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

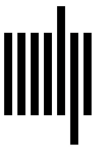

The MIT Press

6 THE SUM OF ALL PARTS: MODELING RELIABILITY WITH REDUNDANCY

It is well to moor your bark with two anchors.
—Publilius Syrus

6.1 SYSTEM-LEVEL SAFETY

"IMPOSSIBLE" FAILURES

On June 24, 1982, 37,000 feet over the Indian Ocean, the pilots of British Airways Flight 009 noticed the power begin to drop on each of the Boeing 747's four engines. Their concerns were soon heightened when pungent smoke began creeping into the flight deck. And they intensified even further when passengers reported seeing a numinous blue glow around the engine housings. Just minutes after the first indication of trouble, one engine shut down completely. Soon a second failed; then a third; and then, upsettingly, the fourth. With no understanding of why, the pilots found themselves gliding a powerless jetliner, 80 miles from land and 180 miles from a viable runway.

Despite a memorably laconic address from the captain,[1] the glowing, silent engines and creeping smoke understandably caused some consternation among the 248 passengers. Several scribbled heartfelt notes to loved ones as their airplane gradually earned a mention in the *Guinness Book of Records* for "Longest unpowered flight by a non-purpose-built aircraft."[2] Just as cabin crew were preparing them for an unpromising midocean ditching, however, the power abruptly returned. Passengers in life jackets applauded the euphony

of restarting engines, and—after some last-minute drama with an occluded windscreen[3]—the pilots deposited them safely in Jakarta (Tootell 1985).

The underlying cause of this aeronautical scare—now widely remembered as the "Jakarta incident"—was not a mystery for long. Investigators found the engines were full of ash particles and quickly deduced that the airplane had inadvertently flown through the plume of a volcano erupting nearby, the ash from which had clogged the engines and sandblasted the cockpit windows. The last-minute reprieve had come when the plane lost enough altitude to leave the plume, whereupon airflow had cleared some of the ash.

Although easily explained, the Jakarta incident unnerved the aviation community, which was (and largely remains) committed to the belief that no commercial jetliner, especially a four-engine 747, could lose propulsion from all its engines simultaneously. The 747 can fly relatively safely on a single engine, so four give it a comfortable margin of security. Under normal circumstances, experts consider even a single engine failure to be highly unlikely. So until British Airways inadvertently demonstrated otherwise, they considered the likelihood of all four failing simultaneously to be functionally zero: as near to impossible as makes no practical difference. The industry was so confident of this that pilots were trained to treat indications of quadruple engine failure as evidence that there was something wrong with the indicators, not the engines. Nobody had considered the possibility of volcanic ash.

Flight 009's near-miss was not the first time that the industry found that it had placed too much confidence in redundancy, however, and it would not be the last. Galison (2000), for example, discusses a similarly unanticipated failure in a redundant aviation system, this time with an unhappier ending. United Airlines Flight 232, a DC-10 en route to Chicago in 1989, lost its hydraulic systems after its tail-mounted engine exploded. McDonnell Douglas had explicitly designed the DC-10's hydraulics to resist such trauma. Each channel was triple-redundant, with redundant pumps, redundant reservoirs of hydraulic fluid, and redundant (and differently designed) power sources. And this redundancy had led the aviation community to deem a total hydraulic failure in the DC-10 to be functionally impossible, just as it had done with engine failure. This was "so straightforward and readily obvious," experts had argued, that "any knowledgeable, experienced person would unequivocally conclude that the failure mode would not occur" (Haynes 1991). Nevertheless, the engine explosion managed to sever all three systems, rendering the airplane all but uncontrollable. (Remarkably, there was a credible attempt

at a crash landing. The pilots, together with a flight instructor who happened to be on board, somehow managed to steer the plane to a runway by manipulating the power to different engines. This feat—which nobody was subsequently able to replicate in the simulator—saved the lives of 185 of the 296 passengers and crew.)

These incidents, and others like them, are significant because the reliability of all civil aircraft, and all assessments of that reliability, depend heavily on notions of redundancy. It would be fair to say, in fact, that redundancy is a sine qua non of almost all catastrophic-technological design and assessment.

