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

12 BURDENS OF PROOF: THE HIDDEN COSTS OF POSITIVISM

> In theory, there is no difference between theory and practice; in practice, there is.
> —Walter J. Savitch

> For a successful technology, reality must take precedence over public relations, for nature cannot be fooled.
> —Richard Feynman

12.1 A THEATER OF OBJECTIVITY

PRETENSE AND HUMBUG

In his memorable introduction to the principles of structural engineering, J. E. Gordon (2018 [1991]) speaks of a divide that characterized the discipline in the early- to mid-nineteenth century. During this period, he explains, structural engineering on the European continent was dominated by a theoretical and quantitative "French" tradition, while a more pragmatic and craft-based "British" tradition ruled on the other side of the English Channel. He goes on to note that the relative standings of these two traditions seem incongruous today. The greatest and most prestigious structures of the age were being built by British engineers like Thomas Telford and Isambard Kingdom Brunel, and yet it was the continental Europeans who pioneered many of the mathematical theories and models that have come to define the discipline. From a modern perspective, Gordon argues, the most lauded engineers of the nineteenth century look almost primitive in their disdain for abstract theory and quantitative analysis (Gordon 2018 [1991], loc. 794–800).

Gordon has an interesting interpretation of this incongruity. He doesn't see it as evidence that French engineers were undervalued by their contemporaries. He sees it, rather, as evidence that British engineers understood something fundamental about their discipline that most modern observers have come to forget. He writes:

> We must be clear that what Telford and his colleagues were objecting to was not a numerate approach as such—they were at least as anxious as anybody else to know what forces were acting on their materials—but rather the means of arriving at these figures. They felt that [French] theoreticians were too often blinded by the elegance of their methods to the neglect of their assumptions, so that they produced the right answer to the wrong sum. In other words, they feared that the arrogance of mathematicians might be more dangerous than the arrogance of pragmatists, who, after all, were more likely to have been chastened by practical experience. (Gordon 2018 [1991], loc. 800–807)

He tells this story as a way of contextualizing and qualifying the formula-rich ("French") theory he is about to explain. It introduces a cautionary motif that runs throughout his book, about structural engineering not being wholly reducible to mathematical tools and formal rules. Engineering knowledge still has vital qualitative dimensions, he wants to say, but the emphasis the discipline places on austere calculations and formal models can seduce us into believing otherwise.

Gordon is far from being the first writer to make this point. Indeed, it is a familiar theme of texts written for lay audiences by academic engineers and observers of engineering, especially when those texts reflect on the nature of the discipline itself (e.g., Petroski 1992a; Blockley 2012). So it is, for instance, that Wynne (1988, 150–151) speaks of a "huge contradiction" at the heart of modern technology governance, wherein a "neat and tidy public image" of engineering practice belies a much "messier" reality. Vincenti (1990), notably, makes a similar argument as Gordon, but in relation to aeronautical engineering, again illustrating it with reference to historical wisdom that gradually came to be forgotten. To this end, he quotes a British engineer speaking to the Royal Aeronautical Society in 1922: "Aeroplanes are not designed by science but by art, in spite of some pretence and humbug to the contrary. I do not mean to suggest that engineering can do without science, on the contrary, it stands on scientific foundations, but there is a big gap between scientific research and the engineering product which has to be bridged by the art of the engineer" (Vincenti 1990, 4).

As Vincenti goes on to argue—and as preceding chapters of this book have illustrated and explored at length—it would not be unreasonable to make the same observation about jetliner design a century later. The process of designing and evaluating ultrareliable airplanes undoubtedly draws on deep foundations of theory, measurement, and calculation (manifest in tests and models), but those foundations cannot sufficiently account for the extraordinary failure performance that jetliners achieve in service. As we have seen, that performance is only really intelligible when also understood in relation to various practical resources and qualitative arts: a deep well of examined experience (thousands of heavily interrogated failures); the tacit expertise required to make complex judgments about relevance; an economically incentivized institutional willingness to consistently prioritize reliability over costs in those judgments, and more. This is why the key to unraveling the "aviation paradox," as I have called it, lay in recognizing that experts have not perfected the accuracy of their tests and models to a level that finitists would find implausible, but circumvent those tests and models by leveraging real-world service experience. They don't "design," "calculate," or "measure" ultrahigh reliability so much as they "whittle," "infer," and "consecrate" it.

