Challenger: Fine-Tuning the Odds Until Something Breaks

William H. Starbuck and Frances J. Milliken

Abstract*:* The Challenger disaster illustrates the effects of repeated successes, gradual acclimatization, and the differing responsibilities of engineers and managers. Past successes and acclimatization alter decision-makers' beliefs about probabilities of future success. Fine-tuning processes result from engineers' and managers' pursuing partially inconsistent goals while trying to learn from their experiences. Fine-tuning reduces probabilities of success, and it continues until a serious failure occurs.

Tragedy from the Commonplace

On 28 January 1986, the space shuttle Challenger disintegrated, killing all seven of its crew, including 'America's teacher in space', Christie McAuliffe. Those who watched are unlikely to forget their feelings of disbelief and horror. The disaster destroyed more than a space shuttle: it destroyed a part of America’s vision as well (Schwartz, 1987).

The American public, like NASA's managers, had grown complacent about the shuttle technology. We assumed the 25th launch would succeed because the previous 24 launches had succeeded. NASA had produced a long string of successes in the face of hypothetically low probabilities of success, and one result seems to be that both NASA and the American public developed a conviction that NASA could always succeed. The disaster suddenly reawakened us to the technology's extreme complexity and high risk. The ensuing investigation into the causes of the accident reminded us how unrealistic and error-prone organizations can be.

Neither Morton-Thiokol nor NASA could be called a typical organization, but their behaviours preceding the Challenger accident had many characteristics that we find commonplace in organizations. Organizations often communicate imperfectly, make errors of judgement, and provide playing fields for control games. Organizations often interpret past successes as evidencing their competence and the adequacy of their procedures, and so they try to lock their behaviours into existing patterns. Organizations often try to generalize from their experiences. Organizations often evolve gradually and incrementally into unexpected states.

Although these patterns of behaviour do occur commonly, we have good reason to fear their consequences when organizations employ high-risk technologies on a day-to-day basis (Perrow, 1984). In such organizations, these normal patterns of behaviour create the potential for tragedy. At the same time, the normality of Thiokol's and NASA's behaviours implies that we should be able to learn lessons from this experience that apply elsewhere.

Drawing on testimony before the Presidential Commission and reports in newspapers and magazines, this article seeks to extract useful lessons from the Challenger disaster. Because it has been investigated so exhaustively, the disaster affords a rich example that illustrates a variety of issues. But the authors of this article believe that the most important lessons relate to the effects of repeated successes, gradual acclimatization, and the differing responsibilities of engineers and managers. Both repeated successes and gradual acclimatization alter decision-makers' beliefs about probabilities of future success; and thereby, they may strongly influence decisions concerning high-risk technologies. These decisions occur in contexts that shift as people try to extract lessons from experience and in organizational arenas where engineers and managers represent somewhat conflicting points of view.

The next section frames the issues in terms of three theories about the ways past successes and acclimatization alter probabilities of future success. Two ensuing sections portray some effects of repeated successes and acclimatization at NASA: the first of these sections details the evolution of problems with joints in the cases of solid rocket boosters, and the second sketches some long-term changes in NASA's general culture. The fifth section then describes fine-tuning processes that result from engineers' and managers' pursuing partially inconsistent goals while trying to learn from their shared experiences. Fine-tuning reduces probabilities of success, and it goes on until a serious failure occurs. The final section comments on our ability to learn from disasters.

Three Theories about Probabilities of Future Success

Before the launch of 28 January 1986, many so-called experts had attempted to assess the riskiness of shuttle launches. For example, serious estimates of the probability that the solid rocket booster (SRB) would fail had ranged from as high as 1 in 10 to as low as 1 in 100,000 (Diamond, 1986a; Feynman, 1986; Sanger, 1986c). The wide range of such estimates casts doubt on their validity. Probably, no one, including the estimators themselves, really believed firmly in any of the estimates: at best, they amount to hypotheses that incorporate multitude assumptions. Not only do engineers and managers have to use actual experiences to appraise these hypotheses and the assumptions underlying them, but experience produces changes in hardware, procedures, and knowledge that alter the probabilities being estimated.

Faced with such elusive targets, engineers and managers have to frame specific hypotheses about riskiness within overarching theories about the effects of experience. They might plausibly adopt any of three macro theories. Theory 1 predicts that neither a success nor a failure changes the probability of a subsequent success. Theory 2 predicts that a success makes a subsequent success less likely, and that a failure makes a subsequent success more likely. Theory 3 predicts that a success makes a subsequent success more likely, and that a failure makes a subsequent success less likely.

Theory 1: Neither Success nor Failure Changes the Expected Probability of a Subsequent Success

Statisticians frequently use probability distributions that assume repeated events have the same probabilities. For instance, they assume that all flips of a coin have the same probability of turning up heads, or that all rolls of a die have the same probability of yielding sixes. Indeed, at one time, statisticians applied the label 'gambler's fallacy' to the idea that probabilities increase or decrease in response to successes or failures; such pejorative labeling fostered the notion that constant probabilities are not just convenient simplifications but absolute truths.

Engineers or managers who have studied statistics might well apply constant probability theories to the situations they face, and they might look with skepticism upon any interpretations that assume changing probabilities. According to Theory 1, the fact that NASA had launched shuttles successfully 24 times in a row ought to be disregarded when deciding whether to proceed with the 25th launch because the probability of failure by a solid rocket booster, or any other component, would be approximately the same on the 25th launch as on the first launch.

Richard P. Feynman compared shuttle launches to Russian roulette (Presidential Commission, 1986, I-148). Building on this analogy, Howard Schwartz (1987: 61-62) remarked: 'In the case of Russian roulette, with one round in the cylinder, the odds are one in six that a pull on the trigger will fire the round. If the round does not fire on the first pull, and the cylinder is spun, the odds are again one in six for the next pull on the trigger. To some persons unfamiliar with the theory of probability, it may seem that the odds with each successive pull would be greater. This is of course wrong. But it is equally wrong to suppose that the odds will be less with each successive event. This is what the NASA officials appeared to believe. The question is, how can it have happened that NASA officials, knowing full well the laws of probability, could have made such an error?'

Although Schwartz cited an often-used statistical model, no laws compel probabilities to remain constant over time. The probability of an event may rise over time or fall, depending on what changes occur in factors that influence this probability. The probability of a pistol's firing may well remain constant throughout several successive spins of the cylinder if the person spinning the cylinder behaves consistently. But Russian roulette may not be a good analogy for shuttle launches because the shuttle's hardware and personnel and operating procedures do change from launch to launch, and the probability of a successful flight may not stay constant.

For the probability of success with a sociotechnical system to stay constant, either the hardware, procedures, and operators' knowledge have to remain substantially unchanged over time, or changes tending to raise the probability of success have to be offset by changes tending to lower it. When a sociotechnical system's probability of success is low, people rarely leave hardware and procedures alone. Thus, a high-risk sociotechnical system should not have a probability of success that remains constant. Although some changes may well offset each other where numerous changes occur simultaneously, as during a period of initial development, the engineers and managers guiding those changes are expecting to raise the overall probability of success, and so they are unlikely to expect the probability of success to stay constant.

Engineers or managers might, however, hypothesize a constant probability of success for a sociotechnical system that appears nearly certain to succeed. And engineers and managers who have successfully launched 24 consecutive shuttles might well infer that the next flight has a very, very high probability of success, either because this probability has been very high all along or because it has risen over time. A number of the statements by Thiokol and NASA personnel suggest they believed the Challenger's probability of success was already so high that they had no need to raise it further.

