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

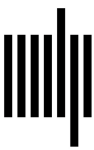

The MIT Press

NOTES

INTRODUCTION

1. The term reliability is being used narrowly here. As, more conventionally, it might also encompass non-catastrophic failures.

2. It is worth noting that Vaughn (2021; 1996) is unusual in having authored signal texts from both sides of Sagan's divide. There is nothing inherently problematic about this, of course, but it does suggest that we should be cautious of categorizing scholars themselves, as commonly occurs.

3. Insofar as a principled disagreement does exist in the scholarship on this question, it probably lies in whether technologies like reactors and atomic weapons can ever be "safe enough": a question that is as much about morality and politics as it is about safety practices.

4. I tried to make this point in an earlier publication (Downer 2011b), but I framed it more as a fundamental difference of principle, which was an error on my part.

CHAPTER 1

1. Take, for instance, estimates of the accident's financial costs. In 2013, these were usually pegged at around $100 billion. By 2017, however, the Japan Institute for Economic Research was putting them in the region of $449–628 billion, simply from decommissioning, decontamination, and compensation (Burnie 2017). This figure does not include indirect costs arising from the accident's massive disruption to the country's energy sector, or wider losses to tourism, agriculture, and industry.

2. In 1959, the RAND Corporation published a top-secret report outlining how the US should manage the publicity fallout from an accidental nuclear warhead detonation. It suggested that authorities aggressively delay and draw out the release of

damaging information because the media's interest would inevitably wane over time (Schlosser 2013, 195). Whether by design or by chance, information about Fukushima has followed a similar trajectory, with RAND's predicted outcome.

3. The "banana equivalent dose" is an informal measurement of radiation exposure premised on the fact that bananas contain radioactive potassium. It is favored by some industry communicators, especially those eager to downplay radiological hazards, but it obfuscates more than it illuminates because it ignores the fact that humans excrete excess potassium (WashingtonsBlog 2013; EPA 2009, 16).

4. The language around this promise can be misleading. Strictly speaking, expert bodies (at least in the US) rarely assert that failures in acute catastrophic technologies are impossible. Such claims are too difficult to defend. Instead, those bodies tend to assert that failures in such systems are too improbable to take seriously, despite their potentially dire consequences. For almost all practical purposes, however, the claims are equivalent. The levels of reliability demonstrably achieved in jetliners, when applied to systems built in much smaller numbers, imply a vanishingly small risk of catastrophic failure. And states have long been comfortable ignoring risks they deem possible but implausible. Few would deny that extra-terrestrial invasions or apocalyptic asteroid strikes are hypothetically possible, but neither contingency is afforded much serious (or even superficial) deliberation.

5. The relative nature of aviation safety is exemplified by a 1920 book titled Aerial Transport, which boasted that "[t]he argument that flying is inherently dangerous cannot in fact be substantiated," citing data showing that 40,000 miles were being flown between fatal accidents (quoted in Chaplin 2011: 78). The same rate today would imply well over a million fatal accidents every year.

6. As of this writing, there are about sixty-five times as many jetliners as there are reactors, and their operational downtimes are more equivalent than is probably intuitive. So, as an extremely rough rule of thumb, if each failed as frequently as the other, then we might expect about sixty-five plane crashes for every reactor meltdown.

7. Atomic weapons are unique in the sense that there are potentially extreme dangers that arise from not building them, as well as from them failing, which complicates any societal risk calculus involving their failure behavior. This tension is evident in Sagan (1993).

8. The NRC and the US Department of Defense, for instance, both explicitly recommend the use of use of qualitative "human performance reliability analysis" in their assessment guidance for critical systems (e.g., DoD 2000, 1999; NRC 1983).

CHAPTER 2

1. Efforts to quantify failure performance predate this shift. In the early 1800s, for instance, railway development inspired extensive studies on the life spans of roller-bearings (Villemeur 1991). In the 1920s, meanwhile, Bell Laboratories, struggling with

the mercurial performance of vacuum tubes, formed a "quality" department that pioneered various quantitative reliability metrics (Zimmel 2004; Shewhart 1931). It was not until World War II, however, that such techniques started to become standard engineering tools. And it was primarily during the Cold War that mathematicians and engineers began to develop the conceptual tools, metrics, and categories that normalized reliability as a quantitatively expressible and manipulable entity. It was also at this time that reliability engineering emerged as a distinct specialty within the discipline and profession (Jones-Imhotep 2002; Regulinski 1984; Kececioglu and Tian 1984).

2. "Service data," in this context, might loosely be understood as referring to the combination of the number of systems in operation multiplied by the duration of their operation. For example, 100 identical systems, each operating for 100 hours, will accrue 10,000 hours of service data, and so will 1,000 identical systems, each operating for 10 hours. As we will see, however, this equivalence is not perfect.

3. Such *ceteris paribus* assumptions are not always straightforward, however. When the US infantry started fighting in Vietnam, for instance, much of the reliability data for its rifles was rendered moot by unexpected operational practices and the humid operating environment (Fallows 1985).

4. Lynch and Cole (2002), for instance, make this argument in relation to the reliability of DNA testing. They point out that the oft-cited one-in-a-billion chance of a false positive is nonsensical, given that the odds of the test being contaminated or tampered with is significantly higher than this.

CHAPTER 3

1. The Airbus A380 carries up to 320,000 liters (roughly 84,500 gallons or 260 tons) of fuel.

2. As of 2017, reportedly the oldest passenger aircraft still in regular service was forty-seven years old. It was a Boeing 737-200, serial number 20335, flying domestic routes for Airfast Indonesia (Smith 2017).

3. These jetliners were operated by 1,397 individual airlines, in 3,864 airports, serving 49,871 routes and carrying 2.97 billion passengers during the course of each year.

