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

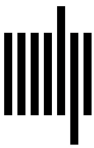

The MIT Press

1 CATASTROPHIC TECHNOLOGIES: THE RISE OF RELIABILITY AS A VARIABLE OF CONSEQUENCE

. . . by slight ligaments are we bound to prosperity and ruin.
—Mary Shelley

1.1 INTRODUCTION

A NEAR-MISS

From halfway around the world, and with the passage of time, the 2011 Fukushima Daiichi nuclear disaster can appear almost unremarkable—just one more entry in the sad ledger of technological failures. This is understandable. The plant's three reactor meltdowns were dramatic while they lasted, certainly, but the world did not stop in 2011, and even in most areas of Japan, daily life rebounded relatively quickly. But this apparent normalcy is deceptive. Understood properly, Fukushima—as I will henceforth refer to the accident—was a near-incomprehensible catastrophe. It should have been a wake-up call to a world that has stumbled blindly into an increasingly dangerous relationship with its technological creations.

Appreciating Fukushima's full significance is difficult, as many of its ramifications remain opaque. Estimates of its costs—human, financial, and ecological—have all swollen considerably in the years since 2011,[1] but in slow and incremental steps that have rendered the rising numbers and their implications all but invisible (Downer 2016).[2] Revised pollution estimates and reassessed expenses rarely sell newspapers, especially when those estimates are unusually complex or contested like Fukushima's. Media coverage

of the disaster has been a disorienting crossfire of claims and counterclaims—a product of the esoteric issues involved and the powerful interests at stake (see, e.g., Hamblin 2007). Even relatively uncontested claims can be rhetorically impotent in this context; the numbers involved are either too large to be relatable—millions and billions occupying a similar space in the collective consciousness, (a phenomenon that psychologists sometimes refer to as "scope neglect" [e.g., Kahneman 2000])—or too abstract to be meaningful. (No measure of radiological pollution makes for an intuitive headline, whether expressed in "rads," "rems," "roentgens," "becquerels," "grays," "sieverts," "curies" or even "bananas.")[3]

To cut through this complexity and contestation around Fukushima, some commenters try to convey the accident's true gravity by focusing on a single, authoritative fact with intuitive significance. For Naoto Kan, the Japanese prime minister at the time of the crisis, that fact was that Japan came bracingly close to losing its capital city.

Proponents of atomic energy robustly dismiss claims that Fukushima threatened Tokyo as hyperbole. Japan's capital is about 150 miles (241 kilometers) from the site of the accident, and few experts contend that the meltdowns could ever have released enough pollution to jeopardize its residents. However, in Kan's telling, and that of many other experts, the meltdowns were not the main event.

Kan's account of how the city narrowly escaped catastrophe centers instead on one of the plant's spent-fuel pools. The pool in question, one of several at the site, was housed above an inoperative reactor in unit 4 of the plant. Resembling a 45-foot (14-meter)–deep swimming pool, its purpose was to store radioactive materials safely. At the time of the accident, it contained years' worth of new and spent reactor fuel: 1,535 assemblies, each consisting between fifty and seventy individual rods. These rods had to be kept submerged in cold water lest they combust spontaneously. Under normal circumstances, this would not have been challenging. Four days into the crisis, however, a large explosion in the reactor building severely damaged the pool and its supporting structures. At this point, the pool began to leak, and keeping the rods submerged in water became extremely challenging.

Plant workers battled desperately to refill the water faster than it drained, but the circumstances were punishing. Desperately understaffed because most of their colleagues had fled the site's soaring radiation, and woefully

underequipped, they were simultaneously managing three concurrent meltdowns amid the destruction and disruption of a historic tsunami and earthquake.

The stakes, however, could not have been higher. Had the pool's fuel rods been exposed for long, they would have ignited and burned unquenchably, liberating an unprecedented volume of radionuclides into the atmosphere. (Among other radioactive materials, the pool contained ten times the volume of cesium-137 released by the Chernobyl accident in the mid-1980s.) This would have forced a full evacuation of the plant, which in turn would have led to the loss of its other pools, some of which were considerably larger. The hazards of this "devil's chain reaction," as one commenter would later refer to it, would have been far more extreme than those of the meltdowns alone. Extreme enough to have very plausibly necessitated the evacuation of Tokyo City (Gilligan 2016; Sieg and Kubota 2012; Lochbaum, Lyman, and Stranahan 2014, 80–85; Osnos 2011; Dvorak 2012a, 2012b; Matsumura 2012; Goodspeed 2012; Fackler 2012; Lean 2012; Kan 2017).