Theory 2: Success Makes a Subsequent Success Seem Less Probable, and Failure Makes a Subsequent Success Appear More Likely

A series of successes, or even a single successful experience, might induce engineers or managers to lower their estimates of the probability of a future success; and conversely, failures might induce engineers or managers to raise their estimates of the probability of a future success. Schwartz alluded to such a theory when he conjectured that a player might expect a slot machine that has not paid off recently to become ready to pay off. Such a player might also expect a slot machine that has just paid off to have a bias against paying off again in the immediate future.

Applied to sociotechnical systems, Theory 2 emphasizes complacency versus striving, confidence versus caution, inattention versus vigilance, routinization versus exploration, habituation versus novelty. Successes foster complacency, confidence, inattention, routinization, and habituation; and so human errors grow increasingly likely as successes accumulate. Failures, on the other hand, remind operators of the need for constant attention, caution, and vigilance; and so failures make human errors less likely. For instance, Karl Weick (1987: 118-19) pointed out: ‘When people think they have a problem solved, they often let up, which means they stop making continuous adjustments. When the shuttle flights continued to depart and return successfully, the criterion for a launch – convince me that I should send the Challenger – was dropped. Underestimating the dynamic nature of reliability, managers inserted a new criterion – convince me that I shouldn't send Challenger.’ Similarly, Richard Feynman interpreted NASA's behaviour according to Theory 2: after each successful flight, he conjectured, NASA's managers thought ‘We can lower our standards a bit because we got away with it last time’ (Presidential Commission, 1986, I-148).

Failures also motivate engineers and managers to search for new methods and to try to create systems that are less likely to fail, and successes may induce engineers and managers to attempt to fine-tune a sociotechnical system – to render it less redundant, more efficient, more profitable, cheaper, or more versatile. Fine-tuning rarely raises the probability of success, and it often makes success less certain. Because fine-tuning seems to be a very important process that has received little attention, a later section of this article looks at it again.

The participants in sociotechnical systems often espouse Theory 2 after failures, conjecturing that past failures will elicit stronger efforts or greater vigilance in the future. However, participants find it difficult to use Theory 2 to interpret their own responses to successes. One reason is that participants may not recognize that repeated successes nurture complacency, confidence, inattention, routinization, and habituation. Another reason is that, when they do notice such changes, participants tolerate them on the premise that they are merely eliminating unnecessary effort and redundancy, not making success less probable. Indeed, because accusations of complacency and inattention seem derogatory, participants might punish a colleague who voices Theory 2. Thus, when applied to successes, Theory 2 is more an observer's theory than a participant's theory. Although bosses might use Theory 2 when appraising their subordinates' actions, they would probably not apply it to themselves.

Theory 3: Success Makes a Subsequent Success Appear More Probable, and Failure Makes a Subsequent Success Seem Less Likely.

The participants in sociotechnical systems espouse Theory 3 readily, because it is easy to believe that success demonstrates competence, whereas failure reveals deficiencies.

Expected probabilities of success are not well-defined facts, but hypotheses to be evaluated through experience. Even if engineers or managers believe that a probability of success remains constant for a long time, they need to revise their estimates of this probability as experience accumulates. Engineers or managers with statistical training might, for example, use hypothetical computations to formulate an initial estimate of a probability of success and then apply Bayes' Theorem to compute successive estimates of this probability: if so, each success would raise the expected probability of success, and each failure would lower this expected probability.

Furthermore, experience with a technology may enable its users to make' fewer mistakes and to employ the technology more safely, and experience may lead to changes in hardware, personnel, or procedures that raise the probability of success. Studies of industrial learning curves show that people do perform better with experience (Dutton and Thomas, 1984). Better, however, may mean either more safely or less so, depending on the goals and values that guide efforts to learn. If better means more cheaply, or quicker, or closer to schedule, then experience may not raise the probability of safe operation.

Explaining that experience produces both advantages and disadvantages, Starbuck (1988) commented:

These learning mechanisms – buffers, slack resources, and programs – offer many advantages: they preserve some of the fruits of success, and they make success more likely in the future. They stabilize behaviors and enable organizations to operate to a great extent on the basis of habits and expectations instead of analyses and communications. They reduce the complexity of social relations and keep people from disobeying or behaving unpredictably. They minimize needs to communicate or to reflect, and they conserve analytic resources. They also give organizations discretion and autonomy with respect to their environments. Organizations do not have to pay very close attention to many of the demands currently arising from their environments, and they do not have to formulate explicit or unique responses to most of these demands. Thus, organizations gain human resources that they can devote to influencing their environments and creating conditions that will sustain their successes in the future.

But these learning mechanisms also carry disadvantages. In fact, each of the advantages has a harmful aspect. People who are acting on the basis of habits and obedience are not reflecting on the assumptions underlying their actions. People who are behaving simply and predictably are not improving their behaviors or validating their behaviors' appropriateness. Organizations that do not pay careful attention to their environments' immediate demands tend to lose track of what is going on in those environments. Organizations that have discretion and autonomy with respect to their environments tend not to adapt to environmental changes; and successful organizations want to keep their worlds as they are, so they try to stop social and technological changes. Indeed, buffers, slack resources, and programs make stable behaviors, current strategies, and existing policies appear realistic by keeping people from seeing problems, threats, or opportunities that would justify changes.

Theory 3 in Action

Theory 3 offers a very plausible characterization of the beliefs of managers at Thiokol's Wasatch Division and NASA's Marshall Space Flight Center (SFC) as they tried to evaluate the risks posed by joints in the shuttle's solid rocket booster (SRB). As successful launches accumulated, these managers appear gradually to have lost their fear of design problems and grown more confident of success. One must understand their story in some detail, however, in order to appreciate the complexity and ambiguity of the technical issues, the managers' milieu, and the slow progression in their beliefs.

Thiokol's engineers based the design of the shuttle's SRB on the Air Force's Titan III because of the latter's reliability. The Titan's case was made of steel segments, with the joints between segments being sealed by rubber O-rings. The Titan's O-rings had occasionally been eroded by the hot gases inside the engine, but Thiokol's engineers did not regard this erosion as significant. Nevertheless, to make the shuttle's SRB safer, Thiokol's engineers put a second, presumably redundant O-ring into each joint.

However, a 1977 test of the SRB's case showed an unexpected 'rotation' of the joints when the engine ignited: this rotation decompressed rather than compressed the O-rings, making it more difficult for the O-rings to seal the joints, and increasing the chance that hot gases would reach the O-rings. This alarmed NASA's engineers, so they asked for a redesign of the joints. Thiokol did not redesign the joints qualitatively, but did enlarge the O-rings to 0.028 inches diameter and thicken the shims that applied pressure on the O-rings from outside. In 1980, a high-level review committee reported that NASA's specialists had 'found the safety factors to be adequate' and the joints 'sufficiently verified with the testing accomplished to date' (Presidential Commission, 1986, I-125). The joints were classified as Criticality 1R: the 1 denoted that joint failure could cause a loss of life or the loss of a shuttle; the R denoted that the secondary O-rings provided redundancy. That is, the secondary O-rings served as a back up for the primary O-rings.

Eight full-scale tests of SRBs yielded no sign of joint problems, nor did the first shuttle flight. During the second flight in November 1981, hot gases eroded one O-ring, but this event made little impression: NASA's personnel did not discuss it at the next flight-readiness review and they did not report it upward to top management. The three flights during 1982 produced no more evidence of O-ring problems.