4. This claim admittedly deserves some caveats. There were nonfatal jetliner incidents, such as Air France Flight 66, an A-380 that suffered an uncontained engine failure and made an emergency landing in Canada. There was a fatal jetliner cargo flight: Turkish Airlines Flight 6491, a B747 that crashed in Kyrgyzstan while trying to land in thick fog, killing four crew members and thirty-five people on the ground. There were also a few accidents with turboprop airliners (as opposed to jetliners), such as West Wind Aviation Flight 280, which crashed while taking off in Canada, killing one passenger.

5. The last US fatal accident prior to 2018 had been Colgan Air Flight 3407, which crashed on approach to Buffalo on February 12, 2009. The streak was sadly undone

on April 17, 2018, when Southwest Airlines Flight 1380—a Boeing 737 en route to Dallas from New York—suffered an engine failure that peppered the fuselage with shrapnel, killing one passenger.

6. The data tell the same story when broken down by specific airframes. Take, for example, the Boeing 747, a once-popular airframe, but by no means the most common. Introduced in 1970, Boeing had built almost 1,500 747s by 2010, which had made over 17 million flights, covering a cumulative distance of over 42 billion nautical miles. During that period, they were involved in only thirty-four separate incidents resulting in multiple fatalities (Airsafe 2010).

7. As already discussed, the fact that jetliner crashes are less consequential than (for instance) reactor meltdowns is offset, in terms of the reliability they require, by the fact that there are many more airplanes than reactors.

CHAPTER 4

1. Engines and airframes are meaningfully distinct in this context. They are manufactured and maintained by different companies and regulated under separate codes (albeit codes governed by very similar principles). Airlines also purchase them separately from the airframes themselves, there usually being multiple engine options for any given airframe.

2. These offices—located in Seattle, Denver, and Los Angeles—are overseen by the FAA's Aircraft Certification Service through its Transport Airplane Directorate based in Renton, Washington. As of 2006, the Transport Airplane Directorate and its three Aircraft Certification Offices employed about 250 technical people to assist in certification activities (NTSB 2006, 68).

3. Prior to digitization, at least, this was literally true. As early as 1980, Lockheed was estimating that in the course of certificating a new Jetliner, it would submit approximately 300,000 engineering drawings, 2,000 engineering reports, 200 vendor reports, and about 1,500 letters (NAS 1980, 29).

4. This is a simplification, albeit a common one. Title 14 of the Code of Federal Regulations actually includes other rules applicable to certifying a transport-category airplane. These include Part 21, "Certification Procedures for Products and Parts"; Part 33, "Airworthiness Standards: Aircraft Engines"; Part 34, "Fuel Venting and Exhaust Emission Requirements for Turbine Engine Powered Airplanes"; and Part 36, "Noise Standards: Aircraft Type and Airworthiness Certification."

5. Advisory Circulars explain acceptable ways to comply with certification requirements. Notable among them for the argument that follows is "AC 25.1309-1: System Design and Analysis," which outlines the rationale for FAR-25's reliability requirements and explains, in broad terms, how manufacturers should demonstrate a system's compliance (FAA 1988, 2002a, 2002b).

6. So it is, for example, that FAA Advisory Circular 25.1309-1 has an appendix called "Background Information for Conducting Failure Analyses," which directs readers to "NUREG-0492," an NRC publication, and "MIL-HDBK-217," a Department of Defense standard.

7. Prior to this, US civil aviation operated without specific federal oversight. In principle, a purblind inebriate could fly loops over Times Square without breaking any specific laws. Perhaps surprisingly, calls for federal regulation were largely driven by the fledgling airlines themselves. They saw regulation as essential to commercial success, in part because more established and trusted modes of public transport, like ferries and railways, had long been subject to oversight (Hansen, McAndrews, and Berkeley 2008). Their argument was bolstered by the experiences of the US Post Office. The Air Mail Service—by far the most significant US operator of aircraft at the time—had implemented its own oversight measures, and these were demonstrating their value. It was averaging one fatality every 463,000 miles, compared to one every 13,500 miles for other commercial flights (Nader and Smith 1994, 4).

8. At first, this was the Aeronautics Branch of the Department of Commerce, created by the 1926 Air Commerce Act. The Aeronautics Branch was reestablished as an independent unit, the Federal Aviation Agency, in 1958, and in 1966, it was subsumed under the Department of Transport and renamed the Federal Aviation Administration (Briddon et al. 1974; Cobb and Primo 2003, 16). The FAA remains under the department's aegis today, although it largely works autonomously.

9. Most of the systems in an aircraft (essentially anything that draws electricity) are assessed almost entirely probabilistically, which is to say that experts calculate their reliability directly to ensure that it satisfies requirements. Structures (such as wings, windows, and fuselage sections), by contrast, are effectively still governed by proscriptive design requirements (FAA 1999a). The process calculates the reliability of these elements indirectly, incorporating them into reliability assessments of the integrated airframe by assuming that, by virtue of meeting specified design requirements, they have a preestablished failure probability (usually zero). The distinction between "structures" and "systems" is important, therefore, but often opaque. "Fly-by-wire" systems, for instance, integrate complex avionics with structural control elements that were traditionally governed through different processes. The NTSB has argued that ambiguity around this distinction has allowed catastrophic weaknesses to fall through cracks in the assessment process. (Take, for example, Alaska Airlines Flight 261, which crashed off the coast of California in January 2000 because of a stabilizer failure. Investigators apportioned some of the blame for this to the certification analysis, which had treated the failed part as "structure" rather than a "system" [NTSB 2000b; 2006, 18, 52–60]).

10. These tools are represented by a fruit salad of acronyms. They include Failure Mode and Effects Analysis (FMEA), Failure Modes and Effects Summary (FMES), Functional Hazard Assessment (FHA), Preliminary System Safety Assessment (PSSA), System Safety Assessment (SSA), Fault Tree Analysis (FTA), Common Cause Analysis (CCA), Zonal

Safety Analysis (ZSA), Probabilistic Risk Analysis (PRA), and Common Mode Analysis (CMA). Many of these, and others, are outlined in Aerospace Recommended Practice ARP4761 (SAE International 1996).

