Rational Accidents

Reckoning with Catastrophic Technologies

By: John Downer

Citation:

Rational Accidents: Reckoning with Catastrophic Technologies
By: John Downer
DOI: 10.7551/mitpress/8844.001.0001
ISBN (electronic): 9780262377010
Publisher: The MIT Press
Published: 2024

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

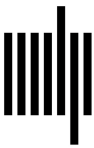

The MIT Press

7 RATIONAL ACCIDENTS: ON FINITISM'S CATASTROPHIC IMPLICATIONS

The best laid plans o' mice an' men / Gang aft agly.
—Robert Burns

7.1 ERROR PLANE

737-CABRIOLET

On April 28, 1988, Aloha Airlines Flight 243, a passenger-laden Boeing 737, left Hilo Airport on a short hop between Hawaiian islands. It climbed gently to the trip's cruising altitude of 24,000 feet. And then it tore apart.

The airplane's pilots would later report a savage lurch, followed by a tremendous "whoosh" of air that tore the cabin door from its hinges. One recalled glancing back through the space where the door used to be and seeing blue sky instead of the first-class cabin (NTSB 1989, 2). Closer inspection, had there been time, would have revealed passengers silhouetted against the emptiness, still strapped to their seats but no longer surrounded by an airplane fuselage. All of them hurtling through the air, unshielded, at hundreds of miles per hour; far above the open ocean.

Unable to communicate over the howling winds, but finding they were still in control of the airplane, the pilots set the emergency transponder and landed gingerly at nearby Kahului. Once safely on the tarmac, the airplane's condition indelibly marked its place in the annals of aeronautics folklore. An eighteen-foot, hemispherical fuselage section—thirty-five square meters of the first-class cabin—had ripped away from the airframe, severing major

FIGURE 7.1
Aloha Airlines Flight 243. *Source:* Hawaii State Archives.

structural beams and important control cables. In the affected section, only the floor and seats remained intact (figure 7.1).

The human toll was less than could reasonably have been hoped. Sixty-five of the ninety passengers and crew had been injured by winds and flying debris, eight seriously. But, by grace and safety belts, only one was lost: senior flight attendant Clarabelle Lansing, who disappeared into the void in the first seconds of the crisis, never to be seen again. The airframe itself was terminal. Never before or since has a civil jetliner survived such a colossal insult to its structural integrity. The incident—henceforth referred to simply as "Aloha"—is still widely remembered by aviation engineers, who, with the dark humor of every profession that grapples with tragedy, sometimes refer to it as the "737-cabriolet."

NESTED CAUSES

When the NTSB published its report into Aloha the following year, it identified the proximate cause to be a fateful combination of stress fractures, saltwater corrosion, and metal fatigue (NTSB 1989). Boeing built the skin of its early 737s from layered aluminum sheets, bonded together with rivets and an epoxy glue. The glue came on a "scrim" tape. Assembly workers would keep the tape refrigerated until it was in place and then cure the glue by

gradually letting it warm. This was a delicate process, however, and if the tape cooled at the wrong speed, the glue would cure incorrectly. Even under the best conditions, it would occasionally bind to oxide on the surface of the aluminum rather than the metal itself (Aubury 1992). It was relatively common for the bonding to be less than perfect, therefore, and the NTSB concluded that this had been a causal factor in the accident. The board's investigation found that imperfect bonding in Aloha's fuselage had allowed salt water to creep between its aluminum sheets. Over time, this had corroded the metal, forcing the sheets apart and creating stress around the rivets. The stress in turn had fostered fatigue cracks in the fuselage, which eventually caused it to fail, suddenly and catastrophically.[1]

The causes of accidents are invariably nested, however, and the NTSB's technical explanation of Aloha raised a series of deeper questions. Structural failures of this magnitude were not supposed to happen so suddenly, whatever the underlying cause. It was not considered possible. Experts had long understood that some 737 fuselages had imperfect bonding, and this could induce stress fatigue, but they also understood that fatigue cracks progressed gradually—slowly enough that routine inspections should have raised alarms long before cracks could cause any kind of rupture.[2] Because of this, the NTSB assigned much of the culpability for the failure to the airline's "deficient" maintenance program, which, when examined closely, proved to be more unruly in practice than it should have been in principle (NTSB 1989, §2.3).

The airline contested this finding, insisting the NTSB report gave a false impression that its maintenance practices were untypical (Cushman 1989). It is easy to be skeptical of such protestations, but the argument seems credible in hindsight. As we have seen, close examinations of technological practice routinely find it to be more untidy and ambiguous in practice than in theory, and this goes as much for maintenance as for design (e.g., Langewiesche 1998b). Few would deny that an airline's maintenance practices could, in principle, be deficient to a point where they became legitimately negligent, but there is little evidence to suggest that Aloha's met this benchmark. Even the NTSB agreed that the airplane was compliant with all relevant FAA mandates, and, tellingly perhaps, the accident led to very few sanctions on the airline.