Industry advocates vigorously dispute this conclusion, but it is supported in substance by an extensive range of highly credible sources. It is now clear, for instance, that the US Nuclear Regulatory Commission (NRC) believed such a scenario to be possible, and at points even probable. Indeed, its chairman later testified to Congress that he believed it had come to pass. An interagency team was convened to plan the emergency evacuation of US citizens from Japan's capital (Lochbaum et al. 2014; Osnos 2011). The official, independent report on the disaster prepared for the National Diet of Japan (NAIIC 2012) is clear about the danger. Its release was accompanied by a flood of chilling testimony. "We barely avoided the worst-case scenario, though the public didn't know it at the time," the chairman of the inquiry told reporters (Lean 2012). "It was extreme luck that Japan managed to avoid experiencing the most disastrous day" said another prominent member (in Lean 2012).

If "extreme luck" seems like perverse commentary on a triple reactor meltdown amid a natural disaster of historic severity, then consider the implications of Japan's unrealized "disastrous day." The gravity of Japan losing Tokyo is difficult to capture in a language that cheerfully wastes its superlatives on breakfast cereals and laundry detergents, but it would undoubtedly be immense. Cities are surprisingly resilient, but radiological pollution

is insidious and difficult to clear. Japan's capital, which once rose phoenixlike from the ashes of horrific US fire bombings of World War II, might have been felled much more permanently by its government's own energy policy; Fukushima's fallout forever poisoning the earth, like the salt that Rome plowed into the soils of Carthage.

Even a temporary evacuation—which the Japanese government began planning in earnest (Gilligan 2016)—would have taken a momentous human and economic toll. Tokyo is a gigantic city, home to over thirty-five million people (more than New York and London combined). Like those cities, it is a command center of the global economy, and it has the highest gross domestic product (GDP) of any metropolis on the planet. In a world where even mispriced mortgage derivatives can instigate a financial meltdown and "Great Recession," the potential financial fallout of such an event almost defies comprehension. At the very least, it would have jeopardized Japan's ability to service its national debt. Kan might not have been exaggerating much when he told the *Wall Street Journal* that his country's "existence as a sovereign nation was at stake" (Quintana 2012).

So it is that Fukushima—already remembered as the signal technological disaster of recent decades—is arguably best understood as a near-miss. Had the explosion in unit 4 been even slightly larger, had individual plant workers been less courageous, or had the area been hit by a significant after-tremor (which are common in the wake of earthquakes), then the pool would likely have been lost, with all the consequences that loss implies. Experts can debate the likelihood of this scenario, but probable or not, the fact that it was even *possible* should cast a long shadow over humanity's relationship to technical ambition in the twenty-first century.

BETTING ON TECHNOLOGY

There are many lessons to be drawn from Fukushima, but perhaps the most fundamental of them pertains to the authority afforded to engineering safety assessments and the purveyors of those assessments. Almost unknowingly, Japan bet its future on the understanding that engineers could speak definitively about extraordinary failure probabilities in a complex technological system, and it almost lost everything.

At some point in our recent history, nations started building technologies that simply could not be allowed to fail catastrophically—technologies with the latent capacity to acutely imperil lives, livelihoods, and lifestyles on a

massive, unconscionable scale. It is important that we govern such technologies wisely. This, in turn, requires that we understand those technologies. Or, perhaps more meaningfully, it requires that we understand the nature of our relationship to them: the capabilities and limitations of the expert bodies that we charge with overseeing them, along with the provenance and credibility of the safety assurances that those bodies provide.

These capabilities and assurances, as well as their fundamental limitations, are the central theme of this book.