As we have also seen, however, jetliner reliability management is not portrayed as functioning in this way. Instead, it is couched in a language of rule-governed objectivity—quantified requirements and calculative practices—that obscures its many subjectivities, indeterminacies, and practicalities, and presents instead an idealized facade of objective rules and deterministic certainty.[1] In this positivist vision of aviation safety, tests and models are capable of revealing the verifiable *truth* of a system's failure behavior, and uncertainty is an aberrant and resolvable condition. By this view, outside experts can *police* the reliability of new jetliners, or any other complex system, to ultrahigh levels, via procedural, rule-governed measurement and analysis. And the public has no more reason to doubt their conclusions than to doubt expert assessments of any other engineering variable.

This is all to say that there is still a lot of "pretense and humbug" in the discourse around jetliner reliability. A tension exists between the way that ultrahigh reliability engineering is commonly understood and the manner in which it actually operates. Jetliners are not safe for the reasons that we believe, and their regulators do not, and cannot, serve the function that we imagine they do. And this tension, in turn, gives rise to a wide range of downstream contradictions and misapprehensions about the function of rules, the abilities of oversight bodies, and the real conditions that foster reliability.

These contradictions and misapprehensions have unanticipated and potentially harmful consequences, which the balance of this chapter will explore. Before that, however, it is worth contextualizing the tension between practice and portrayal a little more.

12.2 PERFORMING OBJECTIVITY

"WHITE-BOXING"

Engineers are far from unusual in overemphasizing the objectivity of their work in public contexts. As a range of literature testifies, it is common for experts to invoke unrealistic abstractions and stylized practices to foster misleading impressions of objectivity and certainty (Lampland 2010; Van Maanen and Pentland 1994; Miller 2003b). This phenomenon is routinely observed in regulatory contexts, where public accountability is important. As Power (2003, 6) puts it, "Regulatory projects are always in some sense visionary . . . possibilities and aspirations for control and order get projected via discussion documents, codes, guidance manuals and the law, often in an ideal form abstracted from the messy realities of implementation."

Nowhere is this more the case than in contexts where scientists and engineers, specifically, explicitly engage with matters of public concern. In these circumstances, it is entirely conventional for technoscientific experts to construct performances of objectivity for lay audiences, hiding their uncertainties and disagreements backstage so as to better present a positivist caricature of their work (e.g., Hilgartner 2000; Rip 1985; Jasanoff 1990, 2003; Wynne 1988, 1989; Collins and Pinch 1998). "[I]f one looks carefully into reassuring public rhetorics of [engineering], one may find hints at a less rule-governed existence for technologies," Wynne (1988, 151) writes, "but, for public consumption, these are highly attenuated and cryptic, if they are expressed at all."

Hence the contradictions outlined here.

Scholars usually understand these stylized performances of objectivity in relation to a modern cultural and legislative milieu that demands formal accountability and idealizes quantitative rules. The emergence of this milieu is explored in an extensive body of scholarship that defies easy summary (e.g., Hacking 1990; Porter 1995; Lampland 2010; Desrosieres 1998;

Verran 2012; Power 1997).[2] Most accounts agree, however, that it has been largely driven by organizational efforts to build legitimacy and authority. Uncertainties invite dissent, and qualitative judgments imply arbitrariness or bias, especially in technoscientific contexts where formal methods have long been idealized. “Non-engineers often seem to stand in awe of the practitioners of applied science and engineering,” as Turner and Pidgeon (1997, 14) put it, “regarding them as inhabitants of a world where rationality reigns supreme, where alternative courses of action can be measured and rigorously compared and where science informs every decision.”