DEFINITIONS AND AMBIGUITIES

Much like reliability itself, "redundancy" is a seemingly intuitive concept that is nevertheless surprisingly, and perhaps revealingly, difficult to pin to the page. As with reliability, the term has a specific meaning in engineering that is both similar to and sometimes at odds with its common English usage. In both contexts, it implies some manner of repeating, but where it usually has negative connotations in its common usage—as something superfluous and excessive—engineers more often equate it with surety. In the latter context, it is understood to be a property of a system's architecture. Broadly, if we imagine a "system" as being composed of many "elements," then an element can be described as being "redundant" insofar as it is part of a system that contains backups to do that element's work if it fails; and a system can be described as "having redundancy" insofar as it consists of redundant elements.

So far, so simple, but note that the definitions given here are inherently contextual. They can operate at multiple levels, such that, for instance, redundancy can take the form of multiple capacitors on a circuit board; multiple circuit boards on a sensor; multiple sensors in a system; and so on, all the way to the duplication of entire sociotechnical networks. (In the US missile programs of the 1950s, for example, the Air Force entrusted mission success to the doctrine of "overkill" and the redundancy of the whole missile and launch apparatus [Swenson, Grimwood, and Alexander 1998, 182]). A further ambiguity in such frameworks is that some redundant elements, such as engines on an airplane, work concurrently (but are capable of carrying the load by themselves if required), whereas others, such as backup generators, usually lie idle, only to awake if and when they are needed.

Jetliners, like most catastrophic technologies, invoke different forms of redundancy on many levels simultaneously. Indeed, the type-certification

process requires them to do so. Redundancy has been a staple design requirement for airframes since the dawn of aviation regulation. (The very first airworthiness standards, for instance, required that biplanes have duplicate flight control cables [FAA 2002c, 23–24]). Its significance to aeronautical engineering and oversight is difficult to overstate. A core tenet of the FAA's safety-philosophy is that no single component-level failure should be catastrophic. Achieving this goal, and formally demonstrating that achievement, would both be unthinkable without redundancy.[4]

RELIABILITY SQUARED

Perrow (1984, 196) equates redundancy with a Kuhnian paradigm, and while invocations of Thomas Kuhn are often hollow, the analogy is uncommonly useful in this context. It highlights the way that redundancy acts as a conceptual lens through which experts "know" artifacts: shaping designs by shaping how those designs are understood.[5]

Redundancy's value to engineers of catastrophic technologies lies in two closely related, but nevertheless distinct, reliability-related functions that it performs. The first is that it allows experts to *design* complex systems that are far more reliable than their constitutive elements. (The realization that complex systems could be made more dependable than their components came surprisingly late, for reasons that I will discuss, but it was a breakthrough nonetheless.) The second relates to what the FAA calls "checkability," its slightly unfortunate term for a system's capacity to be assessed and verified. It is that redundancy allows experts to quantitatively *demonstrate* (as well as achieve) higher levels of reliability than would be possible with tests alone.

To understand why this second function is important, it is useful to understand that even if (contra the previous chapter) we imagined tests to be perfectly representative, they alone would still be unsuitable for demonstrating the extreme reliability required of catastrophic technologies. The precise amount of test data required to statistically establish a defined level of reliability is contested at the margins (e.g., Ahmed and Chateauneuf 2011; Littlewood and Wright 1997), but at the levels required of catastrophic technologies it is unambiguously impractical. Rushby (1993), for instance, calculates that a single test system would need to run failure free for over 114,000 years to demonstrate a mean-time-to-failure in excess of a billion hours. Such times could be reduced by running multiple tests in parallel, but the cost of reducing them enough this way would be gigantic. Even by

certification's own logic, untroubled by the problem of relevance, therefore, jetliner systems would be insufficiently "checkable" if tests were the only means of establishing reliability.

Redundancy offers a solution to this checkability problem. It can serve this second function, in essence, because the way it performs the first (i.e., the manner in which it increases the reliability of a system) can be represented mathematically. Put very simply: where two redundant elements operate in parallel to form a system, then the probability of that system failing can be expressed as the probability of one element failing multiplied by that of the other failing. If tests show that one element has a 0.0001 probability of failing over a given period, for instance, then the probability of two identical elements failing over the same period would be that number squared—that is, $(0.0001)^2$ or 0.00000001. Note that this calculation has demonstrated a *ten-thousandfold* increase in reliability, and has done so without any recourse to more lab tests (Shinners 1967, 56; Littlewood, Popov, and Strigini 2002, 781; FAA 1982, 3).

