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

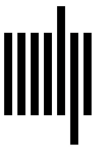

The MIT Press

INTRODUCTION: PURPOSE, SCOPE, AND STRUCTURE

> We live in a society exquisitely dependent on science and technology, in which hardly anyone knows anything about science and technology.
>
> —Carl Sagan

PURPOSE

What's wrong with knowing what you know now and not knowing what you don't know now until later? It's a meditative question, artfully phrased by A. A. Milne for Winnie the Pooh. But the answer—at least if you're designing a complex technological system with the potential to fail catastrophically—is "everything." If somewhere, deep in the bowels of your sprawling atomic weapons arsenal and infrastructure, there lurks an unknown software bug (or an unanticipated corrosion process, or anything else) capable of accidentally instigating a thermonuclear war, that's something you need to know now. Because finding it out later, the hard way, could lead to you having what Milne would probably have called "a bit of a day."

This should be a sobering thought, even for stoics like Pooh. With ever more confidence and enthusiasm we have been building technologies that—like our nuclear weapons infrastructures, and for essentially the same reasons—need to be known with great certainty *now* rather than *later*. These are technologies that we absolutely cannot allow to fail catastrophically, lest we incur intolerable hazards. The only way to be certain that such technologies will not fail, however, is to deeply understand every aspect of their functioning, and there are very convincing epistemological reasons

to believe that this depth of understanding is fundamentally unachievable. The tools that experts would use to interrogate them—the tests, models and calculations by which we know machines—are all inescapably imperfect, and it is rational to imagine that this imperfection should, very occasionally but still too frequently, give rise to unavoidable catastrophes. If knowing the limits of our knowledge is itself a form of wisdom, in other words, then there is a strong argument to be made that we are pursuing some of our technological ambitions in folly: trusting our lives and livelihoods to an implausible promise of technological reliability.[1]

The significance of this conclusion is obvious, but it immediately raises a difficult question because experts demonstrably *have* mastered extreme reliability in at least one highly complex technology: jetliners. In the civil aviation sphere, if in no other, they appear to have transcended the imperfections of their tools, and the limits of their knowledge, by building intricate and demanding machines with failure rates that are as low as we require of our most hazardous systems. We can know this with confidence because, unlike most technologies that require extreme reliability, we operate jetliners in high enough volume to measure their failure rates statistically. And although they do still fail, very infrequently, these failures are diluted by their numbers; such that equivalent failure-rates would imply near perfect reliability in systems like reactors, which we operate in much smaller numbers.

This "aviation paradox," as I will call it, is the problem that drives the argument that follows. The extraordinary reliability of jetliners, together with the accurate predictions of that reliability made by aviation regulators, speaks to an astonishing depth and breadth of understanding of the myriad intricacies of jetliner design and operation. An understanding that is so deep and so broad, I will argue, that it ought to be incommensurable with what philosophers have told us about the limits of engineering knowledge. Civil aviation's achievements in this regard raise fascinating epistemological questions, therefore, and the answers to the questions have important implications regarding our relationship to hazardous technologies more broadly. As we will see, jetliners are not reliable for the reasons we are led to believe, and the way experts achieved this reliability is less transferable than we imagine.

In exploring this question, the chapters that follow introduce a novel way of thinking about technological accidents. They make the case that some accidents—which I will call "rational accidents" in this book (but have elsewhere referred to as "epistemic accidents" [Downer 2011b; 2020])—are

most usefully be attributed to the fundamental limits of what engineers can know about the technologies they build. Such accidents might not be common, I will argue, but they have several noteworthy properties. Much like Perrow's (1999 [1984]) well-known "normal" or "system" accidents, but for meaningfully different reasons, they cannot be avoided, even with ideal organizational practices and engineering analyses. And insofar as they can be considered "blameworthy" at all, it is only in the sense that experts should have understood the limits of their control. Unlike normal accidents, however, they can be managed over time: incrementally reduced, almost (but never entirely) to the point of being eliminated. But only if the organizations managing them have access to certain resources and are willing and able to follow certain practices, the costs of which can be prohibitive.