In 1982, an engineer working for Hercules, Inc. proposed a new joint design: a 'capture lip' would inhibit joint rotation. NASA's engineers thought this proposal looked interesting, but the capture lip would add 600 pounds to each SRB, its practicality was untested, and a more complex joint might harbour unforeseen difficulties. It would take over two years to build SRBs with this design. NASA decided to continue using the old joint design and to award Hercules a contract to develop the new design in conjunction with a new case material, carbon filaments in epoxy resin (Broad, 1986c).

Thiokol too was proposing changes in the SRBs, but these were intended to raise the rockets' efficiency. During 1983, NASA began using SRBs that incorporated three incremental improvements (Broad, 1986c; Marbach et al. 1986). Thiokol made the SRBs' walls 0.02-0.04 inches thinner; they narrowed the nozzles; and they filled the rockets with more powerful fuel. Thinner walls saved several hundred pounds that could be replaced by payloads. More powerful fuel could lift more weight. Smaller nozzles extracted more thrust from the fuel.

These changes, however, made the SRB less durable and exacerbated the joint rotation. More powerful fuel and smaller nozzles raised the SRBs' internal pressures, and thinner walls flexed more under pressure, so the joints developed larger gaps upon ignition. Tests showed that joint rotation could grow large enough to prevent a secondary O-ring from sealing a joint and providing redundancy. Therefore, the R was dropped from the joints' Criticality classification, but the reclassification document, written by a Thiokol engineer, implied the risk was small:

> To date, eight static firings and five flights have resulted in 180 (54 field and 126 factory) joints tested with no evidence of leakage. The Titan III program using a similar joint concept has tested a total of 1076 joints successfully.
>
> A laboratory test program demonstrated the ability of the O-ring to operate successfully when extruded into gaps well over those encountered in this O-ring application (Presidential Commission, 1986, I-241).

The Presidential Commission (1986, I-126) surmised 'that NASA management and Thiokol still considered the joint to be a redundant seal even after the change from Criticality 1R to 1'. Over the next three years, many documents generated by NASA and Thiokol continued to list the Criticality incorrectly as 1R. Neither management really thought that a secondary O-ring might fail to seal a joint. In the view of Joseph C. Kilminster, manager of Thiokol's space boosters programme, 'it had to be a worse-case stack-up of tolerances, which statistically you would not expect' (Bell and Esch, 1987: 45).

Also in 1983, the ninth full-scale test of an SRB and the sixth shuttle flight both produced signs of heat damage. As with the second flight, the NASA personnel did not discuss this damage at the flight-readiness review for the next flight or report it to top management, but this damage may have triggered changes in testing procedures. Up to August 1983, NASA leak-checked both the nozzle joints and the other (field) joints with an air pressure of 50 psi in order to verify that the O-rings had been installed correctly. In August 1983, NASA raised the leak-check pressure for field joints to 100 psi; and in January 1984, they raised it to 200 psi. Similarly, NASA raised the leak-check pressure for nozzle joints to 100 psi starting in November 1983, and to 200 psi starting in April 1985. According to Lawrence B. Mulloy, manager of the SRB project at Marshall SFC, NASA boosted the test pressures in order to force the secondary O-rings into the gaps between adjoining case segments.

NASA and Thiokol finally did review the O-ring problems on flights two and six in February 1984, after the tenth shuttle flight showed erosion of O-rings on both SRBS. At that point, engineers at both NASA and ' Thiokol conjectured that the higher leak-check pressures were creating problems rather than preventing them: the leak checks might be blowing holes in the putty that sealed cracks in the SRBs' insulation and creating paths by which hot gases could reach the O-rings. Laboratory tests suggested, however, that larger holes in the insulating putty might produce less damage than smaller holes, and the tests indicated the O-rings ought to seal even if eroded as much as 0.095 inches. Thiokol's engineers made a computer analysis that implied the primary O-rings would be eroded at most 0.090 inches, just under one-third of their diameter. 'Therefore', concluded the formal report, 'this is not a constraint to future launches' (Presidential Commission, 1986, I-128-32).

Mulloy then introduced the idea that some erosion was 'acceptable' because the O-rings embodied a safety factor (Presidential Commission, 1986, II-Hl). This notion was discussed and approved by NASA's top managers at the flight-readiness review on 30 March 1984.

After the flight launched, on 6 April, 1984 also showed some O-ring erosion, Hans Mark, NASA's second-in-command at that time, asked Mulloy to submit a written report on joint sealing (Sanger, 1986e). Mulloy, in turn, asked Kilminster to conduct this study, and especially to investigate leak checking and the material being used for insulating putty. Thiokol proposed some tests, NASA approved this proposal, but Thiokol did not carry the tests out, and Mark received no report.

Figure 10.1 graphs NASA's observations of joint problems over time. The vertical axis indicates the numbers of joints in which NASA found problems. A short bar below the horizontal axis denotes an absence of evidence. Fractional bars above the horizontal axis symbolize small traces of gas leakage or heat damage. To reflect its seriousness, damage to secondary O-rings is represented by bars that are four times as long as those for damage to primary O-rings or for blow-by (gas leakage).

Figure 10.1: O-ring problems over time

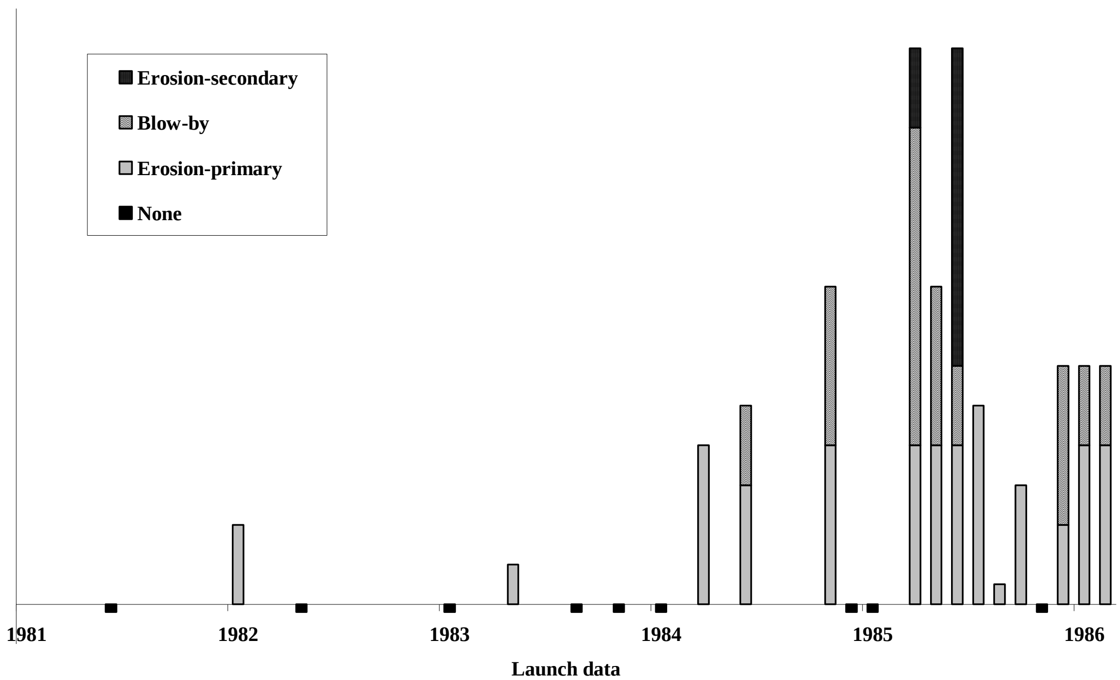

In all, inspectors discovered heat damage to SRB joints after three of the five flights during 1984, after eight of the nine flights during 1985, and after the flight on 12 January 1986. Not graphed are two test firings of SRBs that also produced O-ring erosion. The inspectors found damage in all joints of the SRBs, but the nozzle joints were especially vulnerable, sustaining damage during twelve of the fifteen flights. Not only was heat damage becoming more frequent as time passed, but the inspectors were seeing larger and larger amounts of erosion, including secondary O-rings as well as primary ones. Mulloy and his Thiokol counterparts seem to have grown more and more secure in the face of this evidence. In September 1984, Mulloy spoke of 'allowable erosion'; and in February 1985, Mulloy and Thiokol characterized joint leakage as an 'acceptable risk'.