The official interpretation of the accident as a "maintenance failure" is complicated further by the fact that the 737's fuselage was supposed to be

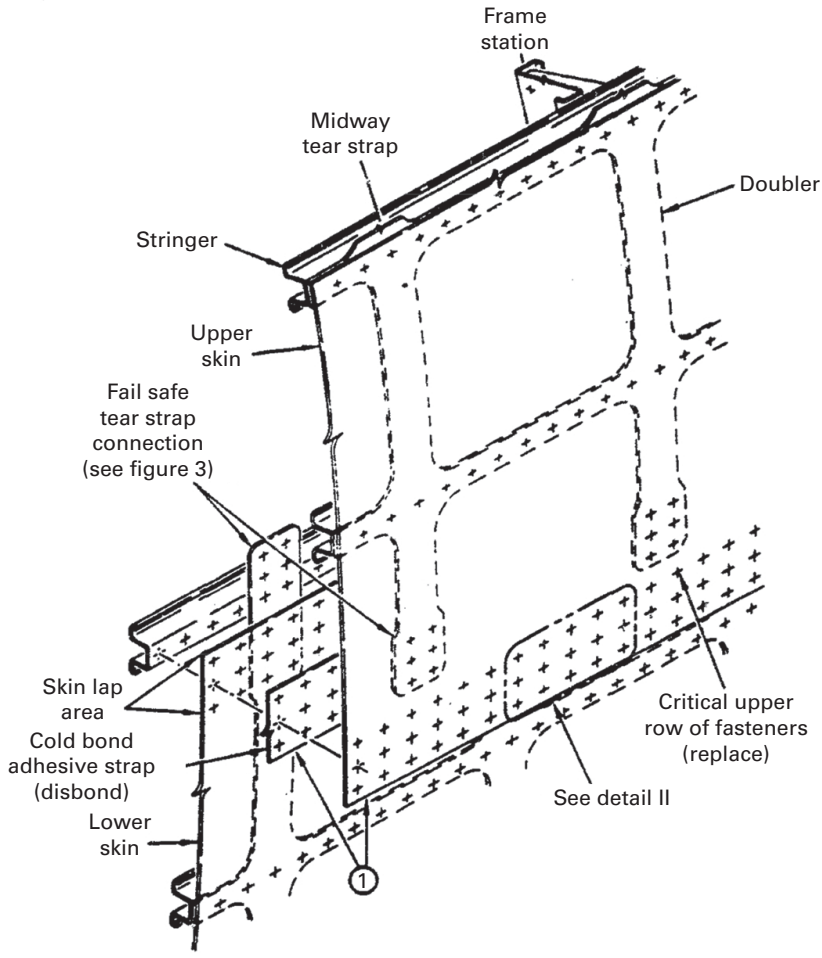

FIGURE 7.2
Boeing 737 fuselage. *Source:* NTSB (1989).

safe even if maintenance inspections fell spectacularly short. This belief was premised on the "fail-safe" design of its metal skin, which was built—per certification requirements—to be tolerant of crack-induced ruptures and breeches. Boeing had achieved this by dividing the skin into small—10-by-20-inch (25.4-by-50.8-cm)—rectangular panels (known as "frame-bays"), each bounded by "tear straps" designed to constrain breaches by channeling and redirecting cracks, much as perforations in paper control tearing (figure 7.2). In theory, therefore, any rupture in the fuselage should have caused

it to "flap" open around a single panel, releasing the internal cabin pressure in a way that was limited, controlled, and—importantly—not threatening to the airframe's basic integrity (NTSB 1989, 34). For extra security, the design allowed for cracks of up to 40 inches that encompassed two panels simultaneously.

Fully understanding Aloha, therefore, requires an explanation that reaches beyond the airframe's maintenance and accounts for the failure of its fail-safe design. And herein lies the accident's most generalizable insights. The story of how this theoretically impossible failure was possible illustrates why the finitist limitations of tests and models translate directly into limits on technological safety. In doing so, it suggests a new perspective on why catastrophic technologies fail despite our best-laid plans, and it lays an important foundation for understanding why jetliners, specifically, fail so infrequently.

Before looking further at Aloha, however, it is worth contextualizing the accident by looking briefly at the broader academic literature around technological disaster.

7.2 THE LIMITS OF CONTROL

FAILURES OF FORESIGHT

If a complex technology like a jetliner fails, especially if it fails in a manner thought to be impossible, then it is easy to interpret the failure as some kind of error: be it in design, manufacture, oversight, maintenance, or operation. A core tenet of what Jasanoff (2005) calls our "civic epistemology" of technological risk—our shared expectations about the kind of problem that it poses, and the broad manner in which it should be governed—is that all technological accidents are (in principle at least) avoidable. This understanding reflects the positivism implicit in mainstream conceptions of engineering knowledge. We believe that there is something ontologically distinctive about *failing* technologies—something that sets them apart from *functioning* technologies, which proper procedure should (or, at minimum, could) always identify in advance if experts were appropriately organized, skilled, and incentivized.

