Blowup

WHO CAN BE BLAMED FOR A DISASTER LIKE THE *CHALLENGER* EXPLOSION? NO ONE, AND WE'D BETTER GET USED TO IT

1.

In the technological age, there is a ritual to disaster. When planes crash or chemical plants explode, each piece of physical evidence — of twisted metal or fractured concrete — becomes a kind of fetish object, painstakingly located, mapped, tagged, and analyzed, with findings submitted to boards of inquiry that then probe and interview and soberly draw conclusions. It is a ritual of reassurance, based on the principle that what we learn from one accident can help us prevent another, and a measure of its effectiveness is that Americans did not shut down the nuclear industry after Three Mile Island and do not abandon the skies after each new plane crash. But the rituals of disaster have rarely been played out so dramatically as they were in the case of the *Challenger* space shuttle, which blew up over southern Florida on January 28, 1986.

Fifty-five minutes after the explosion, when the last of the debris had fallen into the ocean, recovery ships were on the scene. They remained there for the next three months, as part of what turned into the largest maritime salvage operation in history, combing a hundred and fifty thousand square nautical miles for floating debris, while the ocean floor surrounding the crash site was inspected by submarines. In mid-April of 1986, the salvage team found several chunks of charred metal that confirmed what had previously been only suspected: the explosion was caused by a faulty seal in one of the shuttle's rocket boosters, which had allowed a stream of flame to escape and ignite an external fuel tank.

Armed with this confirmation, a special presidential investigative commission concluded the following June that the deficient seal reflected shoddy engineering and lax management at NASA and its prime contractor, Morton Thiokol. Properly chastised, NASA returned to the drawing board, to emerge thirty-two months later with a new shuttle — *Discovery* — redesigned according to the lessons learned from the disaster. During that first post-*Challenger* flight, as America watched breathlessly, the crew of the *Discovery* held a short commemorative service. "Dear friends," the mission commander, Captain Frederick H. Hauck, said, addressing the seven dead *Challenger* astronauts, "your loss has meant that we could confidently begin anew." The ritual was complete. NASA was back.

But what if the assumptions that underlie our disaster rituals aren't true? What if these public postmortems don't help us avoid future accidents? Over the past few years, a group of scholars has begun making the unsettling argument that the rituals that follow things like plane crashes or the Three Mile Island crisis are as much exercises in self-deception as they are genuine opportunities

for reassurance. For these revisionists, high-technology accidents may not have clear causes at all. They may be inherent in the complexity of the technological systems we have created.

This revisionism has now been extended to the *Challenger* disaster, with the publication of *The Challenger Launch Decision*, by the sociologist Diane Vaughan, which is the first truly definitive analysis of the events leading up to January 28, 1986. The conventional view is that the *Challenger* accident was an anomaly, that it happened because people at NASA had not done their job. But the study's conclusion is the opposite: it says that the accident happened because people at NASA had done exactly what they were supposed to do. "No fundamental decision was made at NASA to do evil," Vaughan writes. "Rather, a series of seemingly harmless decisions were made that incrementally moved the space agency toward a catastrophic outcome."

No doubt Vaughan's analysis will be hotly disputed, but even if she is only partly right, the implications of this kind of argument are enormous. We have surrounded ourselves in the modern age with things like power plants and nuclear weapons systems and airports that handle hundreds of planes an hour, on the understanding that the risks they represent are, at the very least, manageable. But if the potential for catastrophe is actually found in the normal functioning of complex systems, this assumption is false. Risks are not easily manageable, accidents are not easily preventable, and the rituals of disaster have no meaning. The first time around, the story of the *Challenger* was tragic. In its retelling, a decade later, it is merely banal.

2.

Perhaps the best way to understand the argument over the *Challenger* explosion is to start with an accident that preceded it — the near disaster at the Three Mile Island (TMI) nuclear-power plant in March of 1979. The conclusion of the president's commission that investigated the TMI accident was that it was the result of human error, particularly on the part of the plant's operators. But the truth of what happened there, the revisionists maintain, is a good deal more complicated than that, and their arguments are worth examining in detail.

The trouble at TMI started with a blockage in what is called the plant's polisher — a kind of giant water filter. Polisher problems were not unusual at TMI, or particularly serious. But in this case the blockage caused moisture to leak into the plant's air system, inadvertently tripping two valves and shutting down the flow of cold water into the plant's steam generator.

As it happens, TMI had a backup cooling system for precisely this situation. But on that particular day, for reasons that no one really knows, the valves for the backup system weren't open. They had been closed, and an indicator in the control room showing they were closed was blocked by a repair tag hanging from a switch above it. That left the reactor dependent on another backup system, a special sort of relief valve. But, as luck would have it, the relief valve wasn't working properly that day, either. It stuck open when it was supposed to close, and, to make matters even worse, a gauge in the control room which should have told the operators that the relief valve wasn't working was itself not working. By the time TMI's

engineers realized what was happening, the reactor had come dangerously close to a meltdown.