1.2 CATASTROPHIC TECHNOLOGIES

A NEW RELATIONSHIP TO FAILURE

Catastrophic technological failure has been a meaningful public concern for as long as there have been engineered structures. In December 1879, outside Dundee in Scotland, the Tay Rail Bridge collapsed during an evening storm, dropping a train full of passengers into the icy waters of the Firth of Forth. In March 1864, outside Sheffield in England, the newly constructed Dale Dyke Dam burst as its reservoir was being filled; the ensuing flood damaged over 600 houses and killed over 240 people. In January 1919, a huge molasses storage tank in Boston erupted at its seams, unleashing a sugary, 35-mile-per-hour tidal wave that fatally engulfed 21 people. There are many such stories, some dating back many hundreds of years. The Koran speaks of the failure of the "Ma'rib dam" in the year 570 or 575, which, historians believe, could have displaced upward of 50,000 people from what is now Yemen.

Despite this history, however, there remains a sense in which modernity's relationship to technological failure changed meaningfully in the mid-twentieth century. As it is only at this juncture that advanced industrial societies began to build complex, dynamic sociotechnical systems—as opposed to static structures—that could almost never be allowed to fail.

I propose to call such systems "catastrophic technologies."

Let us define catastrophic technologies as complex technological systems that—because of the potential hazards (mortal, social, economic, or environmental) of them failing catastrophically (i.e., in a way that leads to a major accident)—would not be viable (economically or politically) unless the probability of such failures was deemed low enough to justify excluding them from all public discourse and decision-making. They are complex systems that require extraordinary, and historically unprecedented, failure

rates—of the order of hundreds of millions, or even billions, of operational hours between catastrophic failures.

This performance requirement is the defining feature of catastrophic technologies—even more than the extreme failure-hazards from which that requirement stems. In most cases, it must be known in advance of these technologies being allowed to operate. At minimum, it must be assumed as a premise of their operation: implicit in discourse and decision-making that never even consider catastrophic failure as a possibility. Every debate and decision regarding catastrophic technologies is grounded in implicit or explicit assumptions about failure being, for all intents and purposes, a solved problem. States would not tolerate such technologies unless they considered catastrophic failures to be functionally impossible, or close enough to impossible as to be beneath serious policy consideration.[4]

Technologies requiring such extreme failure behavior arose as the product of two interrelated trends. The first, straightforwardly, was the emergence of complex systems with the potential to fail in ways that were so unprecedentedly hazardous or costly that even a single, isolated failure might lead to intolerable harm. Reactors exemplify this phenomenon, as Fukushima amply illustrates. The second, slightly more complicated, trend was a sharp rise in the adoption of complex systems that could tolerably be allowed to fail (unlike reactors), but only so long as those failures remained extremely infrequent. Jetliners are exemplary here. Jetliner crashes are tragic, no doubt, but their isolated costs are not equivalent to those of reactor meltdowns. So in the early days of civil aviation—when infrequent flights limited the absolute number of accidents, and smaller aircraft limited the absolute number of people who could die in each accident—the socially tolerable number of accidents per departure could be (astonishingly) high by modern standards. At this time, airliners did not require the extreme mean time to failure of a catastrophic technology. As jetliners emerged and air travel became much more democratic, however, the relative frequency of accidents had to fall precipitously for air travel to remain viable.[5] There were 51 fatal commercial airplane accidents in 1929, roughly one for every million miles flown. This was deemed acceptable at the time, but the same rate today would imply over 7,000 fatal accidents per year: just short of 20 every day (and with many more passengers on each flight). As the absolute number of flights rose, therefore, the failure requirements on airliners became ever more demanding, to the point where they emerged as catastrophic technologies.