It is this capacity for transcending tests that makes redundancy so valuable to catastrophic technology assessment processes like type certification. Although not the only systems-level tool for *increasing* a system's reliability—some elements can be made fail-safe instead, for instance—the manner in which redundancy can be modeled in calculations makes it almost indispensable to the processes by which experts formally *demonstrate* ultrahigh levels of reliability (Downer 2011a; 2009b).

For all their importance and utility, however, redundancy calculations, much like bird-strike tests, are far messier in practice than they appear on paper. This is because models of systems, much like tests of systems, are representations, the fidelity of which raises unanswerable questions. And while redundancy's relationship with reliability might seem straightforward when represented mathematically, its practical implementation involves reckoning with complex uncertainties and judgments that are difficult to capture with quantitative certainty.

6.2 MESSY PRACTICE

IMPERFECT REPRESENTATIONS

Redundancy is undoubtedly a powerful engineering tool. It has allowed engineers to make almost all complex systems—certainly all catastrophic

technologies—more reliable than their constituent elements. Despite its benefits, however, evidence shows that redundant systems are never quite as reliable as formal models would imply (Littlewood 1996). The causes of this deficit are readily apparent when the discourse around redundancy is examined closely, as any such examination reveals routine disagreements about the appropriateness of various assumptions and variables. Such disagreements take many forms, but the following sections sketch out some of their more prominent themes. By outlining issues pertaining to "mediation," "common-cause failures," and "failure propagation," in turn, I will unpack some of redundancy's practical complexities, and show how those complexities give rise to uncertainties that are not easily reduced to definitive numbers.

Let us turn first to redundancy's relationship with mediation.

MEDIATION John von Neumann, the Hungarian émigré polymath, is generally credited as being the first person to propose redundancy as a way of increasing the reliability of complex, tightly integrated technological systems, in a 1956 treatise: "Probabilistic Logics and Synthesis of Reliable Organisms from Unreliable Components." The mid-1950s might seem incongruously late for such a straightforward-seeming insight to gain traction, but the true innovation lay not in the idea of redundancy per se, so much as in envisaging a system that could invoke it without human intervention. Grappling with the headache of aggregating thousands of unfaithful vacuum tubes into a working computer, von Neumann intuited that redundancy could help if it were combined with a managerial system to immediately identify and mediate between failures: recognizing malfunctioning tubes and switching to their backups without interrupting the wider system. It was this management system—and the problem of automatically distinguishing failed from functional tubes—that posed the real engineering challenge.

To resolve the mediation problem, von Neumann proposed a "voting" system. "The basic idea . . . is very simple," he wrote, "Instead of running the incoming data into a single machine, the same information is simultaneously fed into a number of identical machines, and the result that comes out of a majority of these machines is assumed to be true. . . . this technique can be used to control error" (1956, 44). Not every redundant element in a system requires mediation in this fashion—redundant load paths in airframe fuselages don't, for instance—but most do. And by devising a way to

automate mediation, von Neumann made redundancy a staple tool of high-reliability engineering.

His solution was no panacea, however, as mediating systems pose their own costs, challenges, and uncertainties.

One often-underappreciated cost of mediation is that it makes systems more complex. Even before mediation, redundancy increases the number of elements in a system, making unexpected interactions more likely and failure behavior more difficult to verify. (So it is that scholars routinely correlate simplicity with reliability in technological systems [e.g., Perrow 1984, 270; Kaldor 1981, 111; Arthur 2009]). This can be a problem even with unmediated redundancy, but it is greatly exacerbated by adding systems to mediate between elements. So much so, in fact, that experts sometimes argue that mediated redundancy can increase the complexity of a system to a point where it becomes a primary source of *unreliability* (Rushby 1993; Hopkins 1999).