11. To complicate matters further, the operational unit of certification might also be construed as the airplane's "catastrophic failure modes." There is an ambiguity in the guidance between "catastrophic failure modes" and "flight critical systems." Organizationally, the certification process is structured around "critical systems." Confoundingly, however, the ultimate rationalization for the reliability required of those systems is often expressed in terms of "catastrophic failure modes." In these instances, the regulations do not state the maximum frequency of critical system failures but the minimum frequency of catastrophic airplane-level failure conditions (e.g., a loss of power) (FAA 2002, 10–11). The two are routinely equated, however, such that having 100 catastrophic failure modes is understood to imply 100 critical systems. (Thus, this assumes a direct one-to-one relationship between the two that is itself problematic.) Attempting to parse the close logic of certification can be trying, but this is the point.

12. Critical systems are identified in relation to the failure modes that they would create if they failed. So it is, for instance, that redundant elements that would have to fail simultaneously to create a catastrophic failure are grouped together as a single system.

13. This understanding about the relationship between the reliability of individual components to that of the wider system has its roots in the V2 rocket program. Germany's rocket engineers, who struggled enormously with failure, had started out using reliability principles based on a simple and then widely held understanding that "a chain was as strong as its weakest link." They were forced to abandon this assumption when tests proved it to be wildly inaccurate, however, and after consulting the mathematician Eric Pieruschka, concluded that the reliability of a system was actually the product of the reliabilities of its constituent elements. (Thus, the reliability of a system composed of identical elements would be $1/x^n$, where x is the reliability and n is the number of elements [Villemeur 1991, 5].) This, in turn, led to the conclusion already given—namely, that the reliability of individual components and subsystems must be much higher than the desired reliability of the system.

14. This difference is not immediately obvious in the text, which actually states that "a fleet of 100 aircraft of a type, each flying 3000 hours per annum, one or other of the Catastrophic Effects might be expected to turn up once in 30 odd years, which is close to the concept of 'virtually never'" (Lloyd and Tye 1982, 37). (This amounts to 3,000 hours per annum × 30 years = a service life of 90,000 hours.)

15. Anonymous correspondence (March 25, 2005).

16. Anonymous correspondence (March 25, 2005).

17. Anonymous correspondence (March 24, 2005).

18. The guidance appears to have gradually accrued such admonitions and caveats since its initial turn to probabilism. In part, they reflect years of misgivings about quantitative calculations voiced by the NTSB, often in accident reports that have drawn attention to unrealistic and insufficient numbers being used in certification assessments. (After TWA Flight 800 exploded outside JFK Airport in 1996, for example, the NTSB found that "undue reliance" had been placed on formal risk tools when calculating the probability of a fuel vapor explosion, and that such tools "should not be relied on as the sole means of demonstrating [compliance]" [NTSB 2000a, 297]). More directly, however, they reflect protracted debates that arose when the certification process first began to incorporate software (e.g., FAA 1982, 9; 2002, 7).

CHAPTER 5

1. The US Air Force, which routinely flies at lower altitudes, incurs commensurately greater costs relative to its flight hours. In 2003, it estimated that bird strikes were responsible for killing two of its aircrew and downing two of its aircraft every three to five years, costing the service between $50 and $80 million every year (Feris 2003).

2. Birds are poor at avoiding fast-moving objects because they react to proximity regardless of velocity. They usually take evasive action when incoming threats are about 100 feet away: a strategy that works well for them until those threats are traveling at more than about 55 miles per hour. Scientists established this, in part, by driving pickup trucks at unsuspecting turkey vultures (DeVault et al. 2014; 2015).

3. By contrast, the average hatchback, as one major engine manufacturer is pleased to observe, is barely worth its weight in hamburger.

4. Only the largest engines are subjected to the largest birds because it is thought that large birds entering smaller engines tend to strike the cowling and break up before they hit the blades (FAA 1998).

5. Engines certified since November 2007 have been assessed under revised regulations that include an additional class of bird: the "large flocking bird." Tests for this class involve a single bird (either 4, 4.5, or 5.5 pounds [1.8, 2, or 2.5 kilograms] depending on the engine size) that is directed into the fan blades rather than the engine core and require that the engine maintain some thrust (which differs at specific time points) after impact (NTSB 2009, 19, 84).

6. The density of bird traffic decreases exponentially with altitude. The FAA has determined that 71 percent of bird strikes on commercial aircraft occur below 500 feet (from ground level). Above this height, they decline by 34 percent for each 1,000-foot gain in altitude (FAA 2014, xi). Bird strikes are 1.5 times more common during arrival, but 3.4 times more damaging during departure, when the engines are working harder (Dolbeer 2007, 1).

7. Snarge is swabbed and samples are sent to the Smithsonian Institution's Feather Identification Laboratory for DNA analysis, but interpreting the samples is an imperfect

art. As of 2009, for example, the lab was unable to discriminate between multiple birds of the same species and sex (Budgey, 1999; NTSB 2009, 80fn).

8. Only dead birds are used in the tests. Even if engineers were able—and, it should be said, willing—to coax a live bird into one end of a powerful compressed-air cannon, that bird would certainly be dead upon exiting the barrel at 200 knots. A surprising number of people enquire about this.

9. It is not entirely unusual for scaling variables to have nonlinear effects in engineering (see, e.g., White 2016, 338–339). It is worth noting, however, that the JAA's conclusion has far-reaching implications for the underlying logic of modern ingestion standards, which presume a linear relationship between most variables, including mass and damage. The FAA, for instance, presumes a 1.28 percent increase in the likelihood of damage for every 100-gram increase in body mass (FAA 2014, xi).

10. It is usually considered best if the leading fan blades can slice a bird into something resembling salami before it passes through the rest of the engine. Back in the less euphemistic 1970s, the FAA listed "blades which effectively mince birds upon contact" among its "desirable engine features" (FAA 1970, §3).

11. A total of 52 percent of bird strikes occur between July and October, when birds are gathering to migrate (FAA 2014). The US Air Force, which enjoys an authoritative euphemism, often refers to such gatherings as "waves of biomass."

