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

14 FUKUSHIMA REVISITED: REAPING THE WHIRLWINDS OF CERTAINTY

> Eight years involved with the nuclear industry have taught me that when nothing can possibly go wrong and every avenue has been covered, then is the time to buy a house on the next continent.
>
> —Terry Pratchett

14.1 GOVERNING REACTORS

WOEFULLY UNPREPARED

As the sun set on the first day of the Fukushima disaster in 2011, the embattled on-site operators were grappling with mundane supply shortages. Some searched nearby homes for flashlights so they could monitor crucial gauges through the night. The plant did not have its own supply. Others scavenged among the site's flooded detritus for batteries so they could maintain contact with the emergency response center. Their mobile communications units had been equipped with only one hour of reserve power. Nobody had planned for a catastrophe of the scale they were facing. Similar failures of foresight were manifest at every turn. Personal dosimeters were in short supply. On-site dosimeters maxed out at levels far lower than required. The plant's *Severe Accident Manual* lacked appropriate guidance. The designated off-site command post had no air filters or reserves of food, water, or fuel. The region around the plant was even less prepared. Protective medications were scarce, and their distribution was hindered by byzantine rules. Evacuation procedures and official announcements were unrehearsed and poorly considered. There can be no such thing as an orderly nuclear meltdown, but by any reasonable standard, Japan was woefully underprepared for the crisis it was facing (Perrow 2011; Osnos 2011; McNeill and Adelstein 2011; Lochbaum, Lyman, and Stranahan 2014; Kubiak 2011; Onishi and Fackler 2011).

It is tempting to blame Japan for its inadequate preparation, but few if any of its peers were substantially better prepared for such a crisis. In the US, for example, reactor emergency and evacuation planning was (and largely remains) similarly deficient, having been premised almost exclusively on responding to small leaks rather than major meltdowns (Clarke and Perrow 1996; Perrow 2011, 46–47; Kahn 2011). Indeed, the inadequacy of US preparation became evident during the Fukushima crisis, when its tools for assessing radiological threats to its embassy maxed out at fifty miles (Lochbaum et al. 2014, 63).[1] There have been global efforts to improve preparations for major reactor accidents post-Fukushima. Such efforts are limited, however, and it is fair to say that the accident highlighted a blind spot for disaster that still characterizes atomic energy worldwide.

This blind spot is inscribed in the very landscape of atomic energy. Take, for example, the siting of various reactors upwind of major metropolitan areas, such as Indian Point, thirty-five miles north of Manhattan, or Hinkley Point, about the same distance from Bristol in the UK. Or consider that, for economic and political reasons, most nuclear plants are clusters of multiple reactors, meaning that a major failure in one reactor can jeopardize its neighbors, as occurred at Fukushima (Perrow 2007, 136; Osnos 2011, 50). No society would make such choices if it understood Fukushima-scale accidents to be a meaningful possibility. And, having made such choices, no society could adequately prepare for such accidents, even if it resolved to do so. Cities on the scale of Tokyo, New York, or even Bristol are not easily evacuated. Cities of any scale are not easily evacuated. If at first you don't succeed, an old joke goes, then nuclear energy isn't for you.

PRACTICALLY UNTHINKABLE

Seen through the lens of the preceding argument, the global blind spot for reactor catastrophes reflects the positivist promise of implausible reliability,[2] and Fukushima exemplifies the dangers inherent in that promise. Simply stated, states don't plan for such catastrophes because they are confident that such catastrophes will never occur (Rip 1986; Clarke 2005; Fuller 1976; Kahn 2011; Lochbaum et al. 2014; Wellock 2017). The US, for instance, draws a distinction between credible reactor events, such as small radiation leaks, which it deems worthy of public concern, and hypothetical events, such as Fukushima-level meltdowns, which it does not. Japan is no different. This is why Fukushima's Comprehensive Accident Plan focused exclusively on minor incidents, assuring its readers that "the possibility of a severe accident occurring is so small that, from an engineering standpoint, it is practically unthinkable" (Lochbaum et al. 2014, 16).