Take, for example, the engines on a jetliner. It may sound simple to affix more engines to an airplane, but this simplicity quickly dissolves if we consider the elaborate, highly integrated management systems needed to govern them. Failures in redundant jetliner engines instigate an orchestra of rapid and highly automated management interventions. A computer first determines which engine has failed, for instance, and then it indicates the failure to the pilot, cuts the fuel, douses any flames, adjusts the rudder, compensates for the missing thrust, and much more (Rozell 1996). Many of these actions are safety critical—especially during takeoff, when reaction margins are tight and even a momentary loss of power can be fatal—and all require a suite of sensors and computers with high-authority connections to disparate elements of the aircraft. All this inevitably makes the system's failure behavior more difficult to understand and control. It also creates entirely new avenues of failure; a faulty computer that erroneously shut down engines during takeoff could easily fell a jetliner.

This example also illustrates a second dilemma of mediation: that management systems are themselves critical systems that require extreme reliability. Air-accident investigations frequently implicate mediating systems as contributory factors to catastrophic failures. When USAir Flight 427 crashed while approaching Pittsburgh in 1994, for instance, investigators found that the Boeing 737's computer had failed to adequately compensate for the roll generated by a rudder failure (NTSB 2006b, 7). Mediating systems have even

been known to instigate aviation disasters. After Indonesia AirAsia Flight 8501 crashed into the Java Sea in 2014, for instance, investigators identified a failure in the Airbus A320's failure-monitoring systems as a primary cause. They found that a malfunctioning sensor in the rudder had led the crew to reset the computer, which disabled the autopilot, which in turn led them to stall the airplane (BBC 2015; KNKT 2015).[6]

The criticality of mediating systems raises especially challenging questions for the designers and regulators of catastrophic technologies because such systems cannot themselves be redundant. This is because the mediating elements themselves would then need mediating, and so on, in an infinite regress. "The daunting truth," to quote a 1993 report to the FAA, "is that some of the core [mediating] mechanisms in fault-tolerant systems are single points of failure: they just have to work correctly" (Rushby 1993). Engineers usually negotiate this problem by designing mediating systems to be simpler, and thus (hopefully) more reliable, than the systems they mediate. But while this arguably makes the problem tractable from a design perspective, it does little for experts hoping to use abstract models of redundancy to demonstrate ultrahigh levels of reliability.

Experts manage this latter problem—of integrating mediation into models of ultrahigh reliability—in different ways. One is with semantics: certification standards often just don't define mediating systems as "safety critical," so the reliability of those systems is not held to the same standard of proof. Another is by giving mediating functions to human beings. Type certification is an analysis of the airplane, largely unsullied by the vagaries of its operation. (The FAA does assess pilots, but via a separate process; certification essentially treats them as a solved problem.)[7] If pilots are made responsible for mediating between elements, therefore, assessors can effectively "lose" any uncertainties arising from mediation by exploiting the interstices between different regulatory regimes.

The role that humans often play in this regard make them uniquely important to understanding redundancy mediation, and it is worth pausing to examine this role in slightly more depth. From a design perspective, the fact that people sometimes can respond creatively to unanticipated errors and interactions can make them uniquely versatile as mediators between redundant elements (Hollnagel, Woods, and Leveson 2006, 4; Rasmussen 1983). Recall, for instance, the crew of Flight 232, outlined previously, who used the engine throttles to steer a DC-10 after its hydraulics failed. Again,

however, delegating mediating tasks to humans is far from a perfect design solution, and it does little to resolve the problems that mediation poses to reliability calculations (Bainbridge 1983).

There are two basic reasons for this. The first is that relying on human beings rarely negates the need for safety-critical mediating elements. This is because humans can rarely mediate effectively without relying on an elaborate series of indicators and sensors working reliably. (As in the example of Flight 8501, for instance, it is not uncommon for indicator failures to be implicated in accidents.) The second, more fundamental issue is that people are proverbially imperfect. They get ill; they get confused; they run a gamut of emotions from stress to boredom; and for all these reasons, they sometimes make mistakes, disobey rules,[8] and, very occasionally, commit premeditated acts of sabotage (Reason 1990; Perrow 1999, 144–146; Lauber 1989). It would be fair to say, in fact, that human beings are a significant source of failure in jetliners, as they are in all sociotechnical systems (Dumas 1999; Reason 1990; Bainbridge 1983). The NTSB has estimated that 43 percent of fatal accidents involving commercial jetliners are initiated by pilot error (Lewis 1990, 196).