Professional engineering discourse reifies this belief by routinely portraying technological disasters as events that were *allowed* to happen, albeit unintentionally.[3] Media discourse follows the same template: framing catastrophic accidents in terms that emphasize their fundamental preventability

(Downer 2014; Hilgartner 2007). And, as discussed earlier, the academic literature on disaster usually reflects the same implicit conviction, especially in when discussing accidents born of design weaknesses. A lot of this scholarship touches on the difficulties of knowing machines. As early as 1976, for example, Turner (1976, 379) was highlighting, as a practical matter, the limited data and theory with which engineers routinely operate, as well as the implications of this for safety (see also Turner and Pidgeon 1997; Weick 1998; Vaughan 1996). But even in its most nuanced forms, the balance of this literature very rarely "bites the bullet of uncertainty" as Pinch (1991, 155) puts it, by articulating and exploring the *necessary* and *unavoidable* limits of knowledge and the significance of those limits for understanding disaster.

As we saw, however, there is one significant exception to the general assumption that technological failures are theoretically (if not always realistically) avoidable: Perrow's (1986, 1999) Normal Accident Theory (NAT). Perrow explicitly rejects the idea that perfect organizational practices (should they be possible) could yield perfectly reliable machines. And for the purposes of contextualizing the argument that follows, it is worth pausing to examine his argument in more detail.

NORMAL ACCIDENTS

NAT, developed by Yale sociologist Charles Perrow, is most fully articulated in his book *Normal Accidents* (Perrow 1999 [1984]). The text is wide-ranging and accommodates multiple interpretations (see, e.g., Le Coze 2015), but a useful way to understand the theory's core thesis is as an argument about probability and the taming of chance. By this reading, Perrow contends that some accidents in systems with certain properties—his eponymous normal accidents—are fundamentally unforeseeable and unavoidable because they stem from coincidences that are too improbable to identify in advance.

At the heart of this argument are two deceptively simple insights. The first is that accidents can result from fatal one-in-a-billion confluences (which no analysis could ever anticipate) of otherwise unremarkable anomalies (of a kind that no design could ever avoid entirely),[4] which compound each other to create a catastrophe. Perrow argues, for instance, that the 1979 accident at Three Mile Island exemplifies this phenomenon. By his account, the accident began when leaking moisture from a blocked filter tripped valves controlling the plant's cooling system. Redundant backup valves should have intervened, but they were inexplicably and erroneously locked closed. The closed

valves should have been clear from an indicator light, but it was obscured by a tag hanging from a switch above. A tertiary line of technological defense, a relief valve, should then have opened, but it also malfunctioned (which went unnoticed because a different indicator light simultaneously failed, erroneously indicating that the relief valve was functioning). None of these failures or anomalies was particularly noteworthy in itself, he argues, but together they created a catastrophic situation that controllers understandably struggled to comprehend (Perrow 1999).

Perrow's second insight is that these kinds of fateful one-in-a-billion coincidences are statistically probable in systems where there are billions of opportunities for them to occur: specifically, those that are composed of many closely interconnected and highly interdependent elements. (Systems that are "interactively complex" and "tightly coupled," in his terminology.) Because where a system features many interactions between different elements, the number of unanticipated anomalies (and combinations of anomalies) that may occur in it is greatly magnified. And where its safe functioning depends on those elements all working together, the difficulty of managing such unexpected anomalies is greatly increased (Perrow 1999).

So it is, he argues, that certain types of system unavoidably harbor the potential for accidents that escape even the best oversight and control mechanisms: ghosts that lurk in the interstices of engineering risk calculations. Because where potentially dangerous systems have millions of interacting elements that allow billions of unexpected events and interactions, then seemingly impossible billion-to-one coincidences—too remote to register in any engineering analysis—are only to be expected (i.e., they are "normal").[5]

Perrow's argument is exceptional in the context of disaster research because it speaks to the inherent limits of engineering knowledge. It clearly and unambiguously articulates a case for why accidents can occur *without meaningful error*, such that their causes resist all organizational explanations and remedies.[6] NAT, we might say, delineates a conceptually important category of accidents that experts could never even aspire to organize away. It does not claim that all accidents have this character. Indeed, Perrow (1999, 70–71) argues that normal accidents are rare, and that most accidents are potentially avoidable "component failure accidents," characterized by predictable relationships between elements, or by a single, catastrophic technological fault rather than a combination of faults across a system. ("Normal," in this context, is intended to connote "expected" rather than "common.")

Perrow does not believe that Challenger, Bhopal, or Chernobyl were normal accidents, for example. He interprets them in a traditionally positivist fashion: as a product of errors, some of which were culpable, and all of which might, in theory at least, have been organized out of existence.