States are confident in this belief because their experts have assessed new reactor designs and determined them to be ultrareliable. (Westinghouse's new AP1000 reactor, for example, has been calculated to enjoy a "core melt frequency" of no more than once in every 2.4 million reactor-years, and a "large release frequency" of once in every 27 million reactor-years [Sutharshan et al. 2011: 297–298]). And states are comfortable trusting such determinations because they see no a priori distinction between claims about ultrahigh reliabilities in reactors on one hand, and myriad other engineering claims of more demonstrable efficacy on the other, not least those made about the reliability of jetliners, which are ostensively equivalent and undeniably well grounded.

It is common, in fact, for jetliners to be invoked in discussions of reactor safety in ways that imply that the technologies pose comparable challenges regarding their governance (e.g., Bier et al. 2003). To a greater degree than with most catastrophic technologies, therefore, it is worth considering their relationship in more detail.

14.2 JETLINERS AND REACTORS

PARALLEL STRUCTURES

Even more than jetliners, reactors exemplify the approach to governing catastrophic technologies outlined in chapter 1, wherein safety is framed as a function of reliability and reliability is presented as a definitively (and predictively) knowable property of complex systems.[3] It is worth noting, in fact, that this approach actually originated in the US nuclear sphere for the explicit purpose of bounding disasters out of public discourse (Downer and Ramana 2021; Wellock 2017; Rip 1986; Miller 2003a). (To tell this story here would require an excessive detour, but in brief, it was driven by the need to justify the reactor program to Congress, and was fiercely contested by critics at the time [Okrent 1978; RARG 1978].)[4]

What began in the civil nuclear sphere, however, soon colonized civil aviation. The structures and practices through which the US—and, largely as a consequence, most other nations—manages reactors and jetliners have long coevolved, with each influencing the other and both taking similar forms as a result. So it is, for example, that in both spheres, a dedicated regulatory agency promulgates and enforces formal certification standards to ensure the reliability of new designs prior to their operation. The NRC's equivalent to the FAA's type-certification process is its design-certification process, overseen by subdivisions of its Office of New Reactors. Like type certification, it stipulates quantitative failure-performance requirements and assessment practices rather than mandating specific engineering choices, and it relies extensively on delegation to ensure compliance (NRC 2009, 18).

It makes sense that jetliners and reactors would engender similar oversight and governance structures. The two technologies—born within a few years of each other (1952 and 1954, respectively)[5]—appear to pose similar challenges. Both are charismatic, highly complex, tightly coupled sociotechnical systems, designed for near-constant operation over several decades. When in operation, both must actively work to negate powerful countervailing forces (gravity and fission, respectively), and each demands broadly equivalent levels of reliability. (As previously discussed, the consequences of meltdowns are much greater than those of plane crashes, but with many more jetliners in operation than reactors, the reliability required of each is about the same.)[6]

Such parallels have made the nuclear sphere a principal beneficiary of the credibility that civil aviation has lent to the positivist portrayal of ultrahigh reliability management. It is ironic, therefore, that reactors exemplify the ways in which most catastrophic technologies are radically different from jetliners with respect to their reliability.

UNDERLYING DIFFERENCES

Ostensibly very similar when viewed through the lens of rule-governed objectivity, reactors and jetliners look very different when understood in relation to the principles and practices outlined in preceding chapters as being essential to ultrahigh reliability. Perhaps the most fundamental differences in this regard relate to their cumulative service experience. As we have seen, reactors operate in much smaller numbers than jetliners. Tens of thousands of jetliners traverse the skies every day, whereas the World Nuclear Association (2019) was claiming just 450 reactors in operation as of 2019, spread across plants in thirty countries.