The fifteenth flight in January 1985 experienced substantial O-ring damage: hot gas blew by the O-rings in two joints on each SRB, and the heat eroded one O-ring on each SRB. Further, this was the first flight in which a secondary O-ring was damaged. When the flight took off, the ambient temperature at the launch site was only 53 degrees. This event led Thiokol to propose that 'low temperature enhanced the probability of blow-by' (Presidential Commission, 1986, I-136), which was the first time that idea had been introduced. However, even more serious O-ring damage occurred during the seventeenth flight in April 1985, when the temperature at launch was 75 degrees. On this occasion, one primary O-ring eroded 0.171 inches, a substantial amount of hot gas blew by this O-ring, and so its back-up secondary O-ring eroded 0.032 inches. The 0.171 inches represented 61 per cent of the primary O-ring's diameter; and the evidence suggested that the primary O-ring had not sealed until two minutes after launch (Bell and Esch, 1987: 43).

One consequence of these events was that NASA's top management sent two representatives to Marshall SFC to review the O-ring problems, and these visitors asked Marshall to provide a briefing in Washington. This briefing, on 19 August, concluded that 'it is safe to continue flying existing design as long as all joints are leak checked with a 200 psi stabilization pressure, are free of contamination in the seal areas and meet O-ring squeeze requirements' (Presidential Commission, 1986, I- 140). A second consequence was that Mulloy placed a 'launch constraint' on all subsequent flights. In NASA's jargon, a launch constraint is an acknowledgment that a problem of Criticality 1, IR, 2 or 2R might actually occur. The O-rings were still Criticality 1. To Mulloy, this launch constraint evidently did not mean that a flight should not occur; it only meant that each flight-readiness review should include a rationalization for waiving the constraint. During the flight-readiness review on 2 July, Mulloy explained why some erosion of secondary O-rings was also acceptable, and he described the O-ring-erosion problems as 'closed'. Mulloy proceeded to 'waive' this launch constraint for all subsequent flights up to the last one (Bell and Esch, 1987: 43; Boffey, 1986c). Mulloy later reflected: ‘Since the risk Of O-ring erosion was accepted and indeed expected, it was no longer considered an anomaly to be resolved before the next flight. . . . I concluded that we're taking a risk every time. We all signed up for that risk. And the conclusion was, there was *no* significant difference in risk from previous launches. We'd be taking essentially the same risk on Jan. 28 that we have been ever since we first saw O-ring erosion’ (Bell and Esch, 1987: 43, 47).

Kilminster seems to have concurred with Mulloy. In April 1985, NASA reminded Kilminster that Thiokol was supposed to have studied joint sealing. Kilminster then set up an informal task force that, in August 1985, proposed 20 alternative designs for the nozzle joints and 43 designs for the other (field) joints. At that point, Thiokol formalized the task force, but some members felt it was getting insufficient attention. One member, Roger M. Boisjoly, has subsequently said that Kilminster 'just didn't basically understand the problem. We were trying to explain it to him, and he just wouldn't hear it. He felt, I guess, that we were crying wolf' (Bell and Esch, 1987: 45).

Meanwhile, during the Autumn of 1984 and Spring of 1985, Hercules had successfully tested SRBs with carbon-epoxy cases and capture-lip joints. Simultaneously, laboratory tests were demonstrating that 'the capture feature was a good thing' (Broad, 1986c). In July 1985, Thiokol ordered 72 new steel case segments having such joints; the manufacturer was expected to deliver these in February 1987. Marshall's documentation for the briefing in Washington on 19 August 1985 described the capture feature as being a 'potential long-term solution' and projected that this solution might come into effect starting in August 1988.

During a telephone call in early December 1985, someone at Marshall SFC told a low-level Thiokol manager that Marshall's Director of Engineering wanted Thiokol to 'close out' the outstanding problems, especially those that had remained unresolved for six months or more. Thereupon, this Thiokol manager wrote to Thiokol's main liaison with NASA, asking for 'closure' of all O-ring-erosion problems and stating seventeen reasons. Explaining that the problems should be closed because they 'will not be fully resolved for some time', Thiokol's liaison, in turn, relayed this request to Mulloy's immediate subordinate. A NASA 'quality assurance' administrator then marked 'contractor closure received' on all of the Problem Reports itemized in the liaison's letter. As a result, the O-rings were no longer listed as a launch constraint when it came time for the flight-readiness review on 15 January, 1986, and O-rings were not even mentioned in the flight readiness documentation for the doomed Challenger. Indeed, just five days before the Challenger disaster, a NASA administrator marked these same Problem Reports 'problem is considered closed' (Presidential Commission, 1986, I-142-4, II-H3).

Yet on 16 and 17 January 1986, engineers from NASA and Thiokol met to review the O-ring problems and to discuss possible solutions. By that time, NASA or Thiokol personnel had proposed at least eleven alternative hypotheses:

(1) O-rings having round cross-sections did not put enough area against adjacent flat surfaces;
(2) some O-rings were being installed incorrectly, or
(3) some O-rings were smaller than the specified diameter; or
(4) bits of dirt or metal splinters kept some O-rings from sealing, and so leak checking should occur at an air pressure high enough to force the secondary O-rings into the correct positions and to assure that they sealed properly;
(5) high-pressure leak checking was displacing the primary O-rings from their proper positions, so causing them to fail to seal during launches;
(6) high-pressure leak checking was creating holes in the insulating putty;
(7) the primary O-rings were eroding because hot gases leaked through holes in the insulating putty;
(8) the primary O-rings might not seal unless they were pressurized by hot gases that leaked through holes in the insulating putty;
(9) the insulating putty had some unknown deficiencies;
(10) cold temperatures stiffened the insulating putty enough to keep it from responding to the high pressures inside the engine during firing; and
(11) cold temperatures stiffened the O-rings enough to keep them from sealing the joints.

A Can-Do Organization with an 'Operational' System

Theory 3 also affords a plausible description for NASA's general culture. Ironically, participants' belief in Theory 3 may make Theory 2 a more realistic one for observers.

Success breeds confidence and fantasy. When an organization succeeds, its managers usually attribute this success to themselves, or at least to their organization, rather than to luck. The organization's members grow more confident, of their own abilities, of their managers' skill, and of their organization's existing programmes and procedures. They trust the procedures to keep them apprised of developing problems, in the belief that these procedures focus on the most important events and ignore the least significant ones. For instance, during a teleconference on 2 7 January, 1986, Thiokol's engineers said that NASA should not launch the shuttle if the ambient temperature was below 53 degrees because no previous launch had occurred with an ambient temperature below 53 degrees. Lawrence Mulloy protested: '. . . there are currently no Launch Commit Criteria for joint temperature. What you are proposing to do is to generate a new Launch Commit Criteria on the eve of launch, after we have successfully flown with the existing Launch Commit Criteria 24 previous times' (Presidential Commission, 1986, I-96). Mulloy spoke as if he had come to trust the Launch Commit Criteria that had always produced successes.