12. Given the vast legions of chickens that become McNuggets, the handful of birds sacrificed for the betterment of aviation safety seem to elicit an incongruous degree of public anxiety.

13. Engines have failed their bird-strike tests in the past, with meaningful consequences. Rolls-Royce can testify to this. By 1969, the British engine manufacturer had developed, at great expense, a novel fan blade made from an innovative material called "Hyfil." After meeting every other requirement, however, the blades failed their bird-strike tests. The company was forced to abandon Hyfil and return to titanium as a result (Spinardi 2002, 385). This, in turn, made the engines heavier and thus unable to meet Rolls's fuel-consumption guarantees, strongly contributing to its bankruptcy the following year (Newhouse 1982, 174). (It is worth noting, however, that these failures occurred in precertification tests.)

CHAPTER 6

1. "Ladies and gentlemen, this is your captain speaking. We have a small problem. All four engines have stopped. We are doing our damnedest to get them going again. I trust you are not in too much distress," Captain Eric Moody, British Airways Flight 009 (1982).

2. The record was short lived. It would be broken the following year by Air Canada Flight 143, which ran out of fuel mid-flight after ground crew confused imperial and metric units.

3. The eminently quotable captain would later describe the landing as being "a bit like negotiating one's way up a badger's arse" (Faith 1996, 156). From the context, it's clear that he meant by this that it was difficult (the badger presumably objecting to said negotiation).

4. The FAA calls its safety philosophy the "fail-safe design concept." It encompasses a range of principles and rubrics: an emphasis on "margins of safety" is one; "checkability" is another (FAA 2002, 6; NTSB 2006, 86). The plurality of these rubrics depends either directly or indirectly on redundancy.

5. Petroski (1994) argues persuasively for the existence of design paradigms in engineering.

6. Looking beyond aviation for a second, Sagan (1993) reports that false indications from safety devices and backup systems in the US missile-warning apparatus have nearly triggered atomic wars.

7. The type-certification process incorporates the reliability of human beings into its quantitative requirements in a similar fashion to the way that it incorporates that of aircraft structures (as outlined in chapter 4). Designs must meet proscriptive requirements (stipulating minimum cockpit space and sight lines, for instance), which the certification process then assumes will confer mathematically perfect reliability (FAA 1999b). As the FAA (1982, appx 1) puts it, "If the evaluation [of the human system interface] determines that satisfactory intervention can be expected from a properly trained flight crew," and "then the occurrence of the failure condition has been prevented."

8. One example occurred in March 1994, when the pilot of Aeroflot Flight 593 fatefully—and, in the event, fatally—passed the controls of a passenger-laden Airbus A310 to his sixteen-year-old son, Eldar (Norris and Hills 1994, 5).

9. Herein lies another way that human behavior can intrude on redundancy calculations: by undermining the independence of various elements. Investigators of a 1983 multiple engine failure on a Lockheed L-1011, for instance, determined that the engines were united by their maintenance. The same personnel had checked all three engines, and on each one, they had refitted the oil lines without the O-rings necessary to prevent in-flight leakage (Lewis 1990, 207–209).

10. Failures of reserve systems—sometimes known as "latent" or "dormant" failures—can be particularly dangerous because they are more likely to go undetected. Following the crash of USAir Flight 427, for instance, the FAA criticized the Boeing 737's rudder control system on this basis. The system involved two slides: one, which usually did the work, and a second, redundant slide that lay in reserve. The regulator argued that since the system rarely used the second slide, it was prone to fail "silently," leaving the aircraft "a single failure away from disaster" for long periods (Acohido 1996).

11. Anonymous correspondence (July 21, 2006).

12. Thus stricken, the aircraft—Air Transat Flight 236, en route to Lisbon from Toronto with 291 passengers—glided for 115 miles before making a high-speed touchdown

in the Azores that wrecked the undercarriage and blew eight of its ten tires (Airsafe 2008). The incident also has an interesting human dimension, in that the pilot is said to have exacerbated the problem by actively transferring dwindling fuel reserves to the leaking engine.

13. "In most cases," the FAA (2002b, 21) writes, "normal installation practices will result in sufficiently obvious isolation . . . that substantiation can be based on a relatively simple qualitative installation evaluation." The result, as the National Academy of Sciences (NAS 1980, 41) writes, is that "[t]he failure of a neighboring system or structure is not considered within a system's design environment and so is not taken into account when analyzing possible ways a system can fail."

CHAPTER 7

1. This is, by necessity, a slight simplification of a contested and multifaceted engineering explanation, albeit hopefully not a distorted one in respect to the central argument it is being used to illustrate. A fuller account—including a more recently proposed "fluid hammer" hypothesis (Stoller 2001)—would complicate the narrative without changing the underling point.

2. The discipline of structural engineering has a longstanding interest in crack tolerance for this precise reason. "Where human life is concerned," as one classic text on the subject puts it, "it is clearly desirable that a 'safe' crack should be long enough to be visible to a bored and rather stupid inspector working in a bad light on a Friday afternoon" (Gordon 2018 [1991], locs. 1407–1408).

3. The aviation sphere is no exception to this impulse, as is evinced in an oft-excerpted passage from the International Civil Aviation Organization's (ICAO's) first *Accident Prevention Manual*. It states, "The high level of safety achieved in scheduled airline operations lately should not obscure the fact that most of the accidents that occurred could have been prevented. . . . This suggests that in many instances, the safety measures already in place may have been inadequate, circumvented or ignored" (ICAO 1984, 8).

4. The theory rests on the observation that no complex system can function without small irregularities: spilt milk, blown fuses, stuck valves, and so on. Eliminating such irregularities entirely would be impracticable, so engineers construe them as a fundamental design premise: "failure-free" systems being defined by engineers as systems designed to accommodate the many vagaries of normal operation, rather than as systems that *never* experience operational anomalies.

5. It is tempting to see NAT as a restatement of Murphy's Law—the old adage that anything that can go wrong, will—but it is more profound than that. It is a systems-level rethinking of what it means for something to "go wrong," in that it understands failure in terms of a system's organization. Uniquely, it argues that some failures arise from the very fabric of the systems themselves: less a flaw than an emergent property of their underlying structure (Downer 2015).