Compounding this difference in absolute numbers, moreover, are dramatic design variations between reactors, which greatly reduces the relevance of any one reactor's service data to most others. "Similarity" might be impossible to quantify precisely, but even the most cursory taxonomy of reactor architectures makes it clear that the nuclear industry does not practice design stability in anything like the same fashion as civil aviation. At the very broadest level, for instance, the world's reactors can be categorized into a range of types, including pressurized water reactors, boiling water reactors, pressurized heavy water reactors, gas-cooled reactors, and molten salt reactors, which vary even in their fundamental operating principles. Each type can further be divided into multiple generations. (The industry divides pressurized water reactors into four generations, for instance.) And even reactors of the same type and generation are highly variable, partly because there are dramatic variations among the offerings of different manufacturers, but also because each reactor is tailored to suit specific local conditions (such as seismic requirements), a process that creates significant variations between otherwise identical designs (IEE 2005).

These differences in operating volume and design stability mean that the civil nuclear industry's statistically relevant service experience is radically far than that enjoyed by its civil aviation peers. In 2019, again, the World Nuclear Association (2019) was claiming "more than 17,000 reactor-years of experience," a number that pales in comparison to the 45 million-plus years of experience that jetliners accrue every single year, especially given the manifest dissimilarities among the reactors it is counting. This difference in service experience, in turn, dramatically shapes the nuclear industry's relationship to ultrahigh reliability, because "17,000 reactor years of experience" doesn't mean much when modern reactors are expected to melt down less than once in every million years. Even if we imagined that all reactors were identical in their designs, and their years of operation were completely failure free—neither of which is remotely true—that data would be insufficient for statistically testing the industry's reliability claims.

In practical terms, this relative dearth of experience means that experts in the nuclear sphere have not been able to hone reactor designs (and assessments of those designs) in anything like the same way as their counterparts in civil aviation. Modern jetliners have proved that in principle, engineers can leverage the hard-earned lessons of experience to gradually achieve extreme reliabilities in complex systems, given enough time and enough design stability. But engineers working in the nuclear sphere enjoy neither of these resources in anything close to the required amount. The limited number of reactors in operation offer far fewer lessons from which they might learn the limits of their abstractions, and the diversity of reactor designs limits the relevance (and thus the usefulness) of any lessons that they do learn.[7]

The nuclear industry's paucity of service experience also has implications for its incentive structures, which again differ dramatically from those that frame the choices made in civil aviation. With so few reactors in service, those reactors could all be far less reliable than promised and still operate for an extremely long time (probably their entire service lives) before that unreliability manifested in a catastrophic failure. Even though rare, such failures would still be intolerable, given their potential consequences, but the odds of them affecting any specific manufacturer, regulator, or operator would be low. And in the unlikely event that an organization *did* see a catastrophe in one of its reactors, the odds are good that many years would have passed since the failed reactor had been commissioned, designed, approved, purchased, and built. This would mean that the individuals involved in these processes would have moved on to new roles (or retired), and the discourse around reactor safety would be tied to more modern and substantially changed designs.

In combination with other factors—such as statutes that indemnify operators against catastrophic losses[8]—these conditions dramatically limit the costs that organizations can incur for falling short of their extreme reliability assertions. In doing so, they structure the landscape of incentives that act on those organizations. It is easy to imagine this incentive structure weighing heavily on organizational attitudes, cultures, and interpretations across the industry, especially given the longstanding political and economic precarity of atomic energy, which keeps all its associated investment and employment in near-constant jeopardy.

For an insight into this dynamic, we might look to an introspective article by a nuclear engineering consultant in the aftermath of the Chernobyl accident in 1986 (Davies 1986, 59). He writes:

> At some point in their careers most nuclear engineers and scientists have had to resolve in their own minds the awful dilemma posed between the potentially catastrophic effects of a serious accident at a nuclear power plant and the remoteness

> of the possibility of its occurrence. Most have squared the circle by concluding, consciously or otherwise, that the remoteness of the chance of an accident was such that serious accidents likely *would* never happen and even *could* never happen. Even if one did happen, it would not be at your plant.

FUNDAMENTALLY UNSAFE

This is all to say that any symmetries between the reliability problems posed by jetliners and reactors are superficial and misleading. When extreme reliability is understood as depending on things like consecrated service experience and favorable structural incentives, as well as just rigorous analysis, then it becomes clear that these technologies are not alike.