Here, in other words, was a major accident caused by five discrete events. There is no way the engineers in the control room could have known about any of them. No glaring errors or spectacularly bad decisions were made that exacerbated those events. And all the malfunctions — the blocked polisher, the shut valves, the obscured indicator, the faulty relief valve, and the broken gauge — were in themselves so trivial that individually they would have created no more than a nuisance. What caused the accident was the way minor events unexpectedly interacted to create a major problem.

This kind of disaster is what the Yale University sociologist Charles Perrow has famously called a *normal accident*. By *normal*, Perrow does not mean that it is frequent; he means that it is the kind of accident one can expect in the normal functioning of a technologically complex operation. Modern systems, Perrow argues, are made up of thousands of parts, all of which interrelate in ways that are impossible to anticipate. Given that complexity, he says, it is almost inevitable that some combinations of minor failures will eventually amount to something catastrophic. In a classic 1984 treatise on accidents, Perrow takes examples of well-known plane crashes, oil spills, chemical-plant explosions, and nuclear-weapons mishaps and shows how many of them are best understood as normal. If you saw the movie *Apollo 13*, in fact, you have seen a perfect illustration of one of the most famous of all normal accidents: the Apollo flight went awry because of the interaction of failures of the spacecraft's oxygen and hydrogen tanks, and an indicator light that diverted the astronauts' attention from the real problem.

Had this been a "real" accident — if the mission had run into trouble because of one massive or venal error — the story would have made for a much inferior movie. In real accidents, people rant and rave and hunt down the culprit. They do, in short, what people in Hollywood thrillers always do. But what made *Apollo 13* unusual was that the dominant emotion was not anger but bafflement — bafflement that so much could go wrong for so little apparent reason. There was no one to blame, no dark secret to unearth, no recourse but to re-create an entire system in place of one that had inexplicably failed. In the end, the normal accident was the more terrifying one.

3.

Was the *Challenger* explosion a normal accident? In a narrow sense, the answer is no. Unlike what happened at TMI, its explosion was caused by a single, catastrophic malfunction: the so-called O-rings that were supposed to prevent hot gases from leaking out of the rocket boosters didn't do their job. But Vaughan argues that the O-ring problem was really just a symptom. The cause of the accident was the culture of NASA, she says, and that culture led to a series of decisions about the *Challenger* that very much followed the contours of a normal accident.

The heart of the question is how NASA chose to evaluate the problems it had been having with the rocket boosters' O-rings. These are the thin rubber bands that run around the lips of each of the rocket's four segments, and each O-ring was meant to work like the rubber seal on the top of a bottle of preserves, making the fit between each part of the rocket snug and airtight. But from as far back

as 1981, on one shuttle flight after another, the O-rings had shown increasing problems. In a number of instances, the rubber seal had been dangerously eroded — a condition suggesting that hot gases had almost escaped. What's more, O-rings were strongly suspected to be less effective in cold weather, when the rubber would harden and not give as tight a seal. On the morning of January 28, 1986, the shuttle launchpad was encased in ice, and the temperature at liftoff was just above freezing. Anticipating these low temperatures, engineers at Morton Thiokol, the manufacturer of the shuttle's rockets, recommended that the launch be delayed. Morton Thiokol brass and NASA, however, overruled the recommendation, and that decision led both the president's commission and numerous critics since to accuse NASA of egregious — if not criminal — misjudgment.

Vaughan doesn't dispute that the decision was fatally flawed. But, after reviewing thousands of pages of transcripts and internal NASA documents, she can't find any evidence of people acting negligently, or nakedly sacrificing safety in the name of politics or expediency. The mistakes that NASA made, she says, were made in the normal course of operation. For example, in retrospect it may seem obvious that cold weather impaired O-ring performance. But it wasn't obvious at the time. A previous shuttle flight that had suffered worse O-ring damage had been launched in 75-degree heat. And on a series of previous occasions when NASA had proposed — but eventually scrubbed for other reasons — shuttle launches in weather as cold as 41 degrees, Morton Thiokol had not said a word about the potential threat posed by the cold, so its pre-*Challenger* objection had seemed to NASA not reasonable but arbitrary. Vaughan confirms that there was a dispute between managers and engineers on the eve of the launch but points out that in the shuttle program, disputes of this sort were commonplace. And, while the president's commission was astonished by NASA's repeated use of the phrases *acceptable risk* and *acceptable erosion* in internal discussion of the rocket-booster joints, Vaughan shows that flying with acceptable risks was a standard part of NASA culture. The lists of acceptable risks on the space shuttle, in fact, filled six volumes. "Although [O-ring] erosion itself had not been predicted, its occurrence conformed to engineering expectations about large-scale technical systems," she writes. "At NASA, problems were the norm. The word *anomaly* was part of everyday talk.... The whole shuttle system operated on the assumption that deviation could be controlled but not eliminated."

What NASA had created was a closed culture that, in her words, "normalized deviance" so that to the outside world, decisions that were obviously questionable were seen by NASA's management as prudent and reasonable. It is her depiction of this internal world that makes her book so disquieting: when she lays out the sequence of decisions that led to the launch — each decision as trivial as the string of failures that led to the near disaster at TMI — it is difficult to find any precise point where things went wrong or where things might be improved next time. "It can truly be said that the *Challenger* launch decision was a rule-based decision," she concludes. "But the cultural understandings, rules, procedures, and norms that always had worked in the past did not work this time. It was not amorally calculating managers violating rules that were responsible for the tragedy. It was conformity."