6. A different reading of NAT, also supported by the text, construes normal accidents as the product of an organizational irony. Grounded in organizational sociology and drawing on contingency theory, it argues that complex and tightly coupled systems call for contradictory management styles: complex systems requiring decentralized management control (because it takes local knowledge and close expertise to monitor and understand the interactions of a diverse network of elements); and tightly coupled systems requiring centralized control (because systems that propagate failures quickly must be controlled quickly, which requires strict coordination and unquestioned obedience) (Perrow 1999, 331–335). This construal of NAT is "weaker," in the sense that it might, in principle, be resolved if sociologists could identify an organizational framework that reconciled the competing demands of complexity and coupling (e.g., LaPorte and Consolini 1991).

7. "Tensile strength" is a metal's resistance to being pulled apart, usually recorded in pounds per square inch. "Elasticity" is a metal's ability to resume its original form after being distorted.

CHAPTER 8

1. A similar concession was voiced by the deputy director of the FAA's Aircraft Certification Division, who said of the standard that governed avionics software (RTCA DO-178A) that it "recognizes that you can't test every situation you encounter" (Beatson 1989b).

2. Flight protections are also thought to have confused pilots from time to time, being implicated as contributory factors in a few accidents (such as Air France Flight 296 in 1988 [Macarthur and Tesch 1999]). Occasions where they have prevented errors go unrecorded, however, although Langewiesche (2009b) makes a compelling argument that they helped save the "Hudson Miracle" flight.

3. This process would be analogous to that described by Galison (1987) in relation to scientific experiments, wherein he argues that scientists' theories about the world become increasingly useful and veridical over time.

4. Anonymous interview. March 29, 2005.

5. Sociologists have observed the same process of rule refinement in other technological spheres. Wynne (1988, 154), for instance, calls the process of rule-making in technological systems "an ever-accumulating practical craft tradition," which he compares to case law.

6. It is also important in this regard that civil avation accidents were always contained in scope: tragic and costly, but never on a scale that could threaten cities, economies, or more than a few hundred people. (Recall from chapter 1 that jetliners might be thought of as a chronic catastrophic technology, in the sense that their extreme reliability requirements owe as much to the numbers in which they are operated, as they do to the extreme hazards of them failing.)

7. Personal communication, May 4, 2005.

8. One example was the so-called airframe revolution, which saw the introduction of retractable landing gear, aluminum stressed-skin structures, wing flaps, and other innovations that became standard.

9. Engines, which are manufactured and purchased separately from the airframe, have followed a distinct but similarly restrained design trajectory—undergoing a shift from low- to high-bypass turbofans in the early 1970s, whereupon they became much wider.

10. Anonymous correspondence, March 20, 2005.

CHAPTER 9

1. Concorde has an uncommon nomenclature, in that it does not receive an article in common English usage (i.e., it is not referred to as "*The* Concorde").

2. It should be said, however, that the Central Intelligence Agency (CIA) was arguably the true pioneer here. The SR-71 was derived from the A-12, which had been developed for the CIA two years earlier.

3. Manufacturers of much smaller, private aircraft could afford to move more quickly. In 1989, for instance, Beech Aircraft Corporation (later Beechcraft) delivered its first Starship: a six-to-eight-passenger turboprop built almost entirely from composites (Huber 2004; Warwick 1986).

4. Flight 961 was forced to land abruptly in March 2005 after losing its composite rudder while high above the Florida Keys. No lives were lost, but investigators found that it had been a close call. The rupture had dangerously stressed the tail itself, the loss of which would have been unsurvivable (TSB 2007). The formal investigation into the incident was not completed until 2007, but long before the report's publication, the industry concluded that the rudder probably had had an undiscovered stress fracture. To prevent its reoccurrence, Airbus issued a directive requiring operators of A300-series airplanes with composite rudders to inspect their outer surfaces for latent damage with an "acoustic tap test" (NTSB 2006, 2).

The incident did not end there, however. In a twist that speaks to the inescapable dangers of epistemological uncertainty, the problem was considered resolved until eight months later, when an unrelated but serendipitous second incident raised new questions about Flight 961's underlying cause. An accidental blow to the rudder of a FedEx A300-600 led maintenance engineers to examine it more carefully than would usually be required. Their examination revealed a shocking amount of damage, but not from the blow. It turned out that, prior to the accidental insult, hydraulic fluid had been leaking from the plane's control system and attacking the rudder's composites. The materials had deteriorated to a point where the structure might easily have failed catastrophically midflight (TSB 2007; Marks 2006). Since the hydraulic fluid had affected the inner rather than the outer surfaces of the rudder, moreover, its damage

would not have been revealed by standard test protocols, implying that Airbus's safety inspection program (with its tap tests) was inadequate and its airplanes had all been in danger.

The incident had all the hallmarks of a rational accident born of innovation-driven uncertainty. The implications of hydraulic fluid leaking into composite structures went unrecognized because experts never considered the interaction as a possibility (although the basic science was understood). Tests and models had not revealed the danger because nobody had thought it a relevant variable to test or model; and maintenance practices were not primed to look for the kind of damage that it caused for the same reason (TSB 2007; Marks 2006). As with most rational accidents, moreover, it was informative, spurring a joint NTSB-Airbus investigation into fluid contamination, which informed new maintenance tools and revised procedures (NTSB 2006a).

5. Interestingly, the problematic reliability of military airframes plays an important role in the logic around arms sales. It makes any nation that procures advanced US fighters continuously dependent on the US for spare parts and upkeep, without which their air power would quickly degrade.