Reactors are extraordinary engineering achievements, built and assessed by accomplished and dedicated experts. Nobody can deny this. Yet their failure behaviors will always be subject to epistemological and organizational constraints that no amount of effort or analysis could hope to overcome. Making reactors as reliable as jetliners would require the industry to commit to a common design paradigm, it would mean building tens of thousands of nuclear plants, and it would mean enduring many more accidents like Chernobyl and Fukushima. The first two of these conditions are unrealistic on many levels—even if they were incentivized, which they are not—and the last is intolerable; the costs of nuclear accidents will always eclipse the value of any insights they could offer.

Absent these conditions, the experts who design and evaluate reactors are wholly dependent on their abstractions. And the extreme behaviors they are representing—failure performances over billions of hours of operation—demand an unrealistic degree of perfection from those abstractions. Over enough time, therefore, it is reasonable to expect unforeseen failures arising from fundamental ambiguities in the knowledge on which reactors are built—what I have called "rational accidents."

For evidence that the reliability claimed of reactors is indeed unrealistic, we need look no further than the industry's service record. Statistically, proving ultrahigh reliability requires a huge amount of failure-free operation, but *disproving* it only takes a tiny number of catastrophic failures. This means that the history of nuclear energy could be dispositive, if only we allowed it to speak.

The exact number of historical reactor accidents varies depending on the definitions of key variables—whether Fukushima should count as one "plant" meltdown or as three separate "reactor" meltdowns, for example[9]—but relatively conservative estimates put the rate of serious accidents at somewhere between 1 in every 1,300–3,600 reactor-years (Ramana 2011; Taebi, Roeser, and van de Poel 2012, 202–206; Suvrat 2016). Raju (2016, 56) offers an uncommonly sophisticated analysis of this question and comes to an unambiguous conclusion: "It is clear that the results of [nuclear safety assessments] are untenable in the light of empirical data."

Yet we do not allow this record to speak. With very few exceptions, public accounts of reactor accidents interpret them as evidence not of fundamental limits to engineering ambition but of avoidable human failings: correctable aberrations with limited implications for reactor safety more broadly. Chernobyl, for example, is widely remembered as a Soviet accident more than a nuclear accident—the result of failings specific to regional design, oversight, and operating practices, rather than anything more generalizable (Socolow 2011).

Fukushima is understood in much the same way. Official and unofficial accounts of the disaster evince a wide range of narratives. Some highlight specificities of the plant's design or geography, for example, while others point to failures in Japan's specific regulatory system or corporate governance (Downer 2014; Hamblin 2012). Few, if any, frame the accident as evidence that ultrareliable reactors, as well as accurate assessments of those reactors, are fundamentally unachievable goals premised on unrealistic expectations about the knowledge on which they are built. Even though these narratives often highlight real failings, therefore, they routinely miss the larger point: that unavoidable catastrophes lurk in the irreducible uncertainties of their design and operation.

This is a shame, as the story of Fukushima illustrates the point about irreducible uncertainties quite well.

14.3 FUKUSHIMA'S FAILING

FUNDAMENTAL ERRORS

In the most basic and fundamental sense, Fukushima failed because it was built on erroneous assumptions. As with all nuclear plants, its design was framed by a broad set of premises, codified in a set of documents referred to as its "design basis." At the most fundamental level, the design basis of a nuclear plant consists of two interrelated elements. The first is an understanding of the parameters within which the plant must operate: for example, environmental conditions and any hazards it might face. The second is a set of technical requirements that reflect those parameters: for example, the minimum tensile strength of specific structural elements or the minimum separation distance between certain subsystems (Wyss 2016, 31).

So it was that, among many other things, Fukushima's design basis included specifications relating to potential tsunamis: specifying the hazards they posed and how those hazards were to be neutralized. Specifically, it set the maximum height of any tsunami that the plant could face at 3.1 meters (about 10 feet); and to ensure a comfortable margin of safety, it required the plant to have flood defenses capable of withstanding a tsunami almost twice that height: 5.7 meters (19 feet). These measures, its designers and regulators decided, would reduce the risk of tsunami inundation to essentially zero.