The argument explores the implications of these rational accidents. It outlines their significance and logical necessity, it examines the ways in which experts have learned to manage them in civil aviation, and it considers the transferability of their achievement to other technological domains. More generally, however, it explores the epistemological problems posed by extreme reliability in complex systems. And with these goals mind, there are a few ways in which its scope is perhaps broader, and simultaneously narrower, than might immediately be apparent. Before proceeding further, therefore, two points are worth clarifying in this regard: one regarding the argument's theoretical ambitions, the other its practical relevance.

SCOPE

THEORETICAL AMBITIONS The argument that follows is probably narrower with regard to its theoretical ambitions than might be intuitive. In light of this, it is useful to briefly clarify its relationship to two distinct academic literatures: the Science and Technology Studies (STS) literature, and the wider social-scientific literature around technological safety and disaster.

The text has a slightly unorthodox relationship to the STS literature. It is firmly grounded in the STS tradition in the sense that it builds on the epistemological limitations of seemingly objective knowledge and explores the subjectivities that arise from those limitations. Where a more conventional STS account might invoke those subjectivities to historicize and contextualize the knowledge claims it is discussing, however, the emphasis here is more on the practical difficulties and dangers to which they lead. The

intention, in essence, is to explore how (and if) expert organizations successfully navigate the subjectivities of extreme technological reliability; it is not to explain how they came to understand and/or manage extreme reliability in the manner that they do. The latter questions are undoubtedly worthwhile, and their answers make for fascinating stories. The structures and practices through which modern societies govern the reliability of extremely hazardous technologies are more contingent, contested, and socially constructed than we commonly realize (see, e.g., Jones-Imhotep 2017; Wellock 2021; Johnson 2002). They are just not the stories being told here. The STS scholarship that this book draws on most directly, and hopefully most closely resembles (at least in spirit), are texts like Mackenzie (1990; 1996a, b; 2001), Wynne (1988), Collins (1988), Pinch (1991; 1993), and Collins and Pinch (1998); all of which, in different ways, explore the indeterminacies of engineering knowledge and its social implications.

The text has a more complex relationship to the social scientific literature around technological safety and disaster. In exploring the epistemology of extreme reliability, it does speak directly to some core concerns of this literature—the question of what makes technologies safe, for instance, and of why they sometimes fail—but its ambitions in this regard are nevertheless limited. By this, I mean that I am not attempting to grapple with every aspect of what makes technologies with catastrophic potential safe.

Decades of social research into this question has shown that safety is always a holistic accomplishment requiring constant work on a wide range of fronts: as much a function of organization and culture as of engineering epistemology. And it has shown that technological accidents have equally layered and complex underpinnings. Silbey (2009) offers a useful overview of this literature, but Vaughan's (2021) exploration of air traffic controllers, which shines a light on some of the nuanced practices underpinning safe air travel, is exemplary and pertinent (see also, e.g., Perrow 1999; Reason 2016; Dekker 2011; Schulman 1993; Perin 2005; Snook 2000; Vaughan 1996; Turner 1978; Sagan 1993; Weick and Sutcliffe 2001). Seen from the perspective of this literature, the decision to examine technological safety and disaster through the prism of engineering epistemology, as this volume does, might appear reductive.

It is important to understand, therefore, that the account that follows very intentionally explores only a narrow, albeit vital, dimension of technological safety: the question of how it is even *possible* for experts to know

complex systems well enough to achieve and verify of them very extreme reliabilities. The book unpacks this question; it looks at how engineers have managed such tasks in a specific domain, and it considers the wider implications of their success. That is all. It is still a lot, however, in the sense that the narrowness of the central question does not preclude some strong conclusions with far-reaching implications. Ultrareliable design will never be a *sufficient* condition of technological safety, but in the context of technologies that absolutely cannot be allowed to fail it is a *necessary* condition. And while failures attributable to the limits of engineering knowledge cannot sufficiently explain all accidents, they can explain some accidents. (Note that safety and disaster are asymmetrical, in that the former requires everything to go right, while the latter requires only that one thing go fatally wrong.) The many organizational and cultural dimensions of building, maintaining, and operating complex technological systems might be essential to their safety, therefore, but if we cannot also master the epistemological questions then they will never be enough.