In the perceptions of NASA's personnel, as well as the American public, NASA was not a typical organization. It had a magical aura. NASA had not only experienced repeated successes, it had achieved the impossible. It had landed men on the moon and returned them safely to earth. Time and again, it had successfully completed missions with hardware that supposedly had very little chance of operating adequately (Boffey, 1986a). NASA's managers apparently believed that the contributions of astronauts pushed NASA's 'probability of mission success very close to 1.0' (Feynman, 1986, FI). The Presidential Commission (1986, I-172) remarked: 'NASA's attitude historically has reflected the position that “We can do anything”'. Similarly, a former NASA budget analyst, Richard C. Cook, observed that NASA's 'whole culture' calls for 'a can-do attitude that NASA can do whatever it tries to do, can solve any problem that comes up' (Boffey, 1986b).

As Theory 2 holds, success also erodes vigilance and fosters complacency and routinization. The Presidential Commission (1986, I-152) noted that the NASA of 1986 no longer 'insisted upon the exactingly thorough procedures that were its hallmark during the Apollo program'. But the Apollo programme called for vigilance because it was a risky experiment, whereas NASA's personnel believed that the shuttle represented an 'operational' technology. The shuttle had been conceived from the outset not only as a vehicle for space exploration and scientific research, but as a so-called Space Transportation System (STS) that would eventually support industrial manufacture in orbit. According to NASA's formal announcements, this STS had supposedly progressed beyond the stage of experimental development long before 1986. In November 1982, NASA declared that the STS was becoming 'fully operational – meaning that the STS had proven sufficiently safe and error-free to become routine, reliable, and cost-effective. Directives issued in 1982 and 1984 specified 'a flight schedule of up to 24 flights per year with margins for routine contingencies attendant with a flight-surge capability'. NASA had actually scheduled fifteen flights for 1986.

In NASA's conception, an operational system did not have to be tested as thoroughly as an experimental one. Whereas NASA had tested equipment for the Apollo spacecraft in prototype form before purchasing it for actual use, NASA officials assumed that they had learned enough from the Apollo programme that the shuttle required no tests of prototypes. C. Thomas Newman, NASA's comptroller, has explained: 'The shuttle set out with some different objectives. To produce a system of moderate costs, the program was not as thoroughly endowed with test hardware' (Diamond, 1986c, B4). Far from saving money or time, this strategy actually produced a great many revisions in plans, delays that added up to over six years, and operating costs 53 times those projected 'when Congress had approved the programme (Diamond, 1986b). Richard Feynman

(1986) hypothesized that this strategy also contributed directly to the Challenger disaster by making the SRB difficult to test or modify. The fact is, however, that Thiokol's first eight full-scale tests disclosed no joint problems (Sanger, 1986d) – perhaps the tests were intended to prove that the agreed design could function satisfactorily rather than to disclose its limitations and potential deficiencies.

An operational system seemingly also demanded less day-to-day care. As the shuttle became operational, NASA's top managers replaced the NASA personnel who were inspecting contractors' work on-site with 'designated verifiers', employees of the contractors who inspected their own and others' work on NASA's behalf. This increasing trust could reflect improvements over time in the quality of the contractors' work, or reflect an accumulation of evidence that the contractors were meeting specifications, but it could also be interpreted as complacency. Also, NASA cut its internal efforts toward safety, reliability, and quality assurance. Its quality-assurance staff dropped severely from 1689 personnel in 1970 to 505 in 1986, and the biggest cuts came at Marshall SFC, where 615 declined to just 88 (Pear, 1986). These reductions not only meant fewer safety inspections, they meant less careful execution of procedures, less thorough investigation of anomalies, and less documentation of what happened. Milton Silveira, NASA's chief engineer, said: ‘In the early days of the space program we were so damned uncertain of what we were doing that we always got everybody's opinion. We would ask for continual reviews, continual scrutiny by anybody we had respect for, to look at this thing and make sure we were doing it right. As we started to fly the shuttle again and again, I think the system developed false confidence in itself and didn't do the same thing’ (Bell and Esch, 1987: 48).

Fine-Tuning the Odds

The foregoing sections show how repeated successes and gradual acclimatization influenced the lessons that NASA and Thiokol personnel were extracting from their shared experiences. These learning processes involved both engineers and managers, who were representing somewhat different points of view. The traditional differences in the responsibilities of engineers and managers give their interactions an undertone of conflict and make learning partly a process of fine-tuning the probabilities of success. Fine-tuning gradually makes success less and less likely.

Although an organization is supposed 'to solve problems and to achieve goals, it is also a conflict-resolution system that reconciles opposing interests and balances countervailing goals. Suppliers, customers, blue-collar and white-collar employees, executives, owners, neighbours, and governments all contribute resources to a collective pool, and then they all place claims upon this resource pool. Further, every serious problem entails real-world contradictions, such that no action can produce improvement in all dimensions and please all evaluators. For instance, an organization may seek to produce a high-quality product that assures the safety of its users, while also delivering this product promptly and earning a substantial profit. High quality and safety typically support strong demand; but high quality and safety also usually entail costs and slow down production; high costs imply high prices; and high prices and slow production may reduce revenues. Thus, the organization has to balance quality and safety against profit.

Opposing interests and countervailing goals frequently express themselves in intraorganizational labour specializations, and they produce intraorganizational conflicts. An organization asks some members to enhance quality, some to reduce costs, and others to raise revenue; and these people find themselves arguing about the trade-offs between their specialized goals. The organization's members may seek to maintain internal harmony by expelling the conflicts to the organization's boundary, or even beyond it. Thus, both Thiokol's members and NASA's members would normally prefer to frame a controversy as a disagreement between Thiokol and NASA rather than as a disagreement within their own organization. But conflicts between organizations destroy their compatibility, and an organization needs compatibility with its environment just as much as it needs internal cohesion. Intraorganizational conflict enables the organization to resolve some contradictions internally rather than letting them become barriers between the organization and its environment. Thus, on the evening of 27 January 1986, facing a conflict with NASA over the desirability of launching Challenger, Thiokol's Joseph Kilminster asked for a recess of the teleconference with NASA, so that the Thiokol personnel could caucus among themselves.

Thiokol's caucus began with Calvin G. Wiggins, general manager of the space division, asserting: 'We have to make a management decision'. Wiggins appears to have been pointing out that, whereas it had been engineers who had formulated Thiokol's recommendation against launching, the conflict with NASA was raising non-engineering issues that managers should resolve. Two engineers, Roger Boisjoly and Arnold R. Thompson, tried to restate to the managers present why they believed cold weather would make the SRB's joints less likely to seal. After a few minutes, Boisjoly and Thompson surmised that no one was listening to them, so they gave up and resumed their seats. The decision was evidently going to be made in a managerial arena.

The four vice presidents of Thiokol's Wasatch division then discussed the issue among themselves. Kilminster and Robert K. Lund, vice president for engineering, expressed their reluctance to contradict the engineers' position. At that point, Jerald E. Mason, senior vice president and chief executive of the Wasatch operations, urged Lund: 'rake off your engineering hat and put on your management hat'. The four managers then agreed to recommend launching. During this discussion, the managers agreed (a) that the primary O-rings possessed enough safety margin to enable them to tolerate three times the worst erosion observed up to that time and (b) that the secondary O-rings would provide a seal even if a primary O-ring failed. In fact, both assumptions had been contradicted in April 1985, when a primary O-ring had lost three-fifths of its diameter and neither the primary nor the secondary O-ring had sealed for two minutes.

The foregoing scenario illustrates an intraorganizational conflict that crystallizes around the differences between engineers and managers, and shows how these differences may rend a person who plays both an engineering role and a management role (Schriesheim et al., 1977).