6. Reliability problems with military hardware are certainly not exclusive to aviation. The DoD, which leans heavily on technological advantages to achieve its strategic goals, has long struggled with the reliability of its systems. This first became apparent in the early years of the Cold War, when it was discovered that a remarkable portion of the department's expensive, high-tech paraphernalia was down at any given time, and the resources required to diagnose and repair faults was becoming unsupportable. Logistic data from this period testify to the scale of the problem. By 1945, for instance, the navy was annually supplying a million replacement parts to support 160,000 pieces of equipment, and by 1952, it was spending around two dollars every year to maintain each dollar's worth of technology (Coppola 1984, 29). The US defense establishment was acutely concerned by this ever-deepening "reliability crisis," which it saw as jeopardizing its ability to fight wars (Jones-Imhotep 2000, 145).

CHAPTER 10

1. From this perspective, the real mystery of Concorde was not why it disappeared, but why it was built at all. The answer to this question would undoubtedly involve national pride, and the power of what Jasanoff and Kim (2015) would call the Anglo-French sociotechnological imaginary. Concorde, as *Le Monde* once put it, "was created largely to serve the prestige of France. [It was] the expression of political will, founded on a certain idea of national grandeur" (Harriss 2001). Its development is probably better understood in relation to something like the Moon landings than in relation to civil aviation in isolation.

2. A parallel argument could be made about the industry's unusual commitment to interrogating its failures. Recursive practice is expensive. Many organizations struggle to effectively learn from disasters, not least because ambiguities and interpretive

flexibilities allow them to construct casual interpretations that favor their interests. For example, useful insights are often overlooked if they imply expensive redesigns or costly liabilities (see, e.g., Downer 2014; March et al. 1991; March and Olsen 1988; Hamblin 2012). The civil aviation sphere is far from immune to this trend (see, e.g., Perrow 1983). But with a few notable exceptions (which will be discussed elsewhere in this book), it rarely avoids grasping the nettle where design weaknesses are concerned. Its willingness to do this is interesting for the same reason that its commitment to design stability is interesting, therefore, but exploring the latter offers generalizable insights that explain both behaviors.

3. It has long been understood, for instance, that relocating engines to overwing nacelles would improve both the efficiency and acoustic properties of jetliners (Berguin et al. 2018). In a similar vein, "double-bubble" or blended "wing-body" designs would theoretically allow jetliners to carry many more passengers at a significantly reduced cost. Neither this, nor overawing nacelles, are new or untested concepts. Experts have experimented with prototype wing-body designs since the 1940s, for example, and the aforementioned B-2 Spirit bomber testifies to the viability of both innovations in large airframes.

4. It should be noted that scholars and industry observers sometimes argue that manufacturers were reluctant to implement advanced composites because the materials' higher manufacturing and repair costs would render airframes uncompetitive (e.g., Lenorovitz, 1991; Rogers 1996, vii; AWST 1990; Tenney et al. 2009, 2–3; Slayton and Spinardi 2016, 51). This argument is less than satisfying, however, for a range of reasons. First, it doesn't explain the industry's eagerness to adopt the new materials piecemeal over successive generations of airframes. (If the 787's composite airframe was a poor cost proposition, then presumably the B777's composite tail should have been less optimal than a traditional tail for the same reasons, and so on.) Second, it is far from clear that composites increased maintenance and production costs. Abandoning aluminum undoubtedly disrupted established production regimes and maintenance practices, but the financial implications of this were highly contested within the industry. For instance, many experts at the time argued that higher individual maintenance costs were more than offset by a reduced frequency of maintenance, and higher raw material costs were offset by reduced labor costs (e.g., Jones 1999, 48; Raman, Graser, and Younossi 2003, viii; Barrie 2003; Mecham, 2003, 2005). And although the 787 program was indeed beset by manufacturing difficulties and delays, the larger part of these issues arose less from the composites than from software and electrical issues, together with Boeing's expanded outsourcing (Associated Press 2009; Waltz 2006; Marsh 2009). Third, arguments about increased maintenance and production costs rarely frame those costs against the substantial fuel economies that composites promised. Finally, the argument that airframers avoided composites because the materials were uncompetitive is contradicted by the parallel argument—often advanced by the same observers (e.g., Slayton and Spinardi 2016, 51)—that Boeing's decision to embrace them in the 787 was prompted by a perceived need to compete

more aggressively with Airbus. Companies rarely try to compete more aggressively by introducing less competitive products. All these arguments are more complex than this note can reasonably accommodate, and none are conclusive, but collectively they paint a compelling picture.

5. The regulatory implications of outsourcing were evident in 2013, for instance, when the FAA grounded all B787s after two significant battery fires. The batteries were built by Yuasa, a Japanese manufacturer subcontracted by Thales, a French manufacturer that designed and built a specific electrical system. The NTSB report on the incident highlighted inherent deficiencies in this chain of oversight, wherein the FAA oversaw Boeing, which oversaw Thales, which oversaw Yuasa (NTSB 2013; Bonnín Roca et al. 2017).

6. "Tacit knowledge"—a term coined by Michael Polanyi (1958)—is used widely in the STS literature (e.g., Collins 1982, 2010; MacKenzie and Spinardi 1996). It refers to the information, skill, and experience that are vital to a task but get marginalized, ignored, or obscured by formal accounts because they are uncodified, uncodifiable, or both.

7. Anonymous personal communication, May 19, 2009.

8. Note that few of the factors that keep aviation regulators from being able to police jetliner designs are specific to the aviation industry. Regulators in other spheres are in a similar position regarding their expertise and staffing, and most have interests that are aligned, to some extent, with the industries they regulate. The *Report of the Deepwater Commission*, for instance, made a point of emphasizing regulators' failure to curb the industry's safety compromises, arguing that it was unable to exercise the "autonomy needed to overcome the powerful commercial interests that . . . opposed more stringent safety regulation" (NCBP 2011, 67).