So it is that masking uncertainties and judgments behind a performance of rule-governed objectivity plays to societal expectations regarding expertise and authority. In doing so, it bolsters the credibility of expert knowledge claims, and thus the power of the organizations that wield them (March and Simon 1958, 165; Espeland and Stevens 2008). “Following rules may or may not be a good strategy for seeking truth,” Porter (1995, 4) observes, “but it is a poor rhetorician who dwells on the difference. . . . Better to speak grandly of a rigorous method, enforced by disciplinary peers, canceling the biases of the knower and leading ineluctably to valid conclusions.” It is even better, we might say, if those rules are expressed quantitatively, given that numbers—terse, precise, contextless abstractions—are uniquely able to project clarity while simultaneously masking the circumstances of their creation (Espeland and Stevens 2008; Porter 1995).

Note that the argument here is not necessarily that engineers, or experts of any kind, are consciously hiding the unruliness of their work. It is more that a degree of performance has become encoded in the traditions and expectations that frame their stylized discourse. And these traditions and expectations have grown out of longstanding public pressures on organizations to speak authoritatively.

Wynne (1988) neatly captures the essence of this relationship as it pertains to technology oversight in particular, referring to performances of rule-governed objectivity in this context as “white-boxing.” The term is intended to convey the irony of a process wherein the inner workings of a system are occluded from public view—as in “black-boxing”—but via processes that are purportedly designed to make those workings visible and accountable.[3] (It is a useful label, as it highlights both the nature and the purpose of the misrepresentation, and I will borrow it for the discussion that follows.)

NECESSARY FICTIONS?

It is worth recognizing that observers sometimes construe the white-boxing of technology, and performances of objectivity in technoscientific work more broadly, as a relatively benign and functional conceit. Ezrahi (2012), for instance, argues that all democratic societies adopt "necessary fictions" in managing their relationship to technical systems.

Defenders of this view usually argue, in different ways, that white-boxing technoscientific claims helps imbue those claims with the authority they deserve. The logic of this is relatively intuitive. As we saw in chapter 2, for example, the indeterminacies of real engineering practices rarely detract from the usefulness and efficacy of expert engineering assertions. When engineers say a bridge will collapse or an engine will overheat, then we ignore them at our peril. So if misportraying those assertions as grounded in wholly objective and rule-governed practices—in conformity with popular perceptions of, and expectations about, engineering work—helps make them convincing, then it can be difficult to see the harm. It is easy to be skeptical of expert claims that are lubricated by the "oily art" of stage management, as Shapin (1995, 255) puts it, but sometimes even the most creditable claims need lubrication (see also Sismondo 2010, 170–173).

By some accounts, white-boxing can also facilitate the *production* and *application* of technoscientific knowledge, as well as its public reception. The argument here is that experts can use performances of objectivity to create backstage spaces, wherein they can more freely and productively exercise important qualitative judgments without inviting censure for not conforming to idealized expectations about their work (Wynne 1998; Espeland and Stevens 2008; Lampland 2010; Schulman 1993). In this view, the hidden ambiguities and interpretive flexibilities of rules and calculations function as a resource. They create slack in the system that experts can use to navigate the subjectivities of their work behind the scenes, without challenging the conventional mores of modern accountability. (So it is, we might say, that aviation experts more easily manage the debates that permeate bird-strike tests because those debates occur backstage, in esoteric technical discourse, where doubts can be voiced and judgments exercised without inspiring public controversy.)

It is easy to imagine that civil aviation, in particular, might exemplify this phenomenon. After all, jetliners might not be reliable for the reasons we imagine, but, as we have seen, decades of service data show that they are as reliable as experts claim (or have been in recent decades, at least). And even if it is impossible for regulators to police that reliability directly, the same service data show that airframers have not exploited this inaptitude in ways that undermine the reliability of their designs. In this context, therefore, there is a strong argument to be made that if the industry and its regulator were to subvert public expectations by being fully transparent about their work, with all its inherent messiness, then they could rob the industry of public confidence that it genuinely merits, and undermine practices that are functioning well.