Airframers work hard to mitigate the risks posed by human mediation (many of which might reasonably be attributed to ergonomic factors, such as misleading cockpit displays [Perrow 1983; Rasmussen 1990]). They hone the designs of human-machine interfaces with an eye to making them intuitive and error tolerant ("foolproof," or sometimes "drool-proof" in less reverent industries). They also limit the scope of pilots' actions with elaborate flight protections, which, when enabled, sometimes allow the computer to overrule commands it deems dangerous. They even invoke further redundancy, designing jetliners to be crewed by two pilots.

Such measures are far from perfect, however, and sometimes create epiphenomenal dilemmas (Bainbridge 1983). Redundant personnel can induce overconfidence, for example, or what Sagan (2004, 939–941) calls "social shirking": a mutual belief that the other person will "take up any slack." (This can threaten communications between pilots, for instance [e.g., Wiener et al. 1993]). Automated flight protections can have a similar effect, as in 1988, when the captain of an Airbus A320 wrongly assumed its flight computer would prevent a stall during a low-speed airshow flyover (Macarthur and Tesch 1999). (This was the A320's first passenger flight. Thousands of spectators watched as the airplane, full of raffle winners and journalists, crashed into a forest at the end of airfield.) There are also concerns that such

protections undermine the basic competencies of pilots, who are decreasingly required to exercise even rudimentary flight skills outside the simulator. (The 2009 loss of Air France Flight 447, for instance, is often partly attributed to a basic pilot error, which many observers attribute to the universality of automatic flight protections [Langewiesche 2014]).

COMMON CAUSE A second dilemma of redundancy revolves around the independence of redundant elements. Recall Flight 009, with its dramatic power loss outside Jakarta, and Flight 1549, with its remarkable landing in the Hudson River. Both were precipitated by multiple, simultaneous engine failures caused by a common external pressure: volcanic ash in the former case, and geese in the latter. Engineers refer to such incidents—where redundant elements fail at the same time for the same reason—as "common mode failures," and the FAA (1982, appx. 1) has described them as a "persistent problem" for certification assessments. The accidents that they cause are significant in this context because they highlight a prevalent critique of redundancy calculations: that they often erroneously assume that redundant elements will behave independently with respect to their failure behavior (e.g., Popov et al. 2003).

The elegant mathematics outlined previously—where the reliability of one redundant element can be multiplied by that of another to arrive at the combined reliability of the system—only works insofar as those reliabilities are assumed to be perfectly independent, which is to say that the chances of one element failing are not linked, in any way, to the chances of the other failing. As the 155 passengers who began their flight at LaGuardia and ended it in the river can attest, however, independence is rarely perfect.

There are many principled reasons to doubt the independence of redundant elements in a system. Almost by definition, for instance, redundant elements share a common function. Many are also colocated, share a common design, and draw on a shared resource, such as fuel or electricity. Such commonalities inevitably create common vulnerabilities. They mean that a cloud of ash or a flock of birds is likely to stress all the engines on a jetliner simultaneously, for instance, as might a fuel leak or even an imperfect maintenance operation (Eckhard and Lee 1985; Hughes 1987; Littlewood 1996; Littlewood and Miller 1989).[9] Where redundant elements operate at the same

time (as engines do), moreover, then they are likely to fatigue in similar ways, further aligning their failure behavior. (Although it is also entirely possible for external pressures to stress even idle elements as they wait in reserve [Acohido 1996].)[10]