RATIONAL ACCIDENTS

Understanding the basic contours of NAT is useful for framing Aloha because, construed in Perrow's terms, Aloha—the failure of a single critical element (the fuselage) from a known cause (fatigue)—was a quintessential component failure accident. Even according to NAT, therefore, it should have been avoidable. If engineers were adequately rigorous with their tests, thorough with their inspections, and assiduous with their measurements, NAT suggests, the accident would not have happened.

Understanding the accident through a finitist lens, however, suggests a different view. The preceding chapters of this book have explored and illustrated the argument that expert understandings of technical systems necessarily hinge on qualitative interpretations and judgments (e.g., about the representativeness of tests and models). These interpretations and judgments can be better or worse, considered or rash, skillful or inept; what they cannot be, however, is "knowably perfect." There are infinite ways in which a test or model might be unrepresentative, and it is impossible to examine them all. If even the most rigorous engineering analyses are imperfect, however, then it cannot be true that flawed technologies are always distinguishable from flawless technologies. This is axiomatic. If experts cannot know the accuracy of their tests and models with certainty, then they cannot use those tests and models to know (and thus predict and/or control) a system's failure behavior with certainty. It is logical to conclude from this that accidents could result from conditions caused by (and unrecognized because of) erroneous, but nevertheless *rationally held*, engineering beliefs. And that, like normal accidents, these accidents would be fundamentally unavoidable.

Elsewhere, I have referred to such accidents as "epistemic accidents" (Downer 2011b; 2020). "Epistemic" can be an onerous word, however, so in this volume I will refer to them instead as "*rational accidents*." (The change in label is not intended to connote any meaningful change in definition.)

Rational accidents can be defined as accidents that occur because a technological understanding proves to be unsound, even though there were rational reasons to hold that understanding before (although not after)

the event. Unlike normal accidents, which arise at the system level, from the indeterminacies of interactions between elements, rational accidents can arise from uncertainties embedded in a single element. Like normal accidents, however, they are inherently unpreventable and unpredictable. Investigations into their causes would find unproven assumptions underpinning the flawed system's design, but so too would investigations into fully functional systems, so it is wrong to imagine that the former are ontologically distinct from the latter. (To paraphrase Bloor [1976], our understandings of each should be "symmetrical.")

With this idea in mind, let us now return to Aloha, which I will argue is an exemplary rational accident. It occurred because expert understandings of the airplane's fuselage were incomplete, and this incompleteness is more appropriately attributed to the epistemological limits of proof than to any insufficiency of effort, rigor, or logic.

7.3 ALOHA REVISITED

MULTIPLE SITE DAMAGE

To understand how Aloha might be construed as a rational accident, it helps to begin by understanding more about Multiple Site Damage (MSD): the almost imperceptible fatigue-cracking that ultimately led the fuselage to fail so spectacularly.

MSD was a known problem at the time of the accident, but neither Boeing nor the FAA considered it a major safety issue. Their lack of concern in this regard stemmed, in large part, from a belief—conventional across the industry—that no MSD crack could grow from a microscopic level to 40 inches (the level to which the fail-safe tear-panels had been tested) in the period between maintenance inspections. As it was understood at the time, MSD should always have manifested as cracks that were detectable (or, at absolute minimum, controllable by the tear panels), long before they became dangerous.

As Aloha demonstrated, however, this understanding of MSD was dangerously flawed. Specifically, it failed to recognize that in certain areas of the fuselage, and under certain conditions (specifically, where there was significant disbonding between fuselage sheets, combined with a corrosive saltwater environment and a significant passage of time), MSD had a tendency to develop along a horizontal plane between a line of rivets (figure 7.3). And,

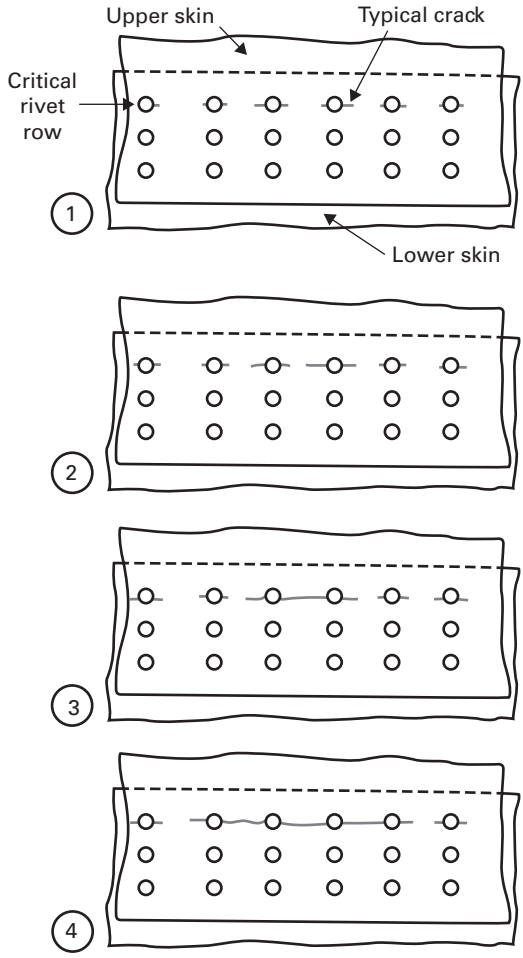

FIGURE 7.3
MSD cracks. *Source:* NTSB (1989).

further, that such a string of almost indiscernible damage could abruptly coalesce into one huge crack, longer than 40 inches, that would nullify the fail-safe tear straps and, in Aloha's case, almost bring down an aircraft (NTSB 1989: §2.3). "The Aloha accident stunned the industry," as the FAA's "Lessons learned" website puts it, "by demonstrating the [unforeseen] effects of undetected multiple site damage."

