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

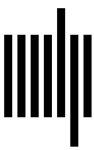

The MIT Press

3 THE AVIATION PARADOX: ON THE IMPOSSIBLE RELIABILITY OF JETLINERS

The bulk of mankind is as well equipped for flying as thinking.
—Jonathan Swift

Which is now a more hopeful statement than Swift intended it to be.
—Will Durant

3.1 UNDENIABLY RELIABLE

A TRIUMPH OF MUNDANITY

In late October 2015, Norwegian Flight 7015—a Boeing 787 Dreamliner: red and white with a decal of a Scandinavian luminary on its tail—departed London en route to New York. Shortly after leaving the tarmac, its captain greeted his passengers over the public address system, wishing them a pleasant flight. He made no mention of an Atlantic storm that loomed portentously in the airplane's projected path. About an hour into the journey, as the skies began to darken, the flight ran into turbulence. A second, terser announcement instructed passengers to fasten their seatbelts. Soon afterward, the airplane's advanced in-flight entertainment system seized and had to be restarted. A small child wailed the bitter, mournful howl of a banshee heralding terrible portents.

About two movies later, the airplane touched down gently at JFK Airport—another wholly unremarkable journey, with all the romance and wonder of the modern jet age: small discomforts, trivial glitches, and officious

security hurdles. It presumably flew around the storm; I was watching the movies rather than the flight tracker.

Excuse the melodrama. The point is simply to highlight the fascinating mundanity of modern air travel. The chapters that follow are punctuated by jetliner terrors and tragedies. Passages that begin with a date and a flight number invariably end with a mortality figure. When reading these stories, however, it is important to remember that such tragedies have become extraordinarily rare. The bigger picture is that every day, in almost every weather condition, private companies around the globe pack millions of people into giant machines, lift them high above the clouds, and then return them gently to Earth hundreds, or sometimes thousands, of miles from their point of departure. They have done this millions of times a year, for decades, and in recent decades they have done it with vanishingly few catastrophic failures. And—incredibly—this impresses almost nobody. The days when passengers would routinely applaud safe landings have long passed.

Understood properly, however, the wholly unremarkable, taken-for-granted safety of civil aviation is among the greatest engineering achievements of the last century—greater perhaps than the Moon landings. As outlined earlier, jetliners are inherently challenging technologies. The largest of them integrate upward of seven million components and over 170 miles of wiring. These components are highly constrained by weight and volume, often forcing engineers to work much closer to design tolerances than they would otherwise prefer. Combined, they make integrated systems of staggering complexity, which are expected to carry large volumes of flammable liquid[1] high into a stochastic operating environment that constantly threatens to punish any loss of power, control, or structural integrity. And they are expected to do this for decades, in a broad range of climates, and with minimal downtime for inspection and maintenance,[2] without ever failing catastrophically (see, e.g., Morris 2017).

As chapter 2 explained, this expectation ought to be quixotic. By all rights, it should be impossible, epistemologically, to know a system of such complexity well enough to achieve such reliability. That would require anticipating—and accurately assessing—every possible failure condition that could occur over billions of hours of enormously varied operation. The mundanity of modern air travel is fascinating, therefore, because it directly challenges this conclusion. Seemingly in defiance of some of finitism's core tenets, our experience with modern jetliners suggests that they are, in fact, ultrareliable.

KNOWABLY RELIABLE

Jetliners are not unique in being publicly understood to be ultrareliable. Indeed, it is our lived experience with most catastrophic technologies that they very rarely fail. The world's baroque atomic arsenals are yet to erupt in a single accidental missile exchange. Meltdowns—nuclear and financial—are once-in-a-generation events. And, in all cases, experts claim extraordinary reliabilities of these systems.

As we have seen, however, it is plausible to argue that the experts must be wrong with respect to most catastrophic technologies because, in most cases, our lived experiences with those technologies can be highly misleading. If a system operates for years without a major accident, that intuitively feels like compelling evidence of extreme reliability. But if that system is expected to operate for hundreds of thousands of years between failures, and there are only a modest number of instances of that system in operation, then those years of operation don't prove much. It could be far less reliable than was claimed and decades might still pass between catastrophic failures. Most catastrophic technologies are like this. They accrue service experience so slowly that the mean-times-to-failure that experts assert of them cannot be tested against empirical data. And in such circumstances, expert assertions about their reliabilities depend on abstractions—lab tests and theoretical models—the accuracy and relevance of which are open to principled critique.