Even if the question being explored here represents only one piece of the safety problem, therefore, it is nevertheless a vital piece, at least in certain contexts. And on this basis, I will speak often about the safety of some technologies being "dependent" on specific conditions and practices that allow experts to push beyond the limitations of their tests and models. Where I do so, however, the intention is never to imply that those practices or conditions are *sufficient* for achieving safety. Researchers seeking a more rounded and expansive understanding of how technological safety is accomplished will need to read more widely.

With that caveat, therefore, this volume might reasonably be counted amid the wider social-scientific scholarship on technological safety and disaster—enough so, at least, that it is probably instructive to locate it in relation to that literature. Such positioning is invariably fraught; not least because it is difficult to map the academic discourse on any complex topic without distorting that discourse. Scholars have nuanced understandings of their work's distinctiveness, and understandably chafe at being lumped together into rough-hewn categories. Navigation requires some systemization, however, and even the most flawed maps can be useful. So, at the risk of reifying some contested categories, I think it is enlightening to understand the argument in relation to an old but influential schema that divides the academic safety literature along an axis that speaks to some of its core themes.

This schema originated with Sagan (1993), who proposed that social scientific scholarship on technological safety might usefully be divided into two broad schools, based, in essence, on whether they hold that technologies can ever be made perfectly safe. On one side, Sagan (1993: 13) posited an "optimistic" school, which construed safety as fundamentally perfectible. This school is exemplified, in his telling, by what he calls the Berkeley group (e.g., La Porte and Consolini 1991; Roberts 1989; Rochlin et al. 1987; Schulman 1993; Rochlin et al. 1987), but it also encompasses outliers like Wildavsky (1988), and today would almost certainly include works such as Vaughan (2021), as well as the burgeoning (largely European) "safety science" literature (e.g., Le Coze 2020). On the other side, he posited a "pessimistic" school, which construed safety as fundamentally imperfectable (Sagan 1993: 13). This school is exemplified in his account by Perrow (1999 [1988]), but it also includes texts such as Clarke (1989), Shivastava (1987), and Reason (1990), and today would probably also include works like Vaughan (1996), Snook (2000), as well as Sagan (1993) itself.[2] In the form that Sagan's schema is conventionally remembered (and largely invoked by Sagan himself), however, the different schools are usually narrowed to just his exemplars—the Berkeley group and Perrow—whose arguments are cast as opposing theories of disaster: Berkeley's optimistic High-Reliability Theory (HRT) being set against Perrow's pessimistic Normal Accident Theory (NAT).

Sagan's schema is overreductive, to be sure, especially in its narrowed form. It ascribes an unwarranted degree of coherence to heterogeneous bodies of literature, and a lot of research fits imperfectly into his categories. (It is notable that many of the researchers he describes rejected his interpretation of their position on perfectibility; most notably his optimists [e.g., LaPorte 1994; LaPorte and Rochlin 1994]). For all its limitations, however, his construal of the literature was undeniably influential. In part because his book—compellingly written, with eye-opening accounts of close calls with accidental atomic war—was revelatory and persuasive. But also, and more substantially, because his categories did seem to capture a real, tangible divide in the scholarship itself; a divide that still has relevance today.

To appreciate the lasting value in Sagan's schema, it helps to think of his schools as divided less by the question of perfectibility, which Sagan foregrounds, and more by whether they are primarily oriented towards researching safety or failure. On one side, we might say, is scholarship that takes, as

its primary objective, the question of what makes complex systems safe, and seeks to identify the characteristics of safe systems. This is Sagan's optimistic literature, exemplified in his account by HRT. And on the other side there is scholarship that takes, as its primary objective, the question of why complex systems fail, and seeks to identify the characteristics of unsafe systems. This is Sagan's pessimistic literature, exemplified by Perrow and his normal accidents.