Engineers are taught to place very high priority on quality and safety. If engineers are not sure whether a product is safe enough, they are supposed to make it much safer than they believe necessary. Facing uncertainty about safety, engineers would typically incorporate a safety factor of at least two – meaning that they would make a structure twice as strong as appeared necessary, or make an engine twice as powerful as needed, or make insulation twice as thick as required. Where failure would be very costly or additional safety would cost little, engineers might make a safety factor as large as ten. Thus, Thiokol's engineers were behaving according to the norm when they decided to put two O-rings into each joint: The second O-ring would be redundant if the shuttle's SRB operated much like the Titan's, but the design engineers could not be certain of this in advance of actual shuttle flights.

Safety factors are, by definition, supposed to be unnecessary. Safety factors of two are wasteful, and safety factors of ten very wasteful, if they turn out to be safety factors in truth. To reduce waste and to make good use of capacity, an organization needs to cut safety factors down.

People may cut safety factors while designing a sociotechnical system. Large safety factors may render projects prohibitively expensive or technically impossible, and thus may prevent the solving of serious problems or the attaining of important goals. When they extrapolate actual experiences into unexplored domains, safety factors may also inadvertently create hazards by introducing unanticipated risks or by taxing other components to their limits.

People are almost certain to reduce some safety factors after creating a system, and successful experiences make safety factors look more and more wasteful. An initial design is only an approximation, probably a conservative one, to an effective operating system. Experience generates information that enables people to fine-tune the design: experience may demonstrate the actual necessity of design characteristics that were once thought unnecessary; it may show the danger, redundancy, or expense of other characteristics; and it may disclose opportunities to increase utilization. Fine-tuning compensates for discovered problems and dangers, removes redundancy, eliminates unnecessary expense, and expands capacities. Experience often enables people to operate a sociotechnical system for much lower cost or to obtain much greater output than the initial design assumed (Box and Draper, 1969; Dutton and Thomas, 1984).

Although engineers may propose cost savings, their emphasis on quality and safety relegates cost to a subordinate priority. Managers, on the other hand, are expected to pursue cost reduction and capacity utilization, so it is managers who usually propose cuts in safety factors. Because managers expect engineers to err on the side of safety, they anticipate that no real risk will ensue from incremental cost reductions or incremental capacity expansions. And engineers, expecting managers to trim costs and to push capacity to the limit, compensate by making safety factors even larger. Top managers are supposed to oversee the balancing of goals against one another, so it is they who often make the final decisions about safety factors. Thus, it is not surprising to find engineers and managers disagreeing about safety factors, or to see top managers taking such decisions out of their subordinates' hands, as happened at Thiokol. Hans Mark has recalled: 'When I was working as Deputy Administrator, I don't think there was a single launch where there was some group of subsystem engineers that didn't get up and say "Don't fly". You always have arguments' (Bell and Esch, 1987: 48).

Formalized safety assessments do not resolve these arguments, and they may exacerbate them by creating additional ambiguity about what is truly important. Engineering caution and administrative defensiveness combine to proliferate formalized warnings and to make formalized safety assessments unusable as practical guidelines. In 1986, the Challenger as a whole incorporated at least 8000 components that had been classified Criticality 1, 2, or 3. It had 829 components that were officially classified as

Criticality 1 or 1R – 748 of them classified 1 rather than 1R. Each SRB had 213 of these 'critical items', 114 of which were classified 1 (Broad, 1986b; Magnuson, 1986: 18). Since no administrative apparatus could pay special and exceptional attention to 8000 issues, formalized Criticality had little practical meaning. To focus attention, NASA had identified special 'hazards' or 'accepted risks': The Challenger supposedly faced 277 of these at launch, 78 of them arising from each SRB. But if NASA's managers had viewed these hazards so seriously that any one of them could readily block a launch, NASA might never have launched any shuttles.

NASA's experience with the SRB's O-rings, as detailed above, looks like a typical example of learning from experience. Neither NASA's nor Thiokol's personnel truly understood in detail all of the contingencies affecting the sealing of joints. The Thiokol engineers imitated a joint design that appeared to have had no serious problems in the Titan's SRB, but they added secondary O-rings as a safety factor. The joints were formally classified Criticality 1R, and then 1, despite the Thiokol and NASA managers' conviction that a joint failure was practically impossible. Then actual shuttle flights seemingly showed that no serious consequences ensued even when the O-rings did not seal promptly and when primary O-rings sustained extensive damage and secondary O-rings minor damage. A number of managers surmised that, although an improved joint design should be adopted in due course, experience demonstrated the O-rings to be less dangerous than the engineers had initially assumed. But some engineers, at Marshall SFC as well as Thiokol, were drawing other conclusions from the evidence: Richard Cook told a reporter that propulsion engineers at Marshall had 'said to me, almost in a whisper in my ear, that the thing could blow up . . . one of them said to me, "When this thing goes up, we hold our breath"' (Boffey, 1986b).

The 1983 changes in the SRB also made sense as fine-tuning improvements after successful experience. These looked small at the time: they trimmed the SRB's weight by only 2 per cent and boosted its thrust by just 5 per cent. Similar incremental changes might, in principle, continue indefinitely as people learn and as better materials become available. For instance, NASA was hoping to obtain a further SRB weight reduction by shifting to a graphite-epoxy case (Sanger, 1986a). However, the SRB changes in 1983 illustrate also that small, incremental changes may produce small, incremental effects that are very difficult to detect or interpret.

Thus, some of the key decisions that doomed a shuttle may have occurred in 1982, when NASA endorsed Thiokol's proposed improvements of the SRBs. Thiokol's revised design had more joint rotation than the initial one, and thinner cases might have been more distorted by use. Moreover, other changes reinforced the importance of joint rotation. In particular, used segments of the SRB cases came back slightly out-of-round, so the segments did not match precisely and the O-rings were being expected to seal uneven gaps. Yet, NASA re-used case segments more often over time, and Kennedy Space Center stopped inspecting O-rings.

The Presidential Commission (1986, I-133-4) focused attention on a different sequence of fine-tuning changes: the increases in leak-check pressures from 50 psi to 200 psi. The Commission pointed out that the test pressures correlated with the frequency of O-ring problems. Using the same damage estimates as Figure 10.1, Figure 10.2 arrays NASA's observations of joint problems as functions of both leak-check pressures and the ambient temperatures at the launch site. Because the nozzle joints were tested at different pressures from the other joints, one flight appears in both Figure 10.2a and Figure 10.2b, and seven flights appear in both Figure 10.2b and Figure 10.2c. Figure 10.2d aggregates the problems across all three test pressures: every flight launched at an ambient temperature below 66 degrees had experienced O-ring problems. The launch that ended in disaster began at an ambient temperature around 28 degrees, 15 degrees lower than any before.