4.

There is another way to look at this problem, and that is from the standpoint of how human beings handle risk. One of the assumptions behind the modern disaster ritual is that when a risk can be identified and eliminated, a system can be made safer. The new booster joints on the shuttle, for example, are so much better than the old ones that the overall chances of a *Challenger*-style accident's ever happening again must be lower, right? This is such a straightforward idea that questioning it seems almost impossible. But that is just what another group of scholars has done, under what is called the theory of *risk homeostasis*. It should be said that within the academic community, there are huge debates over how widely the theory of risk homeostasis can and should be applied. But the basic idea, which has been laid out brilliantly by the Canadian psychologist Gerald Wilde in his book *Target Risk*, is quite simple: under certain circumstances, changes that appear to make a system or an organization safer in fact don't. Why? Because human beings have a seemingly fundamental tendency to compensate for lower risks in one area by taking greater risks in another.

Consider, for example, the results of a famous experiment conducted several years ago in Germany. Part of a fleet of taxicabs in Munich was equipped with antilock brake systems (ABS), a technological innovation that vastly improves braking, particularly on slippery surfaces. The rest of the fleet was left alone, and the two groups — which were otherwise perfectly matched — were placed under careful and secret observation for three years. You would expect the better brakes to make for safer driving.

But that is exactly the opposite of what happened. Giving some drivers ABS made no difference at all in their accident rate; in fact, it turned them into markedly inferior drivers. They drove faster. They made sharper turns. They showed poorer lane discipline. They braked harder. They were more likely to tailgate. They didn't merge as well, and they were involved in more near misses. In other words, the ABS systems were not used to reduce accidents; instead, the drivers used the additional element of safety to enable them to drive faster and more recklessly without increasing their risk of getting into an accident. As economists would say, they *consumed* the risk reduction, they didn't save it.

Risk homeostasis doesn't happen all the time. Often — as in the case of seat belts, say — compensatory behavior only partly offsets the risk reduction of a safety measure. But it happens often enough that it must be given serious consideration. Why are more pedestrians killed crossing the street at marked crosswalks than at unmarked crosswalks? Because they compensate for the "safe" environment of a marked crossing by being less vigilant about oncoming traffic. Why did the introduction of childproof lids on medicine bottles lead, according to one study, to a substantial increase in fatal child poisonings? Because adults became less careful in keeping pill bottles out of the reach of children.

Risk homeostasis also works in the opposite direction. In the late 1960s, Sweden changed over from driving on the left-hand side of the road to driving on the right, a switch that one would think would create an epidemic of accidents. But, in fact, the opposite was true. People compensated for their unfamiliarity with the new traffic

patterns by driving more carefully. During the next twelve months, traffic fatalities dropped 17 percent before returning slowly to their previous levels. As Wilde only half-facetiously argues, countries truly interested in making their streets and highways safer should think about switching over from one side of the road to the other on a regular basis.

It doesn't take much imagination to see how risk homeostasis applies to NASA and the space shuttle. In one frequently quoted phrase, Richard Feynman, the Nobel Prize–winning physicist who served on the *Challenger* commission, said that at NASA decision-making was "a kind of Russian roulette." When the O-rings began to have problems and nothing happened, the agency began to believe that "the risk is no longer so high for the next flights," Feynman said, and that "we can lower our standards a little bit because we got away with it last time." But fixing the O-rings doesn't mean that this kind of risk-taking stops. There are six whole volumes of shuttle components that are deemed by NASA to be as risky as O-rings. It is entirely possible that better O-rings just give NASA the confidence to play Russian roulette with something else.

This is a depressing conclusion, but it shouldn't come as a surprise. The truth is that our stated commitment to safety, our faithful enactment of the rituals of disaster, has always masked a certain hypocrisy. We don't really want the safest of all possible worlds. The national 55-mile-per-hour speed limit probably saved more lives than any other single government intervention of the past generation. But the fact that Congress lifted it last month with a minimum of argument proves that we would rather consume the

recent safety advances of things like seat belts and air bags than save them. The same is true of the dramatic improvements that have been made in recent years in the design of aircraft and flight-navigation systems. Presumably, these innovations could be used to bring down the airline accident rate as low as possible. But that is not what consumers want. They want air travel to be cheaper, more reliable, or more convenient, and so those safety advances have been at least partly consumed by flying and landing planes in worse weather and heavier traffic conditions.

What accidents like the *Challenger* should teach us is that we have constructed a world in which the potential for high-tech catastrophe is embedded in the fabric of day-to-day life. At some point in the future — for the most mundane of reasons, and with the very best of intentions — a NASA spacecraft will again go down in flames. We should at least admit this to ourselves now. And if we cannot — if the possibility is too much to bear — then our only option is to start thinking about getting rid of things like space shuttles altogether.

January 22, 1996