These two projects heavily overlap, of course. "Safety" in this context can almost be defined as the absence of failure, so the difference between researching one and researching the other is often a matter of perspective. Perspectives are important, however, and it is fair to say that the differently-oriented research questions tend to shine meaningfully different lights on the same phenomena. They encourage dissimilar questions, assumptions and interpretations, and, as a result, they tend to yield different insights, opinions and conclusions. So it is, for example, that scholars exploring safety—Sagan's optimists—tend to see the potential for organizations to ameliorate human frailties, counteract economic pressures, and learn from experience. Whereas those exploring failure—Sagan's pessimists—are more likely to see the potential for organizations to be undermined by human frailties, perverse incentives, and imperfect learning. Where one side might portray a "glass of safety" as "99 percent full," as Sagan (1993: 48) puts it, the other is more likely to portray it as "1 percent empty."

This difference in perspective, when extended to Sagan's core question of the perfectibility of safety, gives rise to contrasting pictures, even if it can be difficult to fix points of fundamental disagreement. (When pushed, both sides tend to agree that organizations are capable of developing impressive stratagems for safety that are always imperfect at the margins, and that both the impressiveness and imperfections are worthy of scrutiny.) As Sagan intuited, the distinction becomes most acute in the context of Perrow (1999 [1988]) specifically. This is because Perrow, almost uniquely, makes a sustained and principled argument about the inherent unavoidability of certain kinds of accidents, and then explores the implications of this inevitability. (These are his "normal accidents," the logic of which chapter 7 will explain in more detail.) It is also because Perrow, like Sagan himself, focuses much of his analysis on a technology with such extreme failure hazards that even a 1 percent safety deficit takes on enormous significance (reactors in Perrow's case, and

atomic weapons in Sagan's). It is still difficult to find points of fundamental disagreement. Perrow is not claiming that all, or even most, accidents are unavoidable, and safety-oriented scholars are not explicitly claiming the opposite. At the same time, however, it is rare for the latter to highlight or unpack the implausibility of perfect safety, or to grapple with its implications for extremely hazardous technologies.[3] As with the engineering discourse around such systems, most safety scholarship frames technological failure as a *problem to be solved* rather than an *inevitability to be confronted*.[4]

Sagan's portrayal of the literature offers an instructive backdrop against which to understand the argument of this book, not because it easily accommodates that argument, but because it doesn't. The argument that follows addresses many of the core issues he identifies—the evitability of accidents; the perfectibility of extreme safety; the potential for organizational learning—but in ways that cut squarely across his basic divide.

It would be easy to pigeonhole this volume as belonging squarely in Sagan's disaster-oriented, pessimist, tradition. The title is an allusion to Perrow (1999 [1988]), after all, and there are several clear parallels with that text. It is primarily focused on technologies with catastrophic potential, for instance, wherein even tiny shortfalls in the "glass of safety" become extremely important. And the opening two parts (seven chapters) examine why experts inevitably struggle to fill that glass. As with Perrow's text, moreover, this examination cumulates in the proposition of a new type of accident—different from but analogous to his "normal accident"—that is fundamentally unavoidable, in the sense that even perfect organizations (if they existed) would have no way of preventing it. More broadly, the argument also reckons with the limits of organizational rationality: emphasizing the power of incentive structures to shape practices and priorities, even at the expense of safety.

For all this, however, a case could also be made that much of the text belongs in the safety-oriented, optimist, side of Sagan's ledger. This is because the argument effectively switches orientation in its second half, moving from examining the causes of failure to exploring the achievement of safety. This is to say that, having explored the epistemological dilemmas of extreme reliability, it then turns to explore how civil aviation manages to transcend these limits. It argues that the accidents to which those limits give rise are only fundamentally unavoidable in certain circumstances, and explains how those accidents—if tolerated and interrogated over long periods—can

serve as a foundation on which to build extreme reliability. Then, in chapters that might almost be considered High Reliability Theory, it unpacks the specific practices and structures that, with appropriate incentives, enable organizations to leverage that foundation.

Perhaps unsurprisingly, therefore, the conclusion sets a slightly ambiguous tone, especially with regard to the question of how optimistic we should be about technological safety. In the vast majority of technological contexts, I see no reason for undue pessimism on this front. As with Perrow, I am not suggesting that all (or even most) technological disasters are unavoidable, or that organizational safety practices could not be honed to reduce the frequency of such disasters. And, to the extent that some failures are unavoidable, I do not believe that these failures are necessarily very significant in most contexts. Outside a narrow range of technologies with catastrophic failure potential, failures do not always imply disasters, and it is often reasonable to believe that Sagan's glass can be made sufficiently full, even if it can never be filled entirely.

