomalies. These conditions will vary across organizations, so we cannot assume that all anomalies have trajectories or that all technologies of control have ritualistic properties with social psychological consequences.

A second consideration is that I dichotomized the risky technologies that these agencies regulate - airplanes and space shuttles - as either certain or uncertain. This conceptualization was developed inductively during the analysis, dictated by the characteristics of the organizations that I compared. It is a fruitful distinction for these cases. In other cases, however, the organization may not be assessing the risk of some technology. Under these circumstances, variation in the certainty/uncertainty of tasks alone is probably sufficient. For example, the technologies of control employed by hospitals diagnosing patients and caseworkers interpreting information about possible abuse of foster children would seem to have much in common with each other and with the NASA case when it comes to the problem of identifying and interpreting possible signals of potential

danger. In such comparisons, other aspects of the social organization of risk assessment need to be taken into account: I found, for example, important differences in this comparison in the time available for interpreting information and the division of labor.

Third is the question of the origin of technologies of control. The FAA and NASA are public agencies, further distinguished by their status as monopolies. The technologies of control they employ cannot so readily be imported from other organizations. Instead, their technologies of control have been crafted specifically to suit the needs of their particular risky vehicles. In some organizations, however, the technologies of control may be imported from other organizations with similar tasks. We know little about the frequency of this importation process or what happens when one organization "borrows" a technology of control from some other organization. Institutional theorists raise the question of how organizational fields come to look alike, arguing that organizations respond to normative, mimetic, and coercive forces in their environments by changing structure and process in order to attain legitimacy (Powell and DiMaggio 1991). However, these adaptations may not be aligned with organizational goals of efficiency and effectiveness (Meyer and Rowan 1977). Joyce's (2000) analysis of hospital Magnetic Resonance Imaging technology is an example. Joyce found that MRI technology was adopted by hospitals often as a symbol of legitimacy rather than for its superior diagnostic properties, and currently is used in numerous hospitals even in the absence of evaluation research that compares its effectiveness and hazards with other technologies of control (e.g., clinical diagnosis, ultrasound) used in the identification, measurement, risk assessment, and the formal categorization of anomalies. When technologies of control are imported, what kinds of re-organization ensue, how do they mesh or not mesh with existing strategies, and what might be the unanticipated consequences at the social psychological level for people with the responsibility for interpreting anomalies?

A fourth consideration is the rituals of risk and error invoked *in the wake of accidents*. Historically, the FAA and NASA have had similar post-accident rituals. Their mistakes are public and immediately known. When accidents happen, national accident investigation teams are brought in. Their charge is to gather evidence, identify the cause or causes, and write an accident investigation report. Both the investigation and whatever

changes are implemented as a consequence are also technologies for control! These post-accident investigations look for possible indications of technical or human failure. The search for the latter is oriented by Human Factors Analysis, which targets individuals: was the controller tired, was the pilot poorly trained, was the manager incompetent, did the engineer or technician make a flawed decision or an incorrect manoeuvre, and if so, why? Unsurprisingly, then, investigations at both NASA and the NATS have in common the tendency to pay insufficient attention to the effect of institutional or organizational factors on personnel. Operator Error is a common finding in investigations of accidents and near misses (Perrow 1984; Sagan 1993), which Sagan referred to as the “politics of blame”. Causes must be identified, and in order to move forward, the organization must appear to have resolved the problem. The direct benefit of identifying the cause of an accident as Operator Error is that the individual operator is the target of change. The responsible party can be transferred, demoted, retrained, or fired, all of which obscure flaws in the organization that may have contributed to mistakes made by the individual who is making judgments about risk (Vaughan, 1999). The dark side of this ritualistic practice is that organizational and institutional sources of risk and error are not systematically subject either to investigation or other technologies of control.

Here is how scholars can be helpful. All organizations that seek to identify anomalies and reduce the possible harmful outcomes of failure might benefit from research examining their technologies of control and evaluating the consequences - anticipated and unanticipated - for risk and error. Research can take several directions. First, we need more systematic data on signals of danger and their interpretation, as the situation is defined at the moment, not post-accident. Following Turner’s identification of long incubation periods littered with early warning signs that precede accidents, more micro-level research needs to examine the trajectory of anomalies, how they are identified, officially categorized and defined. The most appropriate method would be case studies where the trajectory of information and the process of definition negotiation can be analyzed. Second, scholars must not study these matters in a vacuum, but examine also the social context: how might the behavior of individuals and the choices they make be affected by the organization itself and its political and economic environment? Third, little research has examined the ritualistic properties of technologies of control

and their possible unintended consequences. Scholars, organization administrators, and other professions dealing with risky technologies have gone far in identifying and analyzing technologies of control in use for both regulators and regulated. However, the organizational system, cultural understandings, and their effects at the social psychological level are so far unacknowledged and thus unexplored.