At the same time, however, any instance where democratic institutions are misleading publics for their own benefit probably merits some introspection. Roads paved with good intentions proverbially lead to dark places, after all. It is important to note, therefore, that even scholars who maintain that performances of objectivity can be functional usually acknowledge that such performances often come with hidden costs. Civil aviation, for example, is in the ironic position of having extraordinarily effective practices for making jetliners ultrareliable, which it misportrays to make them credible. It is difficult to imagine that this is entirely optimal from a governance perspective. As we saw in chapter 11, for instance, airframers we able to exploit this misportrayal to delay survivability measures, invoking the idea that technical questions should have definitive answers to repeatedly call for further research.

With this in mind, let us consider what other hidden costs might arise from performing objectivity.

12.3 THE PRICE OF PERFORMANCE

PRESENTATION COSTS

Any discussion of the costs of white-boxing technological practices would be incomplete without recognizing the straightforward expenditures—of time, effort, and money—that it invariably incurs. Stylized performances of rule-governed objectivity might be useful, but the processes through which they are enacted can be extraordinarily onerous and expensive, especially in catastrophic technological contexts. As we have seen, it is probably fair to say that the work involved in making the reliability of jetliners accountable is substantial, probably costing as much as the work of making jetliners reliable. Understand also that the costs of performing objectivity and formal accountability should not be measured in hours and dollars alone, but also in terms of the less tangible burdens that they place on factors like efficiency, organizational culture, and morale. Such effects are underexplored and difficult to quantify, but a mounting body of scholarship suggests that audit requirements can significantly impede organizations that create complex systems (e.g., Graeber 2015; McCurdy 1993). Vaughan (1996) and McCurdy (1993), for example, both portray NASA as an institution that became progressively weighed down by accountability-driven bureaucracy, to the point where its core mission was jeopardized and it became, in the words of one employee, "the Post Office and IRS gone into space" (Vaughan 1996, 211).

The costs and pressures of formal accountability can also have complex indirect consequences. The National Academy of Sciences (NAS 1998, 47), for instance, found that the expense of establishing a bureaucratically defensible basis for its technical claims had sometimes led the FAA to "work around" its own rule-making processes, addressing safety concerns through nonmandatory guidance or draft recommendations instead of binding rules that would carry more force. (The NAS found this outcome to be less than optimal.) The same report also argued that the slowness of the FAA's procedural determinations sometimes led airlines to delay introducing new safety features, putting them off until they could be sure that their investments would be compliant with forthcoming requirements. The NAS was not impressed by this either, commenting: "The public is not well served by a process that pressures operators to delay implementation of safety enhancements" (NAS 1998, 47). The 737-MAX crisis discussed in chapter 11 also might be construed in this light. Many accounts of the airplane's shortcomings attribute them, at least in part, to compromises born of Boeing's desire to avoid expensive recertification (Nicas and Creswell 2019; Tkacik 2019). Absent regulatory oversight, in other words, it is plausible that Boeing would have chosen to make the MAX *more* reliable.

It is important to tread lightly here, as assessment practices can undoubtedly be functional and some administrative costs are clearly worth paying. Johnson (2001), for instance, draws on the Apollo program to make a compelling argument about the important role that accounting mechanisms can play in structuring and integrating complex technological projects. Hastily made decisions are rarely optimal, meanwhile, and—given the importance of design stability to achieving reliability—it might actually be functional for bureaucratic requirements to stymie innovation. Nobody would seriously advocate removing all the tests, analyses, or documentation that govern catastrophic technologies.

At the same time, however, and without belaboring the point, it is difficult to deny that many technology assessment practices can be dysfunctional; and there are good reasons to believe that this is especially the case in catastrophic-technological contexts. Previous chapters of this book have argued that it is impossible to police the reliability required of jetliners—and, by extension, all catastrophic technologies—with formal rules, and such extreme reliabilities are better understood as being governed by incentives. Insofar as this argument holds, then it is easy to imagine that the extravagant audit requirements these technologies invariably accrue might not always be functional. (Or rather that those requirements serve secondary functions—such as making properties like reliability bureaucratically visible—more than their ostensible function of evaluating reliability and safety.) It is certainly arguable that airframers could, and probably would, make equally reliable jetliners at significantly reduced cost if they did not simultaneously have to make that reliability accountable to an audience that values numbers and the appearance of objectivity. After all, they understand their airplanes better than regulators ever could, and the market ensures that they pay a steep price for unreliability.