This widespread misunderstanding of MSD raises its own questions, however, and itself deserves exploration. After all, it seems reasonable to imagine that engineers should have understood how fuselages fatigue.

DARK ARTS

To understand engineers' relationship to metal fatigue at the time of Aloha, it helps to know a little history. Organized research into metal fatigue began in the mid-nineteenth century, when it first became apparent that seemingly good machine parts were failing as they aged. By 1849, engineers had coined the term "metal fatigue" and were actively working to better understand the problem; and by the 1870s, they had carefully documented the failure behavior of various alloys, even though the causes of this behavior remained opaque (Garrison 2005). Over time, the formal study of metal fatigue grew steadiliy from there, eventually coming to be known as "fracture mechanics." Its practitioners examine the properties (principally the tensile strength and elasticity)[7] of metals, as well as their relationship to different kinds of stress.

Fracture mechanics has long been a central concern of aeronautical engineering, having shot to prominence at the dawn of the jet age when fatigue felled two of the world's first jetliners—de Havilland Comets—within a four-month period of 1954 (Faith 1996, 158–165). "Although much was known about metal fatigue," Geoffrey de Havilland would lament in his autobiography, "not enough was known about it by anyone, anywhere" (quoted in Marks 2009, 20). Having borne witness to the dangers that could hide in uncertainties about fatigue, the aviation industry worked hard to explore the phenomenon. By the time Aloha occurred in 1988, therefore, its understanding of the subject was grounded in extensive experiments and decades of experience with aluminum airplanes. As the accident reaffirmed, however, this research and experience had not perfected that understanding. Even today—over sixty years after the Comet accidents, and over thirty since Aloha—aircraft fatigue management remains an inexact science, one that, in 2005, after fatigue felled a Royal Air Force transport aircraft outside Baghdad, *Air Safety Week* (2005) described as "a dark black art . . . akin to necromancy."

The intransigent uncertainties of metal fatigue are rooted in complexities that resonate with those of bird-strike tests. In the laboratory, where minimally defective materials in straightforward forms are subjected to known stresses, fracture mechanics is a complex but broadly manageable science. Decades of theory and experimentation have led to models that work tolerably well in predicting where, when, and how a metal form will fail, even if the most accurate models must grapple with quantum mechanics. In real, operational aircraft, by contrast, the problem of relevance asserts itself and the representativeness of the laboratory models becomes questionable. In this more practical realm, elaborate geometric forms, replete

with irregularities and imperfections, experience variable and uncertain stresses, the vicissitudes of time, and the insults of human carelessness. All this introduces uncertainty. In these circumstances, even modest and seemingly innocuous design choices can create unanticipated stress points, as can minor deformations like scratch marks or defects in the metal itself. Such variables are so challenging to accurately model in advance that, by some accounts, even the most sophisticated fatigue predictions about real fuselages essentially amount to informed guesses (Feeler 1991, 1–2; Gordon 2018 [1991]: locs. 853–854, 1236–1239).

Compounding this complexity, the properties that make fatigue difficult to predict also make it difficult to monitor. Unlike most metals, which do not fatigue until they reach a specific stress threshold, the aluminum alloys used to make airframes (in the past at least; the industry has recently embraced advanced composites, as we will see) are thought to fatigue at any stress level. This precrack fatigue was once believed to develop at the microscopic level of the metal's crystalline structure, but now (post-Aloha) is understood to accumulate, in a nonlinear fashion, at the atomic level, where it is functionally invisible to maintenance engineers. (As of 2008, "nondestructive evaluation" techniques could detect cracks only as small as 0.04 inch [Maksel 2008]). The result, in the words of one engineer, being that "until cracking begins, it is for all practical purposes impossible to tell whether it will begin in 20 years, or tomorrow" (Garrison 2005).

DUHEM-QUINE

The complexities of fatigue might explain why engineers failed to predict Aloha's vulnerability to MSD and the inadequacies of its tear panels, but it is less obvious why those failed predictions survived the 737's extensive compliance tests. Before the airplane entered service, Boeing, under FAA supervision, extensively tested the airframe's fatigue resilience and decompression behavior; including the efficacy of the tear panels. And years later, as some 737s approached their original design life, the company even acquired and retested an old airframe. Why, then, did the dangers go undetected?