When it comes to technologies that never can be allowed to fail, however, I will argue that we must be more cautious. The levels of reliability that we require of such technologies are actually achievable, given the right conditions and practices. Jetliners prove this. But civil aviation is exceptional in regard to its conditions and practices. Outside of that specific domain, we are not managing our most failure-intolerant systems in ways that would allow the reliability we demand of them. And, in almost all cases, we would not be able to do so. The way that we conventionally understand the reliability of jetliners leads us to think that same reliability should be achievable elsewhere, but a fuller understanding of this achievement implies the opposite conclusion.

PRACTICAL RELEVANCE If the argument is perhaps narrower in its theoretical ambitions than might be intuitive, then it is also probably broader in its applications and implications. Much of its focus is on jetliners and civil aviation, but this is only because they offer a unique window into a wider epistemological problem regarding extreme reliability in complex systems more broadly. It is in this broader context—of what the safety of modern air travel implies for other complex systems with catastrophic potential—that the argument has real purchase. The practices and logics through which we manage ultrareliable technologies closely resemble each other, and all

must grapple with the same epistemological dilemmas, so exploring one illuminates the others.

This is not to say that the argument has no bearing on more mundane technologies with less demanding reliability requirements. This is especially true of the idea at its center, the rational accident, which in principle might apply in almost any technological context. Not every accident is rooted in the fundamental limits of engineering knowledge, of course, and most are better understood through a different lens. But rational accidents are not especially uncommon—as we will see, for instance, they are almost certainly more prevalent than Perrow's normal accidents—and I can imagine the idea potentially being useful in a range of situations.

There are several reasons why, in this volume, I have chosen to explore the idea exclusively in the context of extremely hazardous technologies. One is simply that technologies that cannot be allowed to fail offer an ideal lens through which to explore the mechanisms and implications of inevitable failure. If some accidents are inherently unavoidable—as this argument claims—then this fact has obvious purchase in circumstances where accidents must be avoided at all costs.

Another is that rational accidents are easier to identify and substantiate in this context. Experts scrutinize catastrophically hazardous technologies extremely closely and subject their performance to elaborate assessment and oversight. So on occasions when design shortcomings cause these technologies to fail, it is more difficult to attribute those failures to insufficient effort, rigor or organization, and it becomes more credible to attribute those failures to the fundamental limits of what effort, rigor and organization can achieve. (As a general rule, we might suppose that the proportion of failures best attributed to epistemological limitations rises as technological systems become more reliable overall.)

A third reason for framing the argument in this way is simply that the primary objective of this volume was always to explore the epistemological problem of extreme reliability, more than it was to explore the causes of accidents. Rational accidents are at the heart of its argument and are probably its most generalizable insight, but the intrinsically important problem of extreme reliability is its organizing theme. Insofar as others find the idea of rational accidents useful, therefore, I leave it to them to explore its more diffuse implications and its applications in wider technological contexts. Books need to end somewhere.

STRUCTURE

In keeping with its various turns, the argument that follows has a slightly unusual structure. The heart of the text grapples with the work of making jetliners reliable, but the narrative opens and closes with discussions of nuclear reactors, particularly the 2011 Fukushima meltdowns. This "nuclear sandwich," as one reviewer called it—with reactors as the bread and jetliners as the bacon, lettuce, and tomato—might almost seem like a bait and switch, but it serves a useful purpose. Most directly, it underlines the point that the argument's real significance lies outside of civil aviation. Nuclear reactors are not the only other complex technology from which we demand extreme reliability, as we will see, but they make an excellent counterpoint to jetliners. More clearly than any other system, they exemplify why civil aviation's reliability achievements matter to our lives and livelihoods, and illustrate why the nature of those achievements should be a cause for concern about other technological domains. If understanding jetliners offers us insight into extreme reliability, we might say, then understanding reactors gives that insight meaning.