This is all to say that that, over roughly the same period (from the 1950s to the 1970s), reactors and jetliners both came to require extraordinary levels of failure performance—making them catastrophic technologies—but for slightly different reasons. Reactor meltdowns were (and remain) far more consequential than jetliner crashes, so they needed to be commensurately less probable. No state would tolerate a Fukushima-scale accident on its soil in exchange for atomic energy, but all would accept a tiny number of jetliner crashes in exchange for mass air travel. With far fewer reactors in operation than jetliners, however, both systems needed similar failure probabilities to satisfy these different requirements. The absolute number of plane crashes it would take to jeopardize the political viability of jetliners is clearly far higher than the absolute number of meltdowns it would take to jeopardize the political viability of reactors, in other words, but we require broadly equivalent levels of failure performance from both technologies because there are many more jetliners than there are reactors.[6]

(In light of the distinction made here, and for reasons that will become apparent as the argument progresses, it is useful to divide catastrophic technologies into two types on the basis of why they require their extreme failure performance. Let us call these types "chronic" and "acute." Where chronic catastrophic technologies, exemplified by jetliners, are those systems that can tolerably be allowed to fail on rare occasions but require their extreme failure performance because of the volume at which they operate. And acute catastrophic technologies, exemplified by reactors but representing the majority of all catastrophic technologies, are those that operate at much smaller volumes but require extreme failure performance because *any* failures are intolerable. Both chronic and acute have equivalent failure requirements, we might say, but meaningfully different relationships to individual failures.)

Catastrophic technologies are still quite rare, although this is changing. Reactors and jetliners were early and exemplary examples, but in the years since World War II, modern societies have become increasingly dependent on the near-infallible functioning of various complex technological systems (Perrow 2007). Certain industrial plants, laboratories, and drilling platforms require commensurately extreme failure performance, for instance, as do a significant number of slightly less sociotechnical systems, such as undersea communications cables and some medical devices. Atomic weapons pose exquisitely acute demands in this regard, albeit in a meaningfully different political and organizational setting that—in the context of this argument, at

least—makes them less exemplary than reactors and jetliners.[7] And, insofar as we wish to construe networks as technologies in themselves, as we almost certainly should, then so do their sprawling command-and-control infrastructures, the failure of which could instigate atomic wars (and almost have on multiple occasions [Sagan 1993; Schlosser 2013]). Along the same lines, we might also recognize a range of other networks: from banking computers and electricity grids to air-traffic control and global positioning system (GPS) networks. The internet, in particular, increasingly represents both a catastrophic technology in itself and an acute new source of vulnerability in many other systems (Kaplan 2016). Broadening the definition of "technology" in a different direction (and stretching the terminology of "failure"), a wide range of nonmechanical technologies have similar catastrophic potential, and thus require equivalent assurances about their behavior: some novel biological agents (Downer 2020), for instance, and financial instruments like the derivatives that played a crucial role in the 2007–2008 financial crisis (MacKenzie 2005).

As with most social scientific categories, however, the term "catastrophic technology" defies precise enumeration. "Catastrophic-ness" is neither a binary nor a wholly objective condition, so the exact tally of qualifying technologies will always be contestable. The examples given in this book are in no way intended to represent a comprehensive list, therefore. It is reasonable to assume that others are woven invisibly into the fabric of modern life, and that their numbers will expand in the coming years. (Consider, for instance, autonomous vehicles [Stilgoe 2018], artificial intelligence systems [Bostrom 2014], or even Moon habitats [Benaroya 2018], all of which are under active development.) The picture is further muddied by domain-specific terminologies and preconceptions. (We rarely speak about biological or economic disasters in terms of "systems failure," for example.) For all the ambiguity and heterogeneity of catastrophic technology as a category, however, I would argue that the mutual need for extreme failure performance that defines it creates enough meaningful commonalities between different systems—in their governmentality, and (especially) their epistemology—for the idea to be useful.

Catastrophic technologies have common properties, therefore, but exploring those properties requires depth and detail, which in turn demand focus. For this reason, the argument that follows will focus on a narrow range of electromechanical systems: primarily jetliners and reactors, with most of

the emphasis on the former. The logic for this was outlined in the preface and will become more evident as the argument progresses. Broadly, however, these two specific technologies have been chosen for three reasons. The first is that neither is an ambiguous or borderline example of a catastrophic technology; both exemplify the category and its essential qualities. The second is that the structures—organizations, laws, rules, practices, terminologies, norms—through which both are governed are similar enough to make comparisons intuitive. And the third is that, despite their ostensible similarities, the underlying constraints on each technology are distinct enough to make comparisons useful. They might be governed via ostensibly similar structures, but they are subject to very different epistemological limitations, and this fundamentally alters the way those structures must operate. So it is, I will argue, that contrasting the reliability practices around jetliners with those around reactors offers revelatory and widely generalizable insights into the governance of all catastrophic technologies.