Just as engineers have practical design techniques for managing the complications of mediation, so they have ways of maximizing the independence of redundant elements. Most rely on what is known as "design diversity": the practice of designing redundant elements differently while keeping their functions the same. The idea is to create elements with dissimilar weaknesses, and thus more independent failure behavior. Manufacturers pursue diversity in varying ways. Some leave it to evolve spontaneously, contracting isolated different engineering teams (often from separate contractors) to design redundant elements and trusting that a lack of central authority will create sufficient variation (Littlewood and Strigini 1993, 9; Bishop 1995). The A320's redundant flight computers are supplied by different vendors for this reason, for example, as is the software they run (Beatson 1989). Others actively try to force diversity by explicitly requiring different teams to use divergent approaches, solutions, and testing regimens. Software manufacturers, for instance, sometimes require teams to program in different languages (Popov et al. 2003, 346). An elaboration of this approach, known as "functional diversity," requires that engineers design elements to use different inputs, in the hope that conditions that challenge one will not challenge another (Littlewood, Popov, and Strigini 1999, 2; Beatson 1989). So it is, for instance, that jetliners simultaneously use pressure, radar, and global positioning systems (GPS) to determine altitude. (The same principle is at work when pilots on the same flight are required to eat different meals to protect against simultaneous food poisoning.)

All these strategies can be useful for improving the reliability of a system, but, again, there are many principled reasons to believe they can never be perfect. The idea that different groups, when left to their own devices, will design the same system differently is bolstered by studies highlighting the contingency of technological designs (e.g., Bijker, Hughes, and Pinch 1989), for example, but the same research suggests that their ideas will likely converge when designers come from similar professional cultures and have problems specified in similar ways. (It has been found, for example, that programmers asked to independently design different versions of the same

software tend to make similar mistakes and produce code with coincident failure behaviors [Knight and Leveson 1986]).

"Forced diversity" might be stronger, but it too has theoretical shortcomings. It demands a well-defined notion of "dissimilarity," which in turn poses questions about what constitutes a "meaningful difference." But "difference," much like "representativeness," always has a bounded and socially negotiated meaning (Collins 1985, 65). Like truth, beauty, and contact lenses, it inevitably rests in the eye of the beholder and must be restricted to a finite number of variables before engineers can enforce it on designs. (Should diverse elements be forced to use wires of different gauges, for example? Should they both use wires? Should they both use electricity?)

"Functional diversity" shares the same underlying problem and has its own distinctive limitations. Designing systems to operate on different inputs might help in some circumstances, but the exogenous pressures on a system often come from sources that are unrelated to its inputs, such as a fire, collision or lightning strike. It is also true that seemingly different and separate inputs are often interrelated; extremes of temperature, for instance, correlating (at least loosely) with extremes in pressure. And also that, even when building functionally dissimilar systems, engineers are usually still working from a similar definition of what constitutes a "normal" or "routine" environment (Littlewood and Strigini 1993, 10).

For all these reasons, it would be misleading to imagine that design diversity, in whatever form, produces mathematically perfect failure independence. (As with mediation, it has even been the *source* of failures. The first Space Shuttle launch, for instance, was initially scrubbed after its primary and backup software showed a 40-millisecond time skew. The problem was caused by different programming priorities chosen by their—purposefully different—programmers: IBM and Rockwell [White 2016, 350–351]). To the extent that diversity does provide additional independence, moreover, it would be misleading to imagine that experts could precisely quantify the increase. To do so would require testing elements together as a single system to measure the rate of coincident failures: an endeavor that would quickly run into the practical limits that experts use redundancy to transcend. The precise degree of independence between elements is inherently uncertain, therefore, and techniques like diversity do little for those who would use redundancy to quantify (as opposed to maximize) the reliability of a critical system.

PROPAGATION A third complication of redundancy arises from the fact that failures in energetic and interdependent systems rarely keep to themselves. An important dimension of technical malfunction that redundancy models struggle to capture is that failures have a tendency to propagate across elements in what engineers call "cascades" (NAS 1980, 41; Zdzislaw, Szczepanski, and Balazinski 2007).