As the closing chapter of this volume will explain in more detail, reactors exemplify this relationship. Even decades after the dawn of nuclear energy, the reliability attributed to even the oldest nuclear plants cannot be demonstrated statistically. In essence, this is because they are few in number and highly varied in their designs. As of 2016, there were only 444 reactors in operation worldwide (NEI 2016), no two of which were exactly alike and many of which differed very substantially from each other. In combination, these conditions are an actuary's nightmare: the small number of operating reactors limits the rate at which they collectively accrue service data, and the diversity of their designs limits the statistical relevance of each reactor's service data to that of the wider group. (Just as it would be misleading to invoke the service history of a 1972 Ford Escort to cast doubts on the reliability of a 2023 Tesla Model Y, so it would be misleading to invoke the performance of an early UK molten salt reactor to validate the reliability of a third-generation US pressurized water reactor.) The result is that, even if there had never been a single catastrophic meltdown—which, of

course, there has been—statisticians would still be unable to demonstrate, actuarially, that reactors achieve the ultrahigh reliabilities ascribed to them (Raju 2016). The only way for experts to know that reliability, therefore, is through elaborate predictions, the validity of which can be challenged on epistemological grounds.

Jetliners are fundamentally different in this respect. They operate in much larger numbers than most other catastrophic technologies, and (as we will see at some length) with far less variation among their designs than is probably intuitive. Airframers often sell thousands of a given jetliner type, and airlines operate those jetliners almost continuously for years; sometimes upward of eighteen hours a day. This combination of operating volume and design commonality allows jetliners to quickly generate huge volumes of relevant service data. In 2014, for example, there were 25,332 commercial jetliners in regular service.[3] Together these jetliners averaged about 100,000 flights every single day, accruing over 45 million flight-hours over the course of the year (ATAG 2014).

Under these conditions, the service data on jetliners eventually become statistically significant, even relative to the ultrahigh levels of reliability required of them. So it is that the predicted failure performance of jetliners—unlike that of reactors and almost any other catastrophic technology—can be examined actuarily by looking at how often they have actually failed in service. Aviation experts must still assess the reliability of new designs predictively, via tests and models, as a condition of those designs being allowed to operate in the first place. But, uniquely, that reliability—and thus the accuracy of the predictive assessments—can then be examined against experience. And somewhat confoundingly, at least from a finitist perspective, jetliners in the past have performed about as well as experts predicted they would—better even.

SAFETY IN DATA

The accident statistics for modern civil aviation are truly remarkable. In 2017, for example, no passenger-carrying commercial jetliner was involved in a fatal accident anywhere in the entire world (Calder 2018).[4] None. In the course of that year, airlines moved more than 4 billion people on almost 37 million flights, with a total hull-loss rate—including nonfatal losses, cargo flights, and turboprop aircraft—of 0.11 per 100,000 flight hours (or 1 in every 8.7 million departures). To be fair, this was a record year, but broader statistics

still speak almost as eloquently to the safety of modern air travel. For example, the five-year global hull-loss rate from 2012 to 2016 was only marginally less perfect, at 0.33 per 100,000 flight hours, and over the ten-year period from 2002 to 2011, there were 0.6 fatal accidents for every 1 million flights globally; or 0.4 accidents (and 12.7 fatalities) per million hours flown (CAA 2013).