Between its nuclear bookends (or bread slices), the argument is organized methodically, as a series of subarguments that build on each other progressively. Some of these elements could stand alone, but the wider argument does not lend itself especially well to readers who would skip around. With this in mind, I have labored to make the structure of the wider argument as transparent as possible. To this end, for instance, the chapters that follow are organized into parts. The opening chapters, grouped in part I, outline the core conceptual problem and establish its significance: they make a principled argument that ultrahigh reliability should not be possible in complex technologies, and then problematize this claim by establishing that jetliners are, in fact, ultrareliable. The chapters in part II unpack that problem: they look at the processes by which jetliners are ostensibly made reliable, illustrating the epistemological limitations of those processes and invoking those limitations to articulate the idea of unavoidable "rational accidents." Those in part III resolve that problem: they look behind civil aviation's ostensible reliability processes to explain how experts actually manage the epistemology of jetliner reliability in practice. And those in part IV explore the wider implications of that resolution: they examine the transferability of civil aviation's reliability achievement and reflect on what it means for our relationship to other technical systems with catastrophic potential.

Such structuring might seem heavy-handed at times, but it will hopefully help readers keep the argument's larger trajectory in mind even as the narrative itself sometimes navigates more winding roads. For those who might become lost along the way, however, I will close this introduction with a more detailed chapter breakdown to which they might refer.

CHAPTER BREAKDOWN

PART I

- ***Chapter 1*** invokes the 2011 meltdown at Fukushima to introduce the idea of "catastrophic technologies": complex, sociotechnical systems requiring ultrahigh reliability (i.e., mean-times-to-failure in the region of billions of hours of operation). It argues that the proliferation of these technologies has made the need for experts to achieve, and predictively verify, such extreme levels of reliability a consequential and underappreciated dimension of modern governance.
- ***Chapter 2*** draws on epistemology from STS and the philosophy of science to argue that the extreme reliability required of catastrophic technologies should be impossible for experts to achieve and to verify. It proposes that the processes that experts use to know these technologies contain too many indeterminacies—too many qualitative judgments—to support useful predictions of a complex system's failure behavior over billions of hours of operation to a satisfactory degree of certainty.
- ***Chapter 3*** problematizes the argument of chapter 2 by finding that experts do, demonstrably, achieve and accurately verify ultrahigh levels of reliability in jetliners. It calls this contradiction the "aviation paradox," and argues that it makes jetliner reliability practices uniquely interesting.

PART II

- ***Chapter 4*** outlines the structures, logics, and practices through which US regulators ostensibly govern the extreme reliability of civil jetliners. It finds that the indeterminacies of these practices are visible as ambiguities in the reliability requirements to which jetliners are held.
- ***Chapter 5*** uses engine bird-strike testing as a case study through which to illustrate the inherent limitations of testing for reliability more broadly. It highlights a range of uncertainties about the representativeness of

bird-strike tests, and shows how these uncertainties give rise to doubts about their meaning. It argues that all technological tests necessarily grapple with the same dilemma, and that this should preclude assertions of ultrahigh reliability.

- ***Chapter 6*** follows a similar pattern as chapter 5, but in relation to theoretical models (as opposed to empirical tests) of a jetliner's reliability. It uses redundancy calculations as a case study through which to illustrate the inherent uncertainties of reliability calculations and the doubts to which they give rise. Tests and models might serve distinct purposes at the regulatory level, it argues, but they are subject to the same underlying epistemological constraints.
- ***Chapter 7*** is the heart of the argument in many ways. It draws on the inevitable indeterminacies of tests and models, outlined previously, to explain why some catastrophic failures—what it calls rational accidents—will necessarily and unavoidably elude even the most rigorous engineering analysis and oversight. It then compares this argument to that of normal accident theory, which similarly says that some accidents are unavoidable. Both perspectives imply inevitable failures, it concludes, and neither precludes the other, but the mechanisms (and thus the implications) of each are meaningfully distinct.