Consider, for example, a 2010 incident onboard Qantas Flight 32: a two-decked Airbus A380 carrying over 450 passengers. The airplane's pilots were taken aback when their cockpit lit up with fifty-four simultaneous failure indications. (The cockpit designers had never envisaged so many coincident malfunctions, which filled its displays with more text than they were able to show [Lowy 2010]). After a dramatic overspeed landing in Singapore—involving an emergency gear drop, four blown tires, and a three-hour dousing by emergency crews—it was discovered that the many-pronged crisis had been instigated by a single misaligned counterbore in one of the airplane's oil pipes. The misalignment had led to a fatigue fracture; the fractured oil pipe had led to an engine fire; the fire had caused one of the engine turbines to shatter, releasing three disk fragments through the engine housing; and the disk fragments had wreaked havoc on the airframe. They ripped through a wing, damaging a structural element and slicing electrical cables; they punctured two fuel tanks, causing leaks and more fires; they disabled the airplane's antilock brakes, one of its hydraulic systems, and the controls for one of the other engines; and they damaged the landing flaps. These many insults, in turn, instigated a range of other downstream failures—for instance, the leaking fuel created a mass imbalance that pilots were unable to redress because of electrical damage to the pumps—which, together with the primary failures, created the unmanageable kaleidoscope of cockpit warnings (ATSB 2013; Lowy 2010).

On one level, the fact that failures tend to instigate more failures is a straightforward observation, but it has complex ramifications for reliability calculations. This is because it means that a complete understanding of the reliability offered by redundancy needs to account, not only for the independence of redundant elements in a system, but also for the independence of these elements from other, functionally unrelated elements in other systems. Accounting for this kind of failure propagation can be extremely challenging, however, because its mechanisms are often difficult to predict, measure, and control.

There are many reasons why failure propagation across systems is difficult to model. One is that it is difficult to always predict *how* elements will fail. "You can't always be sure your toilet paper is going to tear along the perforated line," as one aviation engineer put it.[11] In 2005, for instance, Malaysia Airlines Flight 124—a Boeing 777 taking off from Perth—spontaneously pitched upward, activating stall warnings and startling the crew. On its return to the airport, investigators identified the cause to be a faulty accelerometer. The accelerometer had had a redundant backup in case it failed, but its designers had failed to predict the *way* that it would fail. They had assumed that a failure would always result in an output of zero volts, but in this instance it had produced a high-voltage output that confused the flight computer (ATSB 2007).

Beyond this, however, even very predictable failure mechanisms can introduce difficult uncertainties if they propagate in ways that are difficult to contain. Engineers are well aware that explosions are plausible failure modes in some elements, for example, and that explosions can jeopardize other, functionally unrelated elements, but they are hard pressed to mitigate such effects. (Several commercial aircraft have been lost after their fuel tanks exploded. TWA Flight 800, which blew up in 1996 for undetermined reasons shortly after leaving JFK Airport, is one. Pan Am Flight 214, which blew up after being struck by lightning in 1963, is another.) And, while explosions are exemplary forms of propagation, there are many other ways in which the failure of one element can have unexpected consequences throughout an airframe. A failed element might threaten others simply by virtue of having a destabilizing mass, for instance, or by drawing excessively on a common resource. (This happened in August 2001 when a faulty crossfeed valve near the right engine of an Airbus A330-200 bled fuel until none remained to power the aircraft.)[12]

Accounting for propagation can have counterintuitive implications for redundancy's relationship to reliability. Consider, for example, the safety afforded to an airplane by redundant engines. "Two engines are better than one," writes Perrow (1984, 128), and "four better than two." This seems simple enough, and perhaps Perrow is correct, but his truism is more questionable than it appears. If we recognize that when engines fail, they can do so in ways that propagate and catastrophically damage the airframe—as happened to Qantas Flight 32, for instance, or, more fatally, to United Flight 232 in

July 1989—then it is less clear that four engines are safer than two. Indeed, it is possible that four engines could be significantly *less* safe than two.

To understand how two engines could be safer than four, imagine a hypothetical airplane that can function safely with one engine and is sold in both a two- and a four-engine configuration. For the sake of illustration, let us assume these engines have perfectly independent failure behavior. Now let us imagine that the chances of any given engine failing during a flight are one in ten. (It is an extraordinarily unreliable design!) To make matters worse, however, one out of every five engines that fail also explodes, destroying the airplane. (Of course, the airplane is also lost if all the engines fail in the same flight.) In this scenario, the two-engine aircraft would have a higher chance of an "all-engine" failure, but a lower chance of experiencing a catastrophic explosion. The math works out such that the combined risk of any catastrophic event during flight (an all-engine failure or an explosion) is higher with a four-engine configuration than with two. This is to say that two engines would be safer than four.