Let us turn, then, to examine the structures through reactors and jetliners are governed.

1.3 GOVERNING CATASTROPHIC TECHNOLOGIES

THE PRIMACY OF RELIABILITY

In the US and beyond, the catastrophic technological era has given rise to a distinctive form of governmentality, the properties and priorities of which reflect the unique demands of extreme failure hazards. Most straightforwardly and intuitively, this governmentality is characterized by an unusual emphasis on risk management: risk being a far more prominent concern in the governance of jetliners and reactors than in that of toaster ovens and televisions. Distinctively, however, risk management in this context has come to be dominated by concerns about *failure* risks, which tend to eclipse risks associated with systems' normal operation. This is simply to note, for instance, that the safety of reactors or jetliners is conventionally understood almost exclusively in terms of the probability of them failing catastrophically, even though there are ways for them to be unsafe without failing (by polluting, for instance) (see Rijpma 1997; Wolf 2001).

The convention of construing risk exclusively in terms of failure is made more distinctive by a parallel convention of construing failure exclusively in terms of probability. In most engineering contexts, failure risks are understood

as a function of both their probability and their consequences. (Such that engineers might improve the safety of systems by making failures less hazardous when they occur, as well as less likely to occur in the first place.) In catastrophic technological contexts, however, the potential hazards of failures are all but intolerable, so the probability of failure is king. (So it is, for instance, that we address the risk of accidental nuclear detonations by ensuring that they never occur, not by digging bunkers and relocating populations.) The potential consequences of failure are rarely forgotten entirely in this context—as we will see, for instance, states mandate crash survivability measures in jetliners—but such considerations are always secondary and attenuated. This is all to say that there is little talk of "resilience" in the discourse around catastrophic technologies, as the bureaucracies responsible for managing them place far greater weight on preventing failures than they do on mitigating their hazards. If technological artifacts can be understood as responses to implicit problems, as Baxandall (1985) suggests, then the key problem to which catastrophic technologies are a response is how to avoid accidents, not how to survive them.

In electromechanical contexts, these all-important failure probabilities are usually metricized as "reliability," defined narrowly here as the frequency of catastrophic failures (as opposed to, for instance, the frequency and duration of unplanned downtime), and usually expressed as a mean-time-to-failure. For instance, the viability of both jetliners and reactors hinges on them exhibiting a known mean-time-to-failure north of hundreds of millions of hours (tens of thousands of years), or what I will henceforth refer to as "ultra-high reliability."

Again, this convention of emphasizing reliability is more distinctive and less inevitable than might be intuitive, not least because electromechanical systems can fail catastrophically for reasons that would not traditionally be construed as "reliability": sabotage, for example, or human error. Such considerations are rarely forgotten entirely. Indeed, states assiduously manage issues like security or human performance in reactors and jetliners. Yet high-level discussions of these systems' safety are predominantly framed as engineering problems, wherein reliability measures are widely treated as expressions of the absolute likelihood of catastrophic failure from any cause. "Soft" difficult-to-measure risks, such as might arise from security or human-performance questions, are usually folded into these reliability metrics, either by treating them as solvable problems, which functionally disappear once certain requirements

have been satisfied, or by treating their probabilities as objectively quantifiable, akin to those of mechanical failure.[8]

Here, then, are four noteworthy characteristics of the modern (arguably Western) approach to governing catastrophic technologies:

1. It places a strong emphasis on controlling their risks (more often referred to as "safety").
2. It predominantly construes their risks in terms of their failure behaviors.
3. It largely understands their failure behaviors in terms of probability alone (addressing the likelihood of failures rather than their consequences).
4. It tends to reduce the probability of them failing to a reliability metric.

Later chapters of this book will explore these characteristics in greater detail and examine some of their implications. For now, however, it suffices to recognize the significance of reliability as a metric by which modern states measure catastrophic technologies, and through which they approach the governance of those technologies. Directly or indirectly, almost all public deliberations about catastrophic technologies hinge on expert reliability calculations.