The answer, in essence, is that Boeing's tests of its fatigue predictions were framed by the same imperfect theory that framed the predictions themselves. Engineers use tests to examine the validity of their beliefs about a system's functioning, but, as we have seen, tests are themselves inescapably theory-laden (in that they embody complex ideas about the relevance of different variables, and so on). To test one theory, therefore, is to always stand on

another: there can be no view from nowhere. Epistemologists refer to this dilemma—where theories cannot be tested independently from other theories—as the "Duhem-Quine Problem," and Aloha elegantly illustrates its real-world implications.

The 737's fatigue tests were premised in part on a belief that the key determinant of fatigue in an aircraft fuselage was not its age or its hours of service, but its number of "pressurization cycles" (every takeoff and landing usually constituting one full cycle). This belief was so deeply embedded in the industry's understanding of fuselage fatigue—so much so that the "fatigue life" of an airframe was expressed as a number of "cycles" to failure (NTSB 2006b, 87). It was logical, therefore, that 737's fatigue tests should be framed in the same terms. Testers simulated service experience by pressurizing and depressurizing (known as "cycling") a fuselage half-section 150,000 times (representing twice the airplane's design life). This produced no major MSD cracks and fulfilled all FAA certification requirements (NTSB 1989: §1.17.2).

In Aloha's specific case, however, the airframe's pressurization cycles was not the only factor relevant to understanding its fatigue behavior. The airplane was certainly "highly cycled" (because of its short routes), but a range of other conditions contributed to the MSD that brought it down. One was its operating environment: the warm, saltwater air around Hawaii. Another was the manufacturing flaws in its structure: the imperfect bonding in its fuselage. A third was its sheer chronological age: manufactured in 1969, it was one of the longest-serving 737s in operation. These three factors—environment, manufacture, and age—were all crucial to its failure. The disbonding created gaps that allowed Hawaii's salt water to intrude. And, over a long period of time, that water corroded the fuselage in a way that stressed its rivets and nurtured the cracks that caused it to fail. At the same time, however, all these factors set Aloha's airframe apart from the new, properly bonded fuselage that Boeing repeatedly pressurized in a dry laboratory over a highly compressed time frame. It also distinguished the airframe from the older airframe Boeing used for follow-up tests, which again had a properly bonded fuselage that had not been flying short routes in salty Hawaii. As a result, the tests failed to fully represent Aloha's real-world circumstances. By isolating pressurization as the limiting factor in fatigue, the engineers had unwittingly excluded a range of variables that would have been highly significant to predicting the airplane's failure behavior.

Note the circularity of this error, where theories about the causes of fatigue shaped the tests intended to interrogate those theories, thereby rendering

those tests blind to the kinds of theoretical shortcomings that they were intended to reveal. The logic underpinning the 737's structural integrity enjoyed an internal consistency, we might say, but one that held itself aloft by its own bootstraps. This is a property of knowledge that is often highlighted by finitists (e.g., Collins 1985; Kuhn 1996 [1962]; Quine 1975). Bucciarelli (1994, 92), for instance, speaks of "The incestuous character" of what he calls the "model-making process," wherein "the model [is] designed to verify the field data; the data, in turn, providing a reference for the model."

It is difficult to fault the testers for this circularity. Boeing's (and the FAA's) engineers had run into an intractable dilemma that Collins (1985, 84) calls the "experimenter's regress." They had no way of determining the accuracy of their findings without knowing the representativeness of their tests, and no way of determining the representativeness of their tests without knowing the accuracy of their findings.

Once established, moreover, the imperfect understanding of fatigue born of this regress propagated throughout the airframe's wider design and test regimen. Importantly (and somewhat ironically), for instance, it undermined the fail-safe tear panels that were intended to serve as an ultimate hedge against errors and misunderstandings. This is to say that it led engineers to design the panels around the premise that escaped engine blades, not fatigue cracks, would cause the largest possible fuselage ruptures. And it then hid the consequences of this misconception by shaping the way that the panels were tested. Believing that escaped engine blades posed the most risk, engineers tested the panels by "guillotining" a fully pressurized fuselage section with two 15-inch blades that represented engine fragments. This created punctures smaller than 40 inches, which traversed only one "tear-strap" and led the skin to flap open exactly as predicted (NTSB 1989: §1.17.2). Unfortunately, however, it also left engineers blind to the panels' inability to cope with the kind of MSD rupture that almost felled Aloha. As the closing line of the Comet accident report had put it years earlier: "extrapolation is the fertile parent of error" (Allen 2004, 19).

7.4 IMPLICATIONS

KNOWLEDGE AND DESIGN

The story of Aloha's hidden vulnerability speaks to the unusual symmetry, in safety-critical engineering, between "knowing" and "designing." In

highlighting the catastrophic potential of an unexpected property of metal fatigue—which only became dangerous when an airframe with an uncommon manufacturing defect operated for years in a specific environment—it illustrates how even the most marginal misunderstandings can be significant to a system's failure behavior. Seen in this light, it becomes easier to understand why "technological reliability" begins to converge with "epistemological truth" as demands on the former begin to rise. And, as a result, it becomes easier to appreciate the finitist case against catastrophic technologies. The more reliability required of a complex system, the more prohibitive the problem posed by rational accidents.