This conclusion, that tsunami inundation was functionally impossible, then shaped numerous downstream design and planning choices. Fatefully, for example, it informed a decision to place all the plant's redundant backup generators underground for maximum protection from terrestrial hazards. It also influenced the siting of the plant, the elevation of which engineers lowered by 25 meters (80 feet) so that the reactors would stand on bedrock and be more resilient to earthquakes (Lochbaum et al. 2014).

In the event, however, the tsunami that swamped Fukushima in 2011 far outstripped the levels specified in its design basis. At somewhere between 12 and 15 meters (40–49 feet)—over four times the calculated maximum—it crashed over the plant's seawalls, swamping the entire site and drowning all the underground emergency generators simultaneously in an unanticipated "common mode" failure. Thus began a cascade of failures that quickly cumulated in catastrophe. Faced with a massive regional power outage, the plant was unable to fall back on its nonfunctioning generators to cool its furiously heat-generating reactors. Untempered fission then did as untempered fission does, and the system inexorably spiraled out of control.

Viewing Fukushima's failures in this way—as born of a foundational error in the plant's design basis—muddies narratives about flaws in its design. The disaster's immediate wake saw various embattled industry proponents gamely contending that the plant had not actually failed at all, since it had never been designed to survive the conditions that brought it down (e.g., Harvey 2011; Hiserodt 2011). The American Nuclear Society (2011) even went so far as to argue that Fukushima "could actually be considered a 'success,' given the scale of this natural disaster that had not been considered in the original design." The wider logic of such claims is contestable at best—the plant clearly failed, by any reasonable definition of the word—but the frustration is understandable. The plant's design basis was established at the earliest stages of its conception. Its underlying assumptions about tsunamis must have been all but invisible to the experts who then used it to plan, build, and assess the plant itself. No additional effort on their part would likely have revealed the error. They extensively modeled and tested the plant's design, and those models and tests—which were framed on the same foundational belief about tsunamis—repeatedly validated their conclusions. The plant's superintendent spoke for almost everyone involved in its safety when he told reporters that "we encountered a situation we had never imagined" (Lochbaum et al. 2014, 13).

In light of all this, it is tempting to blame the experts who created Fukushima's flawed design basis for the plant's demise, but this too is problematic. Examined closely, there are few reasons to doubt that the design basis reflected the best understanding of tsunamis that was available when construction on the plant began in 1967. The maximum height calculation was wrong, in essence, because seismologists had substantially underestimated the potential size of earthquakes in the region, but this error would not have been apparent at the time. The Tōhoku megaquake that felled Fukushima was an estimated magnitude—9 or 9.1—that was not even recognized as geologically possible in 1967. "Magnitude 9" only became a generally recognized class of earthquakes around 1980, and for decades afterward, many seismologists continued to believe they were limited to certain types of subduction zones and could not occur near Fukushima. History seemed to confirm this assessment. Prior to the twenty-first century, there was no record of such an earthquake, or a tsunami, ever occurring in the region (Nöggerath, Geller, and Gusiakov 2011, 38–42; Lochbaum et al. 2014, 52).

It is also true, of course, that Fukushima's demise came long after the plant was built, and many critics have pointed out that science's understanding of earthquakes had evolved considerably by 2011. As noted previously, seismologists began to recognize the existence of megaquakes in the late 1970s and early 1980s. And by the early 2000s, new techniques for identifying and interpreting sedimentary rocks had led some "paleo-tsunami" researchers to conclude that at least three such quakes had occurred in the region over the previous 3,000 years (Nöggerath et al. 2011, 41; Minoura et al. 2001). The nuclear community was aware, or at least not universally *unaware*, that these findings had implications for Fukushima Daichi. In the years leading up to the disaster, several bodies had formally questioned the plant's design basis on these grounds. In 2002, for instance, the Headquarters for Earthquake Research Promotion issued revised predictions for the region, which, when modeled by the plant operator (TEPCO), implied a potential tsunami of 51 feet (15.7 meters) at the site (Lochbaum et al. 2014, 52).