*Epilogue:* After I wrote this essay, Space Shuttle Columbia disintegrated in the sky over Texas on February 1, 2003. Immediately, I was called by media representatives looking for expertise on NASA and accidents. In April, I was called to testify before the Columbia Accident Investigation Board, at which point I found myself participating in a post-accident ritual of risk and error! I was surprised and delighted to discover that the Board and its staff (n=117) all had been reading research from social scientists on risk, accident, and safety. They were not bound to the Human Factors model; indeed, they were determinedly investigating NASA's organization system and its political and economic environment, as those factors affected decisions made about the foam debris problem, which seemed then and now to be the most likely technical cause of the accident. As I write, my participation continues, as I am consulting with the board on the data analysis and writing of the report. Thanks to social science research, their report will *not* place the blame on Operator Error. How their recommendations get implemented at NASA, however, is yet to be determined.

## BIBLIOGRAPHY

- Bourrier, Mathilde. Le Nucleaire a l'Epreuve de l'Organisation. Coll. Le Travail Humain. Paris: Presses Universitaires de France. 1999.
- Collins, Harry. Changing Order: Replication and Induction in Scientific Practice. London: Sage, 1985.
- Hammack, J.B., and M. L. Raines. Space Shuttle Safety Assessment Report. Johnson Space Center, Safety Division, 5 March 1981. National Archives, Washington D.C.
- Hawkins, Keith O. Environment and Enforcement: Regulation and the Social Definition of Enforcement. New York: Oxford University Press, 1984.
- Joyce, Kelly A. The Transparent Body: Magnetic Resonance Imaging, Knowledge, and Practices. Unpublished Ph.D. Dissertation, Department of Sociology, Boston College, 2000.
- Latour, Bruno. Science in Action. Cambridge: Harvard University Press, 1987.
- Marcus, Alfred. "Risk, Uncertainty, and Scientific Judgment," Minerva 26, 1988: 138-52.
- La Porte, Todd. "A Strawman Speaks Up," Journal of Contingency and Crisis Management," 1994, 4: 60-72.
- Pinch, Trevor and Weibe Bijker. "The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other." Social Studies of Science 14, 1984: 399-441.
- Reiss, Albert J. "Selecting Strategies of Social Control Over Organizational Life." In Enforcing Regulation, eds. Keith Hawkins and John M. Thomas. Boston: Kluwer-Nijhoff, 1984.
- Sagan, Scott. The Limits of Safety. Princeton University Press, 1993.
- Suchman, Lucy A. Plans and Situated Actions. New York: Cambridge University Press, 1987.
- Turner, Barry A. Man-Made Disasters. London: Wykeham, 1978.

Turner, Barry A. and Nick Pidgeon. Man-Made Disasters. 2nd ed. London: Heinemann- Butterworth.

1997

Vaughan, Diane. The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA.

Chicago: University of Chicago Press, 1996.

Vaughan, Diane. "Rational Choice, Situated Action, and the Social Control of Organizations." Law &

Society Review 32, 1: 1998: 23-61.

Vaughan, Diane, "Signals and Interpretive Work: The Role of Culture in a Theory of Practical Action."

In Culture in Mind: Toward a Sociology of Culture and Cognition, Karen A. Cerulo, ed., New York:

Routledge 2002: 28-54.

Vaughan, Diane. "The Dark Side of Organizations: Mistake, Misconduct, and Disaster," Annual Review of

Sociology, Vol. 25, 1999: 271-305.

Wynne, Brian. "Unruly Technology: Practical Rules, Impractical Discourses, and Public Understanding." Social

Studies of Science Vol. 18, 1988: 147-67.

## NOTES

i. This broad usage follows Staudenmeier's definition of technology as 1) a form of knowledge about artifacts and how to use them, 2) activities associated with technology, and 3) hardware. I have added rules and procedures, an additional consistent with his conceptualization.

ii. Some variation in technology-in-use exists between facilities, however. In the US, surveillance hardware will vary according to whether it is used by a tower, a TRACON, or a Centre. Internationally, between National Air Transportation Systems, hardware technologies of control variation due to different designers, manufacturers and ability and willingness of nation states to invest resources in their systems.