Even if there are virtues to be found in making the reliability of catastrophic technologies accountable, therefore, the benefits of this accountability are easy to overestimate and its costs easy to overlook.

CONTRADICTIONS

Beyond the literal costs of performing objectivity are a wide range of more subtle complications that arise from contradictions that performances of objectivity inevitably create, between the work of achieving technological ends on one hand, and the work of *being seen to achieve* those ends on the other. An extensive social scientific literature dating back to Merton (1936; 1940) testifies to the perverse outcomes that can arise when rules and measures imperfectly represent the subtleties of organizational practice. This is a central theme of the sociology of accounting literature (e.g., Miller and Napier 1993; Munro 2004; Neyland and Woolgar 2002; Strathern 2000; Hopwood and Miller 1994; Verran 2012), but it is also a recurrent finding in studies of engineering (e.g., Sims 1999; Schulman 1993).

In these studies of engineering specifically, problems are often found to emerge because some of the audiences for whom performances of objectivity are staged—legislators, for instance—play an active (if usually indirect) role in managing the technologies and practices being represented. In these circumstances, the performances can give rise to unrealistic expectations that then influence decisions regarding the technologies and practices themselves. (Recall again, for instance, the misleading expectation that the FAA's survivability researchers should resolve *every* uncertainty before it enforces new airframe requirements.) Problems can also emerge because real engineering practices are difficult to separate cleanly from their stylized portrayal, meaning that organizations responsible for managing technologies either have to make engineering compromises to make their performances work, or they have to grapple with accountability dilemmas created by poor performance.

To more fully understand how white-boxing's contradictions can cause problems, it helps to consider three broad ways in which those problems can manifest.

HIDDEN TALENTS One way of envisaging how the contradictions that white-boxing creates can cause problems is to think about what happens when decision-makers lose sight of the subjectivities of technical practice. Consider, for example, the way performances of objectivity shape lay perceptions of qualitative expertise. As discussed, the way we talk about engineering makes it easy to forget that engineers necessarily exercise unquantifiable judgment in their work. This oversight comes at a cost. Many scholars, for instance, have found that the convention of portraying technical knowledge as objectively quantifiable tends to devalue any unquantifiable (or unquantified) expertise, and to delegitimize claims that must appeal to such expertise for justification (Wynne 1988; Sims 1999; Silbey 2009; Schulman 1993; Langewiesche 1998b; Jasanoff and Wynne 1998).

These scholars invariably attribute meaningful harm to this phenomenon. Perhaps most famously, for example, Vaughan (1996, 2004, 2005) argues that the 1986 space shuttle *Challenger* disaster might have been averted if NASA managers had not deemed the misgivings of experienced engineers "too qualitative" to merit serious consideration. "Critical information was lost," she writes, "because engineers were trained that intuition, hunch, and qualitative data were inadmissible—'not real science,' 'an emotional argument,' 'subjective'—and any engineer who argued on that basis lost status" (Vaughan 2004, 332).[4]

In devaluing qualitative expertise, moreover, white-boxing simultaneously devalues many of the qualities—such as experience, judgment, tacit knowledge, and technical intimacy—on which such expertise relies. It fosters an image of technical experts as diligent implementers of rule-governed formulas, not judicious exercisers of meaningful discretion: masters of the checkbox rather than the black box. And this, in turn, can encourage outsiders to overestimate the interchangeability of experienced personnel and to underestimate the ramifications of moving or losing them (see, e.g., LaPorte and Consolini 1991; Hopkins 2009, 2010; MacKenzie and Spinardi 1996). As one former aviation engineer cautioned, "the usual corporate paradigm that all people are essentially interchangeable is a mistake of huge proportion [in civil aviation]. People are not interchangeable, because they hold the institutional memory. Without continuity, safety is sacrificed."[5]