So it is that the catastrophic technological era has elevated technological reliability to a variable of enormous consequence. There is undoubtedly a degree of elision in this elevation. Reliability does not capture every dimension of failure, failure does not capture every source of risk, and risk is not the only meaningful measure of a catastrophic technology. But even if reliability cannot encompass everything—or even everything that is ascribed to it—the emphasis afforded to it is not wholly inappropriate. Safety is the sine qua non of catastrophic technologies, and reliability is a necessary component of that safety. Absent at least an implicit understanding of reliability, any other claim about such technologies becomes moot. Reactors and jetliners can only be "economical" or "environmental" to the degree that they do not explode, melt down, or fall repeatedly from the sky.

OBJECTIVE AND AUTHORITATIVE

For all its newfound prominence, the reliability of a complex system—especially at the ultrahigh levels required of catastrophic technologies—is far from being a transparent variable. One cannot simply examine the reliability of a system as one would its weight, volume, or length. Establishing reliabilities at such levels, making them visible and auditable, is an elaborate process,

fraught with difficulties (as chapter 2 will explain in detail), and requiring extensive expertise. So it was that the emergence of reliability as a variable of political consequence was accompanied by the parallel emergence of expert intermediaries charged with making it publicly accountable. (Sims [1999] calls such bodies "marginal groups" to connote the fact that that they inhabit more than one social world, translating one to the other and performing the boundary work of differentiating credible from untrustworthy knowledge.)

Along with atomic energy and mass air travel, therefore, the mid-twentieth century saw the birth of elaborate reliability-accountancy structures: organizations, metrics, rules, guidance, and practices, all designed for measuring and ensuring ultrahigh reliability. The modern incarnations of these structures have many national and domain-specific idiosyncrasies, but most have substantial underlying commonalities. In the US, for instance, their contours are broadly equivalent across the civil aviation and civil nuclear domains and have close analogs in non-US contexts. Their most visible manifestations are dedicated design assessment and oversight bodies, usually prominent subdivisions of larger regulatory agencies, such as the Federal Aviation Administration (FAA), which oversees US civil aviation, and the NRC, which oversees US atomic energy.

These technology oversight bodies wield more power than is often apparent. As scholars of audit processes have long attested, experts and their calculative practices become highly influential when an important property is knowable only through its assessment (Hopwood and Miller 1994; Power 1997), and reliability is no exception in this context. The inherent inscrutability of extreme reliability, together with its centrality to catastrophic technological discourse, imbue the bodies that account for it with considerable agency, which is bolstered by the fact that audiences conventionally construe the findings of these bodies as being highly authoritative: the product of precise, objective, and rule-governed processes. Publics and policymakers routinely doubt and question the findings of expert economists or sociologists, but once the NRC or FAA officially measures the reliability of a specific reactor or jetliner, then a broad range of institutional actors, from courts to budgetary planning offices, are all but obliged to treat that measure as an established fact.

The convention of treating catastrophic technological reliability assessments as objective facts should not be surprising. It is firmly in keeping with a pervasive cosmology of technoscientific knowledge, wherein the

work of interrogating machines is routinely assumed to be driven by strict methodologies, which, when performed diligently and correctly, yield knowably correct results (Rip 1985; Jasanoff 1986; Wynne 1988; Mitcham 1994; Ezrahi 2008). Modern bureaucracies, we might say, operate within a positivist cosmology wherein catastrophic technological black boxes have knowable properties, the truth of which can be established definitively via the checkboxes of formal audit practices.

This positivist cosmology, with its associated certainties, is far from unique. Indeed, it can be found wherever public policy intersects with the properties of technological artifacts. When applied to the reliability of catastrophic technologies, however, it becomes especially crucial. For here, more than anywhere, it is vital that expert determinations be construed as impersonal, authoritative, and rule-governed. Where the fates of cities are at stake, assertions of technological reliability need to stand on more than the considered opinion of industry insiders. On such questions, if on few others, certainty and objectivity are nonnegotiable.

It should be troubling, therefore, that in this specific context—where reliability claims are made about catastrophic technologies—there are compelling reasons to believe that the positivist cosmology of engineering knowledge is uniquely misleading.

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