So it is that propagation might theoretically create circumstances where added redundancy *detracts* from a system's reliability. This is more than simply academic. The engine example given here uses unrealistic probabilities but it captures a real concern. Boeing has made essentially the same argument about two- versus four-engine aircraft, albeit with less straightforward numbers. Advocating for fewer restrictions on two-engine airplanes, the company claimed that its 777 was safer with two engines because of the reduced risk of one failing catastrophically (Sagan 2004, 938).

As with the other complications of redundancy outlined here, engineers have pragmatic design techniques for mitigating propagation. These often involve "isolating" different elements by physically segregating and/or shielding them from each other. Critical microelectronics are sometimes encased in ceramic, for example, while engines are separated on the wing with cowlings designed to contain shrapnel from broken fan blades.

Much like diversity, however, and for essentially the same reasons, isolation is a necessarily interpretive, ambiguous, and unquantifiable property of systems. As with the measures that engineers use to mitigate redundancy's other complications, therefore, it is a useful but imperfect practice that gives rise to complex disagreements. During the certification of the Boeing 747-400, for example, the FAA and its European counterpart differently interpreted an

identically worded stipulation governing the isolation of redundant wiring (FAA 2002b, 4–6). The European regulator interpreted the word "segregation" more conservatively than the Americans, forcing Boeing to redesign the wiring late in the certification process. Because of this, two different designs of the 747-400 coexisted (GAO 1992, 16).

Engineers charged with certifying aircraft assess isolation with analytical tools such as Failure Modes and Effect Analysis, which maps the relationships between elements. Applying these tools is as much art as science, however, it being impossible to foresee every possible interaction (and to calculate the limits of one's forsight), and they are more useful for design than for assessment (NTSB 2006b; Hollnagel 2006). As with the ambiguities arising from mediation and common cause failures, therefore, the type-certification process ultimately navigates the difficulties of quantifying isolation by fiat. It mandates certain deterministic requirements—such as engine separation—and then unrealistically treats those requirements as offering mathematically perfect independence (FAA 2002b; NAS 1980; Leveson et al. 2009).[13]

6.3 OBDURATE UNCERTAINTY

AN IMPERFECT TOOL

As in chapter 5's discussion of bird-strike tests, none of the complications outlined here should be read as a suggestion that redundancy is not a useful engineering tool for reliability. (Although they do suggest that it isn't always the *optimal* design solution for reliability. As Popov et al. [2003, 346] put it: "Redundancy [is only] a reasonable use of project resources if it delivers a substantial . . . increase in reliability, greater than would be delivered by spending the same amount on other ways of improving reliability.") For our purposes, however, the important insight to be gleaned from this discussion is that redundancy is an *imperfect* tool—one that requires subjective judgments and complex interpretations to implement.

These judgments and interpretations are important. The uncertainties they imply are all potential sources of error, where every error holds the potential for catastrophe. The FAA has conceded as much, noting in 2002 that redundancy "has costs, complexities, and the inherent risk of unforeseen failure conditions associated with it" (FAA 2002c, 23–24). Its type-certification processes are framed around quantitative reliability targets, however, and models of redundancy are all but indispensable for demonstrating compliance

with those targets. So experts elide redundancy's uncertainties in their reliability calculations, thereby making certification possible but doing little to negate the dangers that those uncertainties pose (which, even if marginal, ought to be intolerable in contexts requiring ultrahigh reliabilities.)

As with bird-strike tests, therefore, a close look at the nuances of redundancy modeling does more to illustrate the aviation paradox than to resolve it. Type certification is supposed to be the process through which experts establish the ultrahigh reliability of jetliners, but, epistemologically, its imperfect tools and practices are simply not up to the job. Examined closely, the practical limitations of tests and models are impossible to reconcile with the service record of civil aviation. The ultrahigh reliability of modern jetliners, as well as the fact that regulators accurately predict that reliability prior to them entering service, imply a depth of understanding that is incommensurable with the practices from which that understanding is ostensibly derived. Logically, at least, jetliners should be failing for reasons that slip though even the most rigorous analyses.

And some do exactly that.

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