Understandably, therefore, several accounts of the disaster contend that the plant should have been updated to reflect this new evidence (e.g., Nöggerath et al. 2011; Acton and Hibbs 2012). Critics point to missed opportunities created by the new seismology, observing, for instance, that Japan's Nuclear Safety Commission revised its seismic guidelines in 2006 without mandating significantly upgraded tsunami protections, and it failed to act in 2009 after a senior geologist issued strong warnings about tsunami risks (Nöggerath et al. 2011, 42). These missed opportunities look damning in hindsight, and many claim they constitute a culpable failing that should have been avoided. By this view, Fukushima might have been blameless in 2000, but not by 2011.

In principle at least, the design bases of nuclear plants are living documents that can be revised over time, so the claim that they should have been updated is not inherently unreasonable. Still, however, the "failure to heed new seismology" critique is unsatisfying on at least two significant levels. The first is that it assigns blame but offers no comfort, for even if we allow that Japan's nuclear establishment should have revised or closed Fukushima in light of new seismology, then there was still a period of several decades during which there was no reason to doubt the flood defenses. The accident could easily have occurred in the years prior to the discovery of megaquakes, in other words. And, had that occurred, the plant's design basis would have been wholly congruent with the best science available. Those who claim that Fukushima is blameworthy on these grounds, therefore, are also tacitly conceding that failures—rational accidents—are sometimes inherently unavoidable. When contemplating disasters that might threaten major cities, this is not a reassuring thought.

The second issue with this critique is that it isn't very compelling even on its own merits. This is because hindsight gives a misleading perspective on the science and the way it was received. In debating whether Fukushima's design basis should have been altered, it is important to understand that the new seismological findings were vigorously contested within the expert community itself. TEPCO and its regulator countered them with rival findings, also backed by respected seismologists, based on models and simulations that either placed less emphasis on historical earthquake data (e.g., Yanagisawa et al. 2007) or framed that data in ways that downplayed its tsunami implications (Nöggerath et al. 2011, 42). As in all such debates, moreover, policymakers were ill equipped to judge the relative merits of different positions (Hirano 2013, 18–19). "[C]onfusion reigned in the field of seismic risk," as Lochbaum et al. (2014, 44) put it, with "even state-of-the-art science [being] unable to shed much light on the question of 'how safe is safe enough' when it came to building nuclear power plants in Japan."

So it is that, while the nuclear community's resistance to new estimates of seismological risk was undoubtedly self-serving, and perhaps even contrary to the broader scientific consensus, it was at least defensible. Every new scientific finding is a convoluted and incremental achievement, especially if they materially affect powerful institutional actors (Jasanoff 1990; Salter 1988). Complex research is never accepted promptly by parties that stand to lose from it, and in this instance, the nuclear community (and the state itself) stood to lose a lot. It is doubtful whether it would even have been possible to update Fukushima Daiichi to defend against earthquakes and tsunamis that were substantially more severe than those outlined in its original design basis.

In a world where some major polities still barely recognize the existence of anthropogenic climate change, therefore, it is wholly unrealistic to imagine that a private corporation would promptly condemn, or even substantially rebuild, a nuclear plant when confronted with esoteric and contested new paleo-tsunami findings, or that a state would require it to do so. (It is worth noting that the US, for its part, had done little to revise its reactors as seismology evolved [Dedman 2011; Lochbaum et al. 2014, 115].)[10] Arguments that invoke this expectation to assign blame for the Fukushima disaster are probably better understood as critiques of modernity itself.

Insofar as this expectation is unrealistic, moreover, then Fukushima might reasonably be understood as a rational accident. Seen in this way, its wreckage and fallout are a monument to technological hubris; they represent a failure to see the limits of what engineers can know, and what those limits imply for our technological ambitions.

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