The data are even more impressive when framed in relation to specific operating regimes. To date, for instance, no fatal accidents involving a UK-registered jetliner have occurred since 1989, when British Midland Flight 92 lost power and crashed outside East Midlands Airport. That makes well over three consecutive decades of catastrophe-free air travel. The US record is barely less pristine, especially relative to its much larger volume. In 2017, US-registered jetliners carried roughly 841 million people on scheduled commercial flights, amassing 17,853,752 flight hours. And, by the end of that year, it had been almost a decade since their last fatal accident (Insurance Information Institute 2018; Calder 2018; Lowy 2018; Orlady 2017).[5] Speaking in 1989, Boeing's safety manager captured the industry's already impressive reliability achievements in a colorful statistic. "If you were born on an airliner in the US in this decade and never got off," he said, "you would encounter your first fatal accident when you were 2300 years of age, and you would still have a 29% chance of being one of the survivors" (Orlady 2017, 23).[6]

Remember also that a significant percentage of the accidents in these data have little or nothing to do with failures of the jetliners themselves. They include runway incursions, terrorist attacks, pilot errors, pilot murder-suicides, and other incidents. This makes the industry's reliability record even more impressive, for, as discussed previously, reliability is a *necessary* condition for safe air travel, even if it isn't a *sufficient* condition. It is difficult to find good data on the number of aviation accidents attributable to technical failure—not least because that distinction is highly contestable (chapter 6 will touch on this in more depth)—but that number is necessarily smaller than the total number of accidents.

TRUST IN NUMBERS?

Academic observers of civil aviation, especially its critics, sometimes parse these numbers in ways that complicate the conclusion that civil aviation is safe. Fraher's (2014) broadside against US aviation regulation is exemplary in this regard. Presumably conscious that the accident data might undermine

the coming critique, she opens her book by addressing the industry's record. To this end, she equates the purveyors of aviation safety data with Wall Street analysts prior to the 2008 financial crash: calling the "apparently low aviation fatality rate" an "illusion," attributable to "luck and data manipulation" (Fraher 2014, 14–16). Her argument is misleading, especially as it pertains to the core question of this volume, but it is worth pausing here to consider its main themes. For statistics can indeed be deceptive, especially those pertaining to risk (e.g., Blastland and Spiegelhalter 2013; Power 2007), and the impressive reliability of civil aviation is integral to the argument that follows.

Fraher substantiates her skepticism on three main grounds. First, she argues that past performance is no guarantee of future safety, and might even jeopardize it. (To support this point, she cites evidence of growing complacency in the industry arising from its past success.) Second, she argues that aviation's impressive safety data tend to consider only large jetliners rather than other categories of civil aircraft, such as private aircraft, cargo aircraft, and small turboprops, which have a less robust record. And third, she argues that small numbers are deceptive because a tiny increase in the number of accidents can have a dramatic effect on the overall service record. "[A]viation quants admit the occurrence of just one accident would sway their results." She writes, noting that if the 2009 crash of Air France Flight 447 was included in the 2013 data, then fatality statistics for that year would drop "from 1 in 22.8 million to about 1 in 14 million flights: a 37 percent decline" (Fraher 2014, 16).

All these arguments contain some insight, especially in the context of Fraher's argument about regulation. For our purposes, however, none of them has much purchase.

The first is ultimately irrelevant to the central argument of this volume, which pertains to the fundamental achievability of ultrahigh reliability. It may be true, as Fraher implies, that US aviation safety is about to experience a dramatic decline due to complacency within the industry (this is actually plausible, as we will see in chapter 11). But even if aviation accidents suddenly soared for the reasons she proposes, the industry would nevertheless have proven itself capable of achieving epistemologically confounding levels of reliability over a sustained period of decades. And, as such, it would still pose a dilemma for finitists with important ramifications for technology governance. (It is probably worth noting, moreover, that critics have been making similar predictions about complacency for at least a quarter of

a century without their fears ever manifesting in the statistics [e.g., Schiavo 1997; Nader and Smith 1994]).

The second argument is slightly more substantial. It is true that most civil aviation safety statistics refer only to large commercial jetliners, and the accident rate is higher in other categories of civil aviation. At the same time, however, large jetliners carry far more people than any other category of airplane, even if they constitute only about half the total departures. It also important to note that the safety record of wider civil aviation is only marginally less impressive, especially as it pertains to large jetliners in cargo roles. (FedEx's aircraft are not exactly raining from the sky.) And, crucially, it is the reliability of jetliners, specifically, that forms the crux of the discussion here. As in the previous point, if large passenger jetliners can achieve ultrahigh levels of reliability then that still poses epistemological questions, regardless of how other categories of aircraft are performing.