It is worth noting that engineers in this domain widely recognize that they always have gaps in their knowledge, and that these gaps can be a source of danger to airplanes. "The wit of a man cannot anticipate, hence prevent, everything that could go wrong with an airplane in flight," as Newhouse (1982, 83) puts it: a sentiment that is invoked often by aviation practitioners (e.g., NAS 1980, 41) and academic observers (e.g., Mowery and Rosenberg 1981, 348) alike. Turner and Pidgeon (1997; 1978, 71ff), for instance, speak of "notional normality" in engineering, wherein, "perceptions of risk are sustained by sufficiently accurate individual and organizational beliefs about the world . . . up to the point that those beliefs are challenged by a major disaster or crisis."

Despite this recognition, however, the nature and implications of the relationship between epistemology and disaster are routinely underconsidered. It is worth exploring these implications and their significance, therefore, and a useful way of framing such an examination is to compare briefly the properties of what I have called rational accidents with those of Perrow's normal accidents.

Three properties, in particular, are worth highlighting here. I will call them "avoidability," "vulnerability," and "learning."

AVOIDABILITY As we have seen, both rational and normal accidents are fundamentally unavoidable and (in the specific rather than general sense) unforeseeable. The cause of this is different in each case, however. Normal accidents are unavoidable because engineers cannot wholly predict the multiplicity of possible interactions in a system, whereas rational accidents are unavoidable because the myriad knowledge claims that engineers draw on when designing and evaluating systems are inherently fallible.

One notable implication of this difference, as we also have seen, is that rational accidents and normal accidents define the scope of unavoidability differently, with the former suggesting that its ambit extends further than NAT would allow. Aloha was not a normal accident—there was no unforeseeable, billion-to-one confluence of otherwise foreseeable events, just the failure of a single element (the fuselage)—but Aloha nevertheless has a good claim to being unavoidable on epistemological grounds. The design of its fuselage embodied complex theories about metal fatigue; those theories were built on (and then reified by) tests and models; and those tests and models were themselves inescapably theory laden. This circularity created an irreducible measure of uncertainty: no process, however rigorous, could have guaranteed that every judgment underpinning the airframe's design and assessment was correct. Its design might have embodied flawed beliefs, therefore, but those beliefs were neither lazy nor illogical. It would be unreasonable, in these circumstances, to construe the accident as a failure of foresight.

(This is not to say, of course, that fatal engineering errors and design flaws can *never* be unreasonable or culpable. No epistemologist would deny that there are *better* and *worse* ways of establishing the properties of artifacts, even if there is no *perfect* way. Thus, there are undoubtedly "responsible" and "irresponsible" engineering practices, even if both have socially negotiated definitions and the former can never guarantee safety. The point is that *some* accidents will always be unavoidable on epistemic grounds, even if scholars might debate the extent to which any specific accident should qualify. It goes without saying, therefore, that social scientists should continue to explore the social, psychological, and organizational foundations of error. The social practices underpinning technological safety are undeniably consequential and demonstrably improvable, even if they are not perfectible.)

VULNERABILITY As with normal accidents, there are good reasons to imagine that rational accidents are more probable in some systems than in others. Because these accidents have different causal mechanisms, however, the probability of each is controlled by different variables.

As outlined previously, Perrow's key indicators of a system's vulnerability to normal accidents are "tight coupling" and "high interactive complexity." Both these properties are also likely to make systems more susceptible to rational accidents. Coupling, because the more tightly interdependent the elements in a system are, the more likely it is that any epistemologically

driven failures will instigate catastrophic accidents; complexity, because every extra element in a system, and every extra relationship between elements, represents a new set of potentially fallible knowledge claims. Beyond this, however, rational accidents are likely to vary with properties that normal accident theorists might otherwise ignore.

Consider, for instance, the "variegation" of a system, which is to say the level of differentiation between its elements. Systems can be highly complex without being very variegated. Early computers, for example, consisted of thousands of identical vacuum tubes. And from an NAT perspective, such systems might be no less "complex" than those with high variegation (consisting of differently designed elements, made of dissimilar materials, performing many distinct functions). Significantly, however, highly variegated systems represent a much larger number of knowledge claims, and it follows from this that they would be significantly more vulnerable to rational accidents.

Consider also the "innovativeness" of a system, loosely defined as the extent to which it stretches the boundaries of established theory and prior experience. An airframe panel made from a new material, for instance, is neither complex nor tightly coupled, and as such would not be flagged by normal accident theorists as a source of vulnerability. (Perrow [1999, 128] actually cites the aviation industry's "use of exotic new materials" as a factor that directly contributes to the safety of modern aircraft.) From a rational accidents perspective, however, this innovation looks inherently risky. It deprives the experts charged with designing and assessing the panel of decades of research and service experience on which they might otherwise have drawn when anticipating its failure behavior, creating additional epistemic uncertainty.