The devaluing of qualitative expertise can also create problems for initiatives like the FAA's program of delegating technical tasks to Designated Engineering Representatives (DERs), as outlined in chapter 10. The designee system has limitations, as we saw, but it undoubtedly serves a valuable function. Insofar as the FAA is able to effectively engage with design decisions at all, it is because the DER program allows it to access the judgment of insiders. Given that public-facing regulatory discourse downplays the value of such judgment, however, the FAA struggles to articulate the necessity of the program and repeatedly has to defend it against high-level criticism and investigation. Nader and Smith (1994, 169), for example, skeptically observe that the FAA "believes in the honor system for airline compliance." (See also, e.g., NAS 1980; GAO 2004; Schiavo 1997; Mihm 2019; Robison and Newkirk 2019.)

Relatedly, and perhaps most consequentially, white-boxing tends to hide the vulnerability of expert judgment to external influences. This is to say that when audiences fail to appreciate that technological experts are making qualitative choices in their determinations and not simply applying objective rules, then they are more likely to overlook factors that might shape those choices, like conflicts of interest, biases, and the circumstances that encourage them. (In civil aviation specifically, for instance, few observers recognize the relationship outlined in the previous chapter, wherein the industry's economic incentives tend to favor reliability over survivability.) This is reflected in a widespread belief—organizationally entrenched in many catastrophic technological contexts—that safety can be governed independent of wider concerns like economics (Perrow 1999; Hopkins 2009). So it is, for example, that finance managers in catastrophic-technological organizations are often given little mandate to consider the safety implications of their decision-making (Hopkins 2009, 2010). And, conversely, safety managers are rarely given much say in the kinds of financial decisions that shape economic incentives.

PRISONERS TO PERFORMANCE A different set of tensions and contradictions can arise from the practical, real-world difficulty of reconciling engineering's messy qualitative practices with a performance of formal objectivity. We have seen in this discussion that white-boxing technical practices arguably create slack for experts to exercise important qualitative judgments by hiding them. And this might be true in some circumstances, but that space is inevitably limited by what can be effectively hidden.

To invoke the mores of objectivity for this purpose is to exploit the fact that seemingly rigid rules and calculations are more interpretively flexible than is commonly understood. For example, it utilizes the fact that engineers can discretely exercise discretion in how they choose to apply a rule or quantify a variable. Rules and calculations are not infinitely interpretable, however, especially in circumstances where their observance is closely scrutinized. Just as actors in a literal play must monitor their behavior backstage to avoid disrupting the performance onstage, therefore, so technical experts sometimes must make compromises to maintain the performance of rule-governed calculation.

This is to say that technical experts (and the organizations they represent) can find themselves prisoners of their own performances: the requirements of making their work publicly accountable force them to make choices that they privately believe to be suboptimal. A common manifestation of this problem arises when experts find themselves obliged to chase metrics that they know imperfectly represent the phenomena they purport to measure: what Merton (1936, 1940) called "goal displacement." The following interview excerpt, from a UK Ministry of Defence official responsible for avionics quality assurance, neatly illustrates the nature of the dilemma:

> Bid assessment is becoming very formal, and under regulations for government departments it's worse, to the extent—I mean, if we haven't got a valid excuse for why somebody hasn't got a job, then you can take us to court. In fact, my observation is that bidding is not about selecting the winner, it's about justifying why the losers haven't got it. . . . I want to go to certain people because I know

> that they're good at their job, and I have brick walls set up internally to stop that happening.[6]

The constraints that accountability imposes on expert discretion can even manifest in the designs of systems themselves. Shatzberg (1999) and Bugos (1996) both illustrate how different military airframes came to be shaped, sometimes suboptimally, by bureaucratic accountability requirements. The same observation might be made about jetliners. Recall, for example, the civil aviation industry's relationship to redundancy (outlined in chapter 6). Given the essential role redundancy plays in *demonstrating*—as distinct from *achieving*—ultrahigh reliability, it is reasonable to imagine that the FAA's audit requirements might lead airframers to favor it over other, potentially more optimal design approaches. (As we have seen, adding redundant elements is not inevitably the most effective way for engineers to spend a system's limited weight or volume budget to maximize its reliability. A single, double-strength element might confer more reliability than a redundant backup, for instance, even while it made that reliability harder to demonstrate quantitatively.) No airframer would publicly concede that its designs are anything less than optimal, but Boeing's assertion (in chapter 6) that two engines on its 777 were safer than four might plausibly be read as an implicit admission that there was a period where the airplane's four-engine configuration owed more to regulatory compliance than design optimization.