The third argument is just bad statistics. It is certainly true, as Fraher asserts, that small numbers can have counterintuitive implications for how we should understand the safety of jetliners or any catastrophic technology. Indeed, this is an important insight to use when reckoning with such technologies. (If the world erupted in an accidental atomic war tomorrow, for instance, then decades of zero people being killed by malfunctioning deterrence networks would become statistically meaningless in a single devastating instant.) In the context in which she is invoking it, however, the insight does not count for much.

Take, for example, her specific point that a single accident like Air France Flight 447, had it occurred in 2013, would have changed that year's fatality statistics from "1 in 22.8 million to about 1 in 14 million flights," leading to a "37 percent decline" in aviation safety. This claim isn't inaccurate per se, but it is highly misleading. This is because—somewhat ironically, given Fraher's admonition about small numbers—it overinterprets small numbers. In the parlance of risk calculation, we might say she is invoking relative risk where absolute risk would be more appropriate (Spiegelhalter 2017). This is simply to say that while "37 percent" sounds like a substantial increase, the data are always going to be "lumpy," so to speak, on a yearly basis when dealing with so few accidents every year, especially when expressed in percentage terms. If we accept that there were zero commercial jetliner fatalities in 2017, for instance, then a single fatality in 2018—expressed as a percentage change, as per Fraher's example—would represent an infinite increase in risk

from the year before, while barely altering the overall level of aviation risk for the decade as a whole. While it is true that the loss of Flight 447 would have altered the statistics for 2013, therefore, it would not have meaningfully altered them for the period 2010–2020. Besides, even if we were to allow that aviation risk can be expressed in relation to an isolated year—which we should not—one fatality for every 14 million flights is still a remarkable (and epistemologically puzzling) record.

This is all to say that the data on jetliner service paint a compelling picture, even if they might be contestable at the margins. Relative to many data sets, the records on commercial flight safety are enviably comprehensive: vanishingly few departures, accidents, or fatalities go unrecorded. And while those data do not necessarily imply that civil aviation's structures and practices are beyond improvement, or that its record will not deteriorate in the future, they do speak authoritatively and affirmatively to the specific question of whether it is possible to build and assess ultrareliable jetliners. And, in doing so, they present a paradox.

3.2 A PARADOXICAL ACHIEVEMENT

PROBLEM AND PROMISE

Chapter 2 made a principled argument that ultrahigh reliability should be unachievable and unverifiable in complex technological systems. Jetliners ought to exemplify this argument. They are highly complex and inherently difficult systems. The reliability required of them is extraordinary—as demanding as that required of any system ever built, with mandated mean-times-to-failure north of hundreds of millions of hours.[7] (Chapter 4 outlines the specific requirements in more detail.) And, as with other catastrophic technologies, that performance must be established predictively, before a new type of jetliner can enter service. A finitist, STS understanding of engineering knowledge implies that these conditions should impose impossible demands on the experts responsible for achieving and measuring that performance. The task implies a depth and certainty of technical knowledge that is incompatible with the inherent ambiguities of the tests and models from which that knowledge is supposed to be derived.

So it is that civil aviation's statistically visible service record, together with the ultrahigh reliability to which that record testifies, raise academic questions with far-reaching policy implications. Contrary to core tenets of

the STS literature, the history of modern jetliners strongly suggests that the ultrahigh reliabilities claimed of catastrophic technologies are not in fact unachievable or unverifiable. In this sphere at least, experts seem adept at interrogating complex systems to such an extraordinary degree of accuracy that they can accurately control, and predict, the failure performance of those systems over billions of hours of operation.

Let us call this the "aviation paradox."

This paradox is the reason why jetliners are a uniquely interesting catastrophic technology, and why the structures through which they are managed are worth exploring. It has broad implications for the governance of all catastrophic technologies, and potentially great promise. Because if aviation experts can surmount the hurdles of ultrahigh reliability in jetliners—the only catastrophic technology for which we have statistically meaningful failure data—then it seems intuitive that other experts, using analogous tools and practices, might do the same for other complex systems with similarly extreme reliability requirements. Insofar as we can successfully design and assess ultrareliable jetliners, in other words, then why not reactors, deterrence networks, drilling platforms, banking computers, or anything else?

Let us turn, then, to the question of how civil aviation experts manage the reliability of jetliners.

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