(As with Perrow's complexity and coupling, the innovativeness of a system might be difficult to measure exactly or quantify objectively, as might the degree of variegation between systems, but this does not mean that these terms have no analytical value. Beauty is proverbially subjective, but this hardly negates its existence, its usefulness as an explanatory category, or its tangible consequences.)

LEARNING The aforementioned relationship between experience and knowledge speaks to a third distinction between rational and normal accidents: their different relationship to hindsight and, through it, to learning.

Take, for instance, the extent to which each kind of accident could be said to have "warning signals" that experts missed. Even with the benefit

of hindsight, it makes little sense to speak about normal accidents having warning signals. The kinds of minor anomalies that combine to cause them are *expected* to occur in complex systems; it is only in their unexpected confluence that they become meaningful. So it is that normal accidents seemingly come from nowhere. (They almost have to, because if terrible coincidences built incrementally and identifiably over time, they would no longer be coincidences.) Rational accidents, by contrast, are significantly different in this respect. From the vantage of hindsight, for instance, Aloha looks replete with potential warning signals. There are key aspects of the accident—such as the planar accretion of MSD—that experts today would be expected to identify as a precursor to disaster. Such signals were not obvious at the time, because experts did not know where to look or for what they should be searching, but they are discernible in retrospect and so would constitute warning signals today. (We might even say that those warning signals *existed* only in retrospect.) The fatigue that led to Aloha would probably have been caught by modern maintenance practices, for example, but only because of insights gleaned from the accident itself. "New knowledge can turn normality into hazards overnight," as Beck (1999, 58) puts it.

It follows from this that experts might *learn* from rational accidents, by leveraging hindsight in a way that is not possible with normal accidents.

In their purest from, normal accidents yield few useful lessons. There are two broad reasons for this. The first is that the accidents from the same proximate causes are *highly unlikely to reoccur*. Where there are billions of possible billion-to-one coincidences that can instigate a catastrophic accident, then it is logical to anticipate repeated accidents, but not to anticipate the exact same accident twice. Addressing the specific factors that contributed to one normal accident, therefore, is unlikely to protect against the next. The second reason why normal accidents yield few lessons is that they *do not challenge common engineering understandings and theories about the world*. This is because the factors that combine to produce them are not surprising in themselves. For example, a stuck valve usually reveals little about valves in general and does nothing to challenge the knowledge underlying their design. The surprising aspect of normal accidents lies in the coincidence of different failures compounding each other. So it is that normal accidents only teach one lesson, and it is always the same: *experts can never design out every tragic coincidence*. This insight has important policy ramifications, to

be sure, but it is of very limited value to experts looking to improve systems over time.

Rational accidents are very different in both these respects. Unlike normal accidents, the events that instigate them are likely to reoccur *if left uncorrected*. If Aloha had disappeared over an ocean, cause undiscovered, then—as with the Comets in 1954—other aircraft would likely have failed in the same way, and for the same reason. And, again unlike normal accidents, rational accidents *challenge common engineering understandings and theories about the world*. Aloha revealed fundamental misunderstandings about metal fatigue, for instance. It also revealed meaningful ways in which fuselage tests were unrepresentative of the phenomena they sought to reproduce. These two properties of rational accidents—the fact that they reoccur and the fact that they challenge conventional understandings—mean that they can yield useful design insights. Post-Aloha, for example, experts could revisit their understanding of metal fatigue and its relationship to the 737 fuselage, and this in turn meant that they could ensure that the same accident would not happen again. Engineers understood jetliners better after the accident, in other words, and jetliners are safer as a result.

So it is that a finitist understanding of failure suggests that we understand all new technologies as real-life experiments with dangerously uncertain, but ultimately instructive outcomes. This dynamic, wherein engineers might leverage hindsight for design insights, is important. Indeed, it is key to resolving the aviation paradox.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Sans and Stone Serif by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Downer, John (John R.), author.
Title: Rational accidents : reckoning with catastrophic technologies / John Downer.
Description: Cambridge, Massachusetts : The MIT Press, [2023] | Series: Inside technology | Includes bibliographical references and index.
Identifiers: LCCN 2023002845 (print) | LCCN 2023002846 (ebook) | ISBN 9780262546997 (paperback) | ISBN 9780262377027 (epub) |
ISBN 9780262377010 (pdf)
Subjects: LCSH: Reliability (Engineering) | Aircraft accidents—Prevention. | Risk assessment. | Industrial accidents—Prevention.
Classification: LCC TA169 .D69 2023 (print) | LCC TA169 (ebook) | DDC 620/.00452—dc23/eng/20230202
LC record available at https://lccn.loc.gov/2023002845
LC ebook record available at https://lccn.loc.gov/2023002846