Figure 10.2a: Test pressure of 50 psi

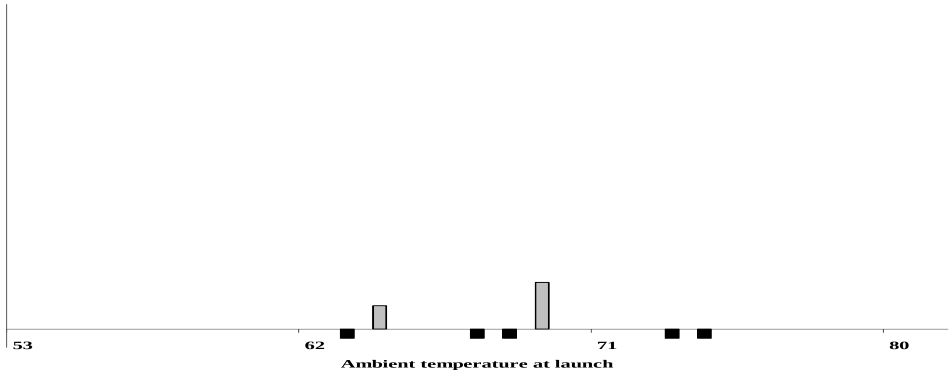

Figure 10.2b: Test pressure of 100 psi

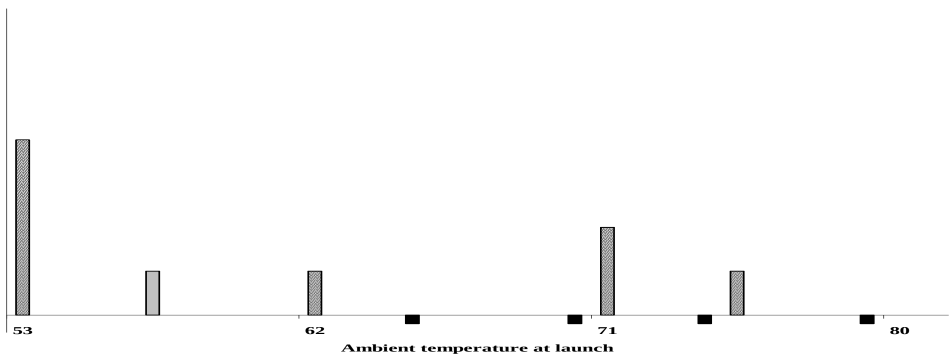

Figure 10.2c: Test pressure of 200 psi

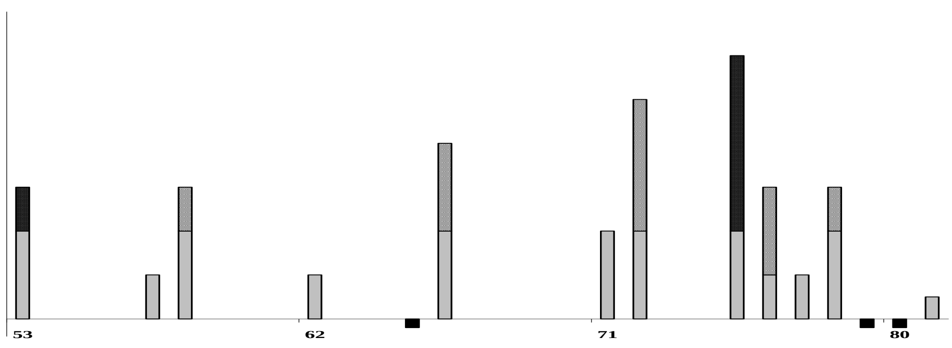

Figure 10.2d: All test pressures combined

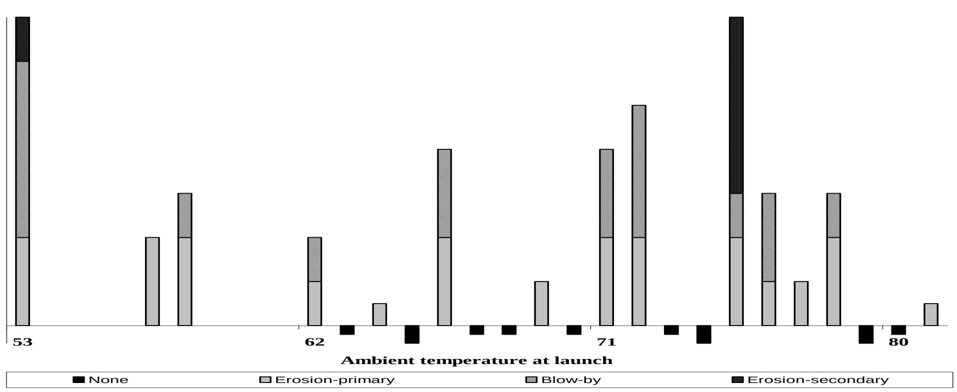

Fine- Tuning Until Failure Occurs

The most important lesson to learn from the Challenger disaster is not that some managers made the wrong decisions or that some engineers did not understand adequately how O-rings worked: the most important lesson is that fine-tuning makes failures very likely.

Fine-tuning changes always have plausible rationales, so they generate benefits most of the time. But fine-tuning is real-life experimentation in the face of uncertainty, and it often occurs in the context of very complex sociotechnical systems, so its outcomes appear partially random. For instance, because NASA did not know all of the limitations bounding shuttle operations, the doomed shuttle might not have been the flight on 28 January 1986. An SRB joint could well have triggered a disaster earlier – say, in January 1985, when five O-rings in four joints suffered damage. Or, a disaster could have happened later, had NASA postponed the launch until the cold weather abated. Such a postponement might have turned the flight into another demonstration that the O-ring problems were not so urgent – or into a disaster at 50 degrees.

Fine-tuning changes constitute experiments, but multiple, incremental experiments in uncontrolled settings produce confounded outcomes that are difficult to interpret. Thus, much of the time, people only discover the content and consequences of an unknown limitation by violating it and then analysing what happened in retrospect. As George H. Diller, spokesman at Kennedy Space Center, said: 'It is only after a problem that you can sometimes go back and see that you had a thread of data that is consistent with the event' (Sanger, 1986b: 1). For example, although the NASA and Thiokol personnel had considered many hypotheses about the O-ring problems, they had not even thought of two factors that may have been important on 28 January. Rain and freezing temperatures before launch could have produced ice in the SRBs' joints, and laboratory tests after the disaster showed that such ice could unseat secondary O-rings. Also, laboratory tests showed the fit between case segments to be as important as the O-rings' temperatures: O-rings failed to seal at temperatures below 55 degrees when the gap between case segments was only 0.004 inches, but they sealed at 25 degrees when case segments fitted with a 0.010-inch gap. The case segments that had to fit this precisely had diameters of 146 inches.

NASA's incremental changes in hardware, procedures, and operating conditions were creeping inexorably toward a conclusive demonstration of some kind. In retrospect, it now seems obvious that numerous launches had generated increasingly threatening outcomes, yet NASA's managers persisted until a launch produced an outcome too serious to process routinely. They seem to have been pursuing a course of testing to destruction.

NASA's apparent insensitivity to escalating threats has attracted criticism, and NASA could undoubtedly have made better use of the available evidence, but NASA was behaving in a commonplace way. Because fine-tuning creates sequences of experiments that are supposed to probe the limits of theoretical knowledge, people tend to continue one of these experimental sequences as long as its outcomes are not so bad: the sequence goes on until an outcome inflicts costs heavy enough to disrupt the normal course of events and to bring fine-tuning to a temporary halt.

Learning from Disasters

We may need disasters in order to halt erroneous progress. We have difficulty in distinguishing correct inferences from incorrect ones when we are making multiple, incremental experiments with incompletely understood, complex systems in uncontrolled settings; and sometimes we begin to interpret our experiments in erroneous, although plausible frameworks. Incremental experimentation also produces gradual acclimatization that dulls our sensitivities, both to phenomena and to costs and benefits. For instance, given the tendencies of NASA's and Thiokol's managers to interpret non-fatal O-ring erosion as evidence that O-ring erosion could be tolerated, it is hard to imagine how a successful flight could have produced O-ring erosion bad enough to persuade the NASA and Thiokol managers to halt launches for two or three years until the new SRB cases would be ready. Indeed, more erosion of secondary O-rings might have induced NASA to boost the leak-check pressure yet again.