PART III

- ***Chapter 8*** resolves the paradox of why jetliners can be so reliable despite the indeterminacies of engineering knowledge and the rational accidents that they imply. The key to understanding jetliner reliability, it argues, is to recognize that civil aviation does not manage reliability in the manner it purports. Rather than using formal assessment practices—tests and models—to examine new machines, it finds that experts have slowly whittled their understanding of jetliners over time, assiduously interrogating failures to hone a common, stable, airframe design-paradigm to extreme levels of reliability.
- ***Chapter 9*** builds on chapter 8 by further substantiating its most controversial claim: that the reliability of modern jetliners depends on them adhering to a common and stable design paradigm. To this end, it looks first at new composite materials, which are often considered evidence of meaningful innovation. It contextualizes this technology and the

opinions held about it, arguing that composites represent a more modest and incremental shift than is immediately apparent, and that opinions about it need to understood in context. From there, it turns to consider instances where airframers have unambiguously embraced radical innovation. It looks at military aviation, arguing that although military aircraft are highly innovative, their reliability is substantially inferior to that of civil jetliners as a result. Finally, it addresses Concorde. It recognizes that the airplane was, without doubt, a radically innovative design, but finds that this was reflected in its counterintuitively dismal reliability record.

- ***Chapter 10*** introduces an organizational problem that arises from the solution to the epistemological aviation paradox. It observes that the stability of airframe designs implies that the companies that design jetliners often forgo, or substantially delay, adopting economically advantageous innovations. And that, in doing so, they consistently trade short-term competitive advantage for long-term safety: behavior that organizational sociologists have long held to be unrealistic. Having outlined and substantiated this new problem, the chapter goes on to resolve it. It looks at regulation but concludes that the ambiguities of technological practice make this an implausible explanation. It then revisits and reassesses civil aviation's structural incentives, arguing that certain characteristics of the industry (primarily its operating volume) give it a unique relationship to failure that incentivizes reliability to a much greater degree than in other catastrophic technological spheres.
- ***Chapter 11*** substantiates and refines the argument of chapter 10, concerning the primacy of structural incentives. It first addresses the 737-MAX crisis. The MAX, it argues, amply illustrates the extreme costs of unreliability in civil aviation—and thus the industry's unusual incentive structure—but it also reminds us that organizations are complex and imperfect, and that (absent periodic shocks) they tend to lose sight of their interests. From there, the chapter turns to consider civil aviation's relationship to crash survivability. It finds that the industry is markedly less proactive about design choices intended to make jetliner accidents less dangerous than it is about those intended to make accidents less frequent. This difference exemplifies the primacy of structural incentives, it claims, as the incentive structures around crash frequency (or reliability) and crash survivability pull in opposing directions.

PART IV

- ***Chapter 12*** begins to explore the wider implications of civil aviation's approach to managing extreme reliability. It proposes that the public portrayal of jetliner reliability management—as grounded in formal analysis via tests and models—is misleading and elides the real practices and conditions on which ultrahigh reliability depends. It then explores the reasons and logic driving this misportrayal, finding that it sometimes can be functional in civil aviation. From there, however, it argues that the misportrayal also gives rise to underappreciated costs and dilemmas, many of which are likely to be less acute in civil aviation than they are in other catastrophic technological domains.
- ***Chapter 13*** focuses on one prominent problem that arises from the misportrayal of civil aviation's reliability practices: the false but widespread belief that its achievements in this regard should be replicable in other catastrophic technological domains. Understood properly, it argues, the ultrahigh reliability of jetliners depends on resources and practices—such as high operating volumes, a commitment to design stability, and a legacy of instructive failures—that are not available, practicable, or achievable elsewhere. The most consequential cost of misportraying jetliner reliability management, it concludes, is that this difference becomes opaque.
- ***Chapter 14***, by way of coda, returns to Fukushima. It invokes the accident to argue that reactors exemplify the essential differences between jetliners and other catastrophic technologies. It argues that experts working in the civil nuclear sphere enjoy few of the resources that their counterparts in civil aviation use to navigate the harsh epistemology of extreme reliability. Although ostensibly governed via equivalent structures and processes to jetliners, it concludes, reactors cannot be as reliable. This difference is difficult to observe because of the small number of reactors in operation, it argues, but the shortcomings of reactors are nevertheless visible if we know where to look—not least in Fukushima itself, which might reasonably be understood as a rational accident.

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