9. It is worth noting that Concorde, again, offers tangible evidence of regulators' powerlessness to enforce stability in defiance of powerful structural interests. Upon welcoming the airplane's inaugural US flight, Undersecretary of Transportation John Barnum declared that it had "met with ease all safety and technical requirements" (Simons 2012, loc. 4717–4722). We could call this a diplomatic truth, but it is probably more accurately described as a lie. US regulators had very serious concerns about the airframe's novel design, which violated a range of several seemingly unambiguous certification rules. (Its engines were not separated on the wing, for example—a factor that became significant in its fatal accident—and it was unable to meet regulations governing minimum fuel reserves [Donin 1976, 54–7]). Concorde had met its US safety and technical requirements, true, but only because those requirements had been warped to accommodate the new design, the regulators' misgivings having been overruled after intervention from the highest levels of government (Simons 2012, loc. 3752–3764, 3967–3976, 3777–3779). The UK and France were both deeply invested in their jetliner's success, and to deny approval would have been to instigate a fierce diplomatic row, and potentially a costly trade war (Simons 2012; Harriss

2001). More broadly, the approval reflected a tangible sense that supersonic flight was the future, and authorities needed to bow to the inevitable lest the US be "left behind" (Simons 2012, loc. 3379–3380).

10. The most notable of these incidents were the losses of American Airlines Flight 191 in May 1979 and United Airlines Flight 232 (the triple hydraulic failure discussed in chapter 6) in July 1989.

CHAPTER 11

1. The MAX was fitted with two AOA indicators, but MCAS took information from only one of them. The airplane could be fitted with an alert to indicate when the two indicators disagreed, but this was an optional extra, for which Boeing charged more.

2. Despite Boeing being the dominant partner in the merger, former McDonnell executives wound up taking leading roles in Boeing's management, giving rise to a longstanding joke about McDonnell effectively buying Boeing with Boeing's money (See, e.g., Useem 2019).

3. The NTSB's recommendation came in response to a DC-8 that overran a runway in Alaska in 1970. The forces involved in the accident were easily survivable, but aircraft had caught fire, and 47 of the 229 people on board died. The postcrash investigation found cyanide in the blood of the bodies recovered and determined that most who died were killed not by the fire directly, but by toxic smoke, which could have been avoidable. The finding led it to recommend that the regulators begin to explore the fire toxicity of cabin interiors (NTSB 1972). In the wake of the NTSB recommendation the FAA established a research effort led by the Civil AeroMedical Institute (CAMI) in Oklahoma City. It concluded that among the many materials that make up a jetliner cabin—everything from seat furnishings to wire claddings—were substances that produced toxic fumes such as carbon monoxide and hydrogen cyanide when heated or burned.

4. Cabin fires are complicated phenomena. Burning conditions can dramatically alter smoke composition and toxicity, for example, even when the same materials are involved (Chatuverdi 2010, 2). So it was that in the decades-long contest over fire-toxicity regulations, airframers questioned the FAA's scientific justifications on a wide number of fronts. The kinds of uncertainties that they raised are well illustrated, however, by a series of simulated cabin-fire tests that the FAA conducted in 1978, with old Lockheed C-133 fuselage. The tests all resulted in "flashovers": explosive conflagrations that would almost inevitably have killed passengers long before toxic fumes became an issue. And experts concluded from this that measures to mitigate toxic fumes were pointless, and regulatory efforts should focus instead on preventing fires in the first place. Others, however, questioned the tests' representativeness. They argued that flashovers occur when smoke and fumes are trapped in a poorly ventilated compartment where they can concentrate and heat until they ignite, and the FAA's

test fuselage—an intact structure, fire-hardened for reuse and largely sealed to outside ventilation—provided an ideal space for this to occur. But in most real aviation fires, they contended, cabins tend to be much better ventilated, partly because the accidents that instigate them often rupture the fuselage, but also, more prosaically, because people open the doors. Because of this, they claimed, flashovers were much less likely than the tests suggested, and toxicity much more important. The accident record spoke for itself, they protested. Fires did not always involve flashovers, and passengers were demonstrably dying with toxic levels of combustion products in their blood. With experts disagreeing on the meaning of its tests, the FAA eventually conceded that more research was required before it could legitimately require the industry to invest in changes, a position that it still held two decades later (Weir 2000, 87–88; Flight Safety Foundation 1994; Chatuverdi and Sanders 1995, 1; Chatuverdi 2010, 2; ETSC 1996).

5. Perrow, ever insightful, comments on this relationship, albeit in passing. The aviation industry is preoccupied with accident prevention over damage mitigation, he notes, because reducing the injury and death rate from accidents "has little or no effect upon [its] economic variables" (Perrow 1999, 163).

CHAPTER 12

1. Elements of the type-certification process seem actively designed to hide uncertainties and disagreements from public view. Take, for instance, the way that it treats "issue papers": documents that codify expert misgivings and points of contention that arise during assessment. If made public, they would offer a window into the indeterminacies and negotiations that underpin key rulings, but the process designates them as "draft" documents and exempts them from disclosure (NTSB 2006, 42–43, 51).

2. Most trace the roots of this milieu at least as far back as the nineteenth century, when a slew of probabilistic tools emerged as both the product and the driver of new efforts to order and manage the economic world. By many accounts, however, formal audit practices really came of age in the early- to mid-twentieth century—metastasizing, rule by rule, metric by metric, into the daunting modern architecture of statutes, practices, and regulations that Power (1997) calls the "audit society." Chroniclers of this process highlight a wide variety of processes and protagonists in telling this story, including but not limited to the Army Corps of Engineers, flood control, Taylorism, operations research, the RAND Corporation, Cold War defense procurement, the financial services industry, and management consultancy. Many locate its primary drivers in the interfaces between organizational structures, construing audit practices as products of efforts to control and communicate (e.g., Hacking 1990; Porter 1995; Lampland 2010; Desrosieres 1998; Verran 2012; Power 1997).

3. Another structural engineer, A. R. Dykes, artfully captured the essence of white-boxing in a 1976 speech to the British Institution of Structural Engineers when he described his discipline as "the art of modeling materials we do not wholly understand

into shapes we cannot precisely analyze as to withstand forces we cannot properly assess, in such a way that the public has no reason to suspect the extent of our ignorance" (Schmidt 2009, 9).