One is reminded of Gregory Bateson's metaphor about a frog in hot water: A frog dropped into a pot of cold water will remain there calmly while the water is gradually heated to a boil, but a frog dropped into hot water will leap out instantaneously.

Because some disasters do inevitably happen, we should strive to make disasters less costly and more beneficial. Failures have to be costly in order for us to judge them disasters, but the Challenger disaster killed far fewer people thin other disasters that have received much less attention. Publicity and extreme visibility made the difference. We saw the Challenger disaster live on television, and we read about it and heard about it for five months, and so we valued those seven lives highly. Also, disasters often seem more costly where the people who died were not those who chose the courses of action. This poses a practical dilemma. On the one hand, our sense of justice says that the actual astronauts should decide whether to launch. On the other hand, the Challenger disaster would probably have received less public attention if the astronauts had participated in the teleconference between NASA and Thiokol on 27 January, and had themselves decided to launch at 28 degrees.

We benefit from disasters only if we learn from them. Dramatic examples can make good teachers. They grab our attention and elicit efforts to discover what caused them, although few disasters receive as much attention as Challenger. In principle, by analysing disasters, we can learn how to reduce the costs of failures, to prevent repetitions of failures, and to make failures rarer.

But learning from disasters is neither inevitable nor easy. Disasters typically leave incomplete and minimal evidence. Complex systems in uncontrolled settings can fail in a multitude of ways; unknown limitations mean that fine-tuning terminates somewhat randomly; and incremental experiments may possess numerous explanations even in retrospect. Retrospective analyses always oversimplify the connections between behaviours and outcomes, and make the actual outcomes appear highly inevitable and highly predictable (Starbuck and Milliken, 1988). Retrospection often creates an erroneous impression that errors should have been anticipated and prevented. For instance, the Presidential Commission found that The O-ring erosion history presented to Level I at NASA Headquarters in August 1985 was sufficiently detailed to require corrective action prior to the next flight', but would the Commission members have drawn this same conclusion in August 1985 on the basis of the information then at hand?

Effective learning from disasters may require looking beyond the first explanations that seem to work, and addressing remote causes as well as proximate ones. With the help of the press, the Presidential Commission did try to do that: it explored quite a few alternative hypotheses, appraised NASA's administrative processes, and pointed to potential future problems. NASA's and Thiokol's reactions are also instructive: they seem to have focused on short-run changes. NASA and Thiokol replaced many managers (Sanger, 1986g). NASA made more funds available for testing, and reviewed and 'resolved' 262 problems involving critical components, but decided not to modify the SRB cases to any substantial degree. With the addition of a third O-ring in each joint and the deletion of insulating putty, NASA's next launch will use the capture-lip cases that Thiokol had ordered in July 1985 (Sanger, 1986g).

Two years after the Challenger disaster, one astronaut observed that it had taught lessons that NASA will probably have to learn again and again.

References

Bell, T. E., and Esch, K. (1987). ‘The fatal flaw in flight 51-L.’ *IEEE Spectrum*, 24(2): 36-51.

Boffey, P. M. (1986a). ‘Space agency image: a sudden shattering.’ *New York Times*, 135(5 February): Al, A25.

Boffey, P. M. (1986b). ‘Analyst who gave shuttle warning faults “gung-ho, can-do” attitude.’ *New York Times*, 135(14 February): B4.

Boffey, P. M. (1986c). ‘Safety assessment hinged on weather and booster seats.’ *New York Times*, 135(20 February), Al, D22.

Box, G. E. P., and Draper, N. R. (1969). *Evolutionary Operation*. New York: Wiley.

Broad, W. J. (1986a). ‘Changes in rocket strained booster's seals, experts say.’ *New York Times*, 135(17 February): A14.

Broad, W. J. (1986b). ‘NASA official orders review of 900 shuttle parts.’ *New York Times*, 135(27 February): D27.

Broad, W. J. (1986c). ‘NASA had solution to key flaw in rocket when shuttle exploded.’ *New York Times*, 135(22 September): Al, B8.

Diamond, S. (1986a). ‘Study of rockets by Air Force said risks were 1 in 35.’ *New York Times*, 135(11 February): Al, A24.

Diamond, S. (1986b). ‘NASA wasted billions, Federal audits disclose.’ *New York Times*, 135(23 April): Al, A14-A15.

Diamond, S. (1986c). ‘NASA cut or delayed safety spending.’ *New York Times*, 135(24 April): Al, B4.

Dutton, J. M., and Thomas, A. (1984). ‘Treating progress functions as a managerial opportunity.’ *Academy of Management Review*, 9: 235-247.

Feynman, R. P. (1986). ‘Personal observations on reliability of shuttle.’ *Report of the Presidential Commission on the Space Shuttle Challenger Accident, Volume II, Appendix F*. Washington, DC: US Government Printing Office.

Magnluson, E. (1986). ‘Fixing NASA.’ *Time*, 127(23, 9 June): 14-18, 20, 23-25.

Marbach, W. D. et al. (1986). ‘What went wrong?’ *Newsweek*, 107(6, 10 February): 32-34.

Pear, R. (1986). ‘Senator says NASA cut 70% of staff checking quality.’ *New York Times*, 135(8 May): Al, B25.

Perrow, C. (1 984). *Normal Accidents: Living with High-Risk Technologies*. New York: Basic Books.

Presidential Commission (1986). *Report of the Presidential Commission on the Space Shuttle Challenger Accident*. Washington, DC: US Government Printing Office.

Sanger, D. E. (1986a). 'Panel questioned shuttle schedule'. *New York Times*, 135(13 February): 2.

Sanger, D. E. (1986b). ‘Shuttle data now emerging show that clues to disaster were there.’ *New York Times*, 135(16 February): 1, 32.

Sanger, D. E. (1986c). ‘Challenger report is said to omit some key safety issues for NASA.’ *New York Times*, 135(8 June): 1, 36.

Sanger, D. E. (1986d). ‘NASA pressing shuttle change amid concerns.’ *New York Times*, 135(23 September): Al, C10.

Sanger, D. E. (1986e). ‘Top NASA aides knew of shuttle flaw in ‘84.’ *New York Times*, 135(21 December): 1, 34.

Sanger, D. E. (1986g). ‘Rebuilt NASA “on way back” as an array of doubts persist.’ *New York Times*, 135(29 December): Al. B12.

Sanger, D. E. (1986f). ‘Shuttle changing in extensive ways to foster safety.’ *New York Times*, 135(28 December): 1, 22.

Schriesheim [Fulk] J., Von Glinow, M. A., and Kerr, S. (1977). ‘Professionals in bureaucracies: A structural alternative.’ In P. C. Nystrom and W. H. Starbuck. (eds.), *Prescriptive Models of Organizations*. Amsterdam: North-Holland, 55-69.

Schwartz, H. S. (1987). ‘On the psychodynamics of organizational disaster: The case of the space shuttle Challenger.’ Columbia Journal of World Business, 22(1): 59-67.

Starbuck, W. H. (1988). ‘Why organizations run into crises…and sometimes survive them.’ In K. Laudon and J. Turner (eds.), *Information Technology and Management Strategy*. Englewood Cliffs, NJ: Prentice-Hall, 11-33.

Starbuck, W. H., and Milliken, F. J. (1988). ‘Executives’ perceptual filters: What they notice and how they make sense.’ In D. Hambrick (ed.), *The Executive Effect: Concepts and Methods for Studying Top Managers*. Greenwich, CT: JAI Press, 35-65.

Weick, K. E. (1987). 'Organizational culture and high reliability'. *California Management Review*, 29: 112-127.