4. A similar sentiment is evident in Burrows's (2010, 560) epitaph on the accident: "Challenger was lost because NASA came to believe its own propaganda." He writes. "The agency's deeply impacted cultural hubris had it that technology—engineering—would always triumph over random disaster if certain rules were followed. The engineers-turned-technocrats could not bring themselves to accept the psychology of machines with abandoning the core principle of their own faith: equations, geometry, and repetition—physical law, precision design, and testing—must defy chaos. No matter that astronauts and cosmonauts had perished in precisely designed and carefully tested machines. Solid engineering could always provide a safety margin, because the engineers believed, there was complete safety in numbers."

5. Anonymous interview, March 29, 2005.

6. Anonymous interview (February 2, 1996), courtesy of Donald MacKenzie.

CHAPTER 13

1. Much the same might be said of Perrow's normal accidents. Indeed the performance of modern jetliners implies that such accidents are extremely rare, even by the standards required of catastrophic technologies.

2. In an absolute sense, it is rare that reactors melt down, drilling platforms explode, financial instruments crash, or errant deterrence networks almost instigate atomic wars, but such disasters and near-misses occur far more often than experts promise and predict (e.g., Ramana 2011; Sagan 1993). Statistically, this record should be damaging to the idea that ultra-high reliabilities are achievable and knowable, but it is not.

3. It is worth restating here that aviation safety ultimately depends on a wide spectrum of further requirements regarding, for instance, pilot training, maintenance operations, and much more. Navigating the epistemology is only one piece of the puzzle, even if it is a vital piece.

4. Such was the conclusion of Parker Jones, a supervisor of nuclear safety at Sandia National Laboratories, who wrote a report on the accident in 1969, which was declassified in 2013 (Pilkington 2013). And even the conclusion of Secretary of Defense Robert McNamara (quoted in Schlosser 2013, 301), who also references a similar incident in Texas.

CHAPTER 14

1. As the disaster unfolded, the US State Department, unlike several of its European counterparts, chose not to evacuate its embassy in Tokyo. It justified this decision by referring to the Nuclear Regulatory Commission's own radiological hazard modeling,

which, it claimed, showed absolutely no threat to the city. What it didn't say is that the hazard model—named RASCAL, and designed to predict small leaks—maxed out at a range far short of the dangers that it was being used to predict, and so it could not have shown a threat to Japan's capital under any circumstances (Lochbaum et al. 2014, 63).

2. In the discourse around Fukushima, this idea has come to be known as "the myth of absolute safety" (e.g., Hirano 2013; Srinivasan 2013; Onishi 2011; Soble 2014; Nöggerathet al. 2011).

3. As in civil aviation, there is now a tension between the practice and the portrayal of reactor safety assessment. In their internal discourse, leading reactor assessment experts speak eloquently about the subjectivities and limitations of failure predictions. But in public contexts, the same experts vigorously promulgate an understanding of those assessments as rule governed and definitive (Apostolakis 1990; 2004; 1988; Wu and Apostolakis 1992, 335; Wellock 2017, 694; Miller 2003, 183–184). See, for instance, the NRC guidelines for external risk communication. (NRC 2004). Intended for internal consumption, the document instructs regulators to avoid any equivocation in their public assertions of reactor safety because professions of uncertainty "reinforce feelings of helplessness and lack of individual control" (NRC 2004, 38). "Avoid making statements such as 'I cannot guarantee . . .' or '[t]here are no guarantees in life'," it advises, because "statements like these contribute to public outrage" (NRC 2004, 38).

4. Chapter 13 suggested that the promise of perfect reliability was a problematic epiphenomenon of the positivist portrayal of ultrahigh reliability management in civil aviation. From an historical perspective, however, it is clear that this promise was the intended purpose of portraying reliability this way (albeit in the civil nuclear sphere).

5. Propellor-driven airliners predate jetliners, of course, but then nonnuclear power stations predate reactors.

6. It is worth noting, in passing, that this distinction—wherein jetliners require ultrahigh reliability because of their numbers, and reactors do so because of their extreme hazards—becomes significant if we consider the implications of externalities such as sabotage. Aviation safety assessment explicitly does not cover sabotage (e.g., bombings or hijackings) (FAA 2002a, 9). And it can afford to do this because sabotaged jetliners are, to a very limited extent, societally tolerable. This is not the case with reactors, however, as even one sabotaged reactor might be catastrophic. And, unlike most failure mechanisms, there is no reason to imagine that the number of sabotage attempts should be a direct function of how many reactors there are in operation. (This is, admittedly, a complex issue, with considerations not easily captured in a note. Reactors are presumably easier to defend against bombs and hijackers than jetliners are, for instance, but then again, hijacked jetliners could pose a threat to reactors.)

7. Like their aviation counterparts, nuclear authorities (and reactor operators) do endeavor to investigate accidents and anomalies for generalizable insights that can be

applied to other reactors (see, e.g., Perin 2005; Schulman 1993, 363). Post-Fukushima, for example, future plants are unlikely to locate all their backup generators underground, where they are vulnerable to flooding. These efforts are severely constrained by the industry's limited service experience and design stability, however, and, especially in the context of catastrophic failure, they often exhibit a spectrum of weaknesses identified by March, Sproull, and Tamuz (1991, 6–7) in an essay on the problems of extrapolating from limited data. (See also Downer 2014, 2015, 2016; Hamblin 2012; Perin 2005; and Rip 1986).

8. The Price-Anderson Act serves this function in the US, but such statutes are common to almost all nuclear states because the insurance industry is unwilling to cover reactors against catastrophic accidents.

9. The industry invariably counts reactors separately when calculating "cumulative safe operational hours," but groups them into plants when calculating "cumulative accidents."

10. A 2010 NRC study reviewed the implications of the revised seismology for US reactors. It found that earthquake risks to some were significantly higher than had been assumed, to the point where potential quakes might exceed some plants' design bases (NRC 2010; Dedman 2011; Lochbaum et al. 2014, 115). The report ranked plants by risk, finding Indian Point Unit 3, thirty-five miles north of Manhattan, to be the most vulnerable.

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