Even when the demands of performing objectivity do not restrict experts' discretion directly, they often do so indirectly. As with the Ministry of Defence official, for example, experts will sometimes act against their own private judgment for fear that their work will be challenged retroactively and deemed noncompliant relative to an unrealistic ideal. This dynamic has long been observed in other contexts. Noting that "following the rules" serves as a useful defense against blame, for example, Rothstein, Huber, and Gaskell (2006, 101) describe a process wherein administrative fictions akin to those involved in white-boxing tend to organically morph over time into meaningful organizational expectations, "colonizing" the practices they were intended to facilitate (see also Luhmann 2005, 13; Rothstein and Downer 2012).

MISPLACED BLAME Concerns about retroactive judgment speak to another dilemma that arises from using performances of objectivity to create backstage spaces: the fact that those performances can collapse, starkly revealing the contradictions they hide and fostering impressions more damaging than those they were designed to assuage. When organizations seek to hide the messiness of real technological practice, they often forget that that messiness has a way of leaking out at inopportune moments. This is because performances of objectivity only go so far. Even the most elaborate attempts to white-box technical work are never inviolable, and—with the possible exception of some military spheres—it usually takes little investigative effort to glimpse the less-rule-governed dimensions of catastrophic technologies. As discussed in chapter 4, for instance, it makes sense to think of civil aviation as having a private discourse, wherein discussions of its more subjective practices are hidden from public scrutiny, but the privacy of that discourse is protected more by recondite bureaucracy and engineering esotericism than by substantive confidentiality measures. Such barriers deter most idle scrutiny, but there are nevertheless occasions where outsiders are motivated to look more closely. The introduction of new systems, for instance, often drives coverage that highlights elements of technical practice—especially disagreements—that might otherwise remain invisible. (Recall the controversies that accompanied the launch of Concorde and the A320 with its fly-by-wire avionics.) By far the most acute drivers of unanticipated scrutiny, however, are accidents. Accidents attract media coverage and demand formal investigations, both of which invariably shine lights on practices and determinations that would usually be private.

Regardless of its motivation, such backstage scrutiny routinely uncovers practices that fail to match the white-boxed public portrayal of those practices, and these contradictions, in turn, frequently lead observers to draw misleading conclusions. Arguably the most favorable, or least unfavorable, of these that observers draw from witnessing such contradictions is to construe them as evidence of organizational errors or shortcomings. Recall again, for instance, the safety concerns expressed by software experts about the A320's avionics. As we saw, these concerns were ultimately misplaced—the A320 has been as reliable as any aircraft in history—but they arose, we might say, because the software community overestimated the objectivity of certification. They understood enough to know that the regulator's formal reliability calculations were implausible but were not familiar enough with broader aviation regulation to recognize that such calculations are *always* insufficient, even in nonsoftware contexts, and were not the true foundation on which expert confidence in the system rested.

More commonly, however, and usually more damagingly, audiences wrongly interpret the contradictions of technical practice as evidence of wrongdoing: rule-breaking, for instance, or some other form of culpable negligence. This is especially common in the context of accident investigations, which feed into long-established disaster narratives that are already predisposed to apportion blame (Jasanoff 1987; Wynne 1988). So it is, as Wynne (1988, 149) observes, that accidents routinely inspire "flips" in public attitudes toward technical organizations, where a misplaced image of competence and control becomes an equally misplaced image of disorder and incompetence. Or, as the National Academy of Sciences has put it: "Following some highly publicized accidents, there is a technically unjustified loss of public confidence" (NAS 1998: 45).

Langewiesche (1998a; 1998b) captures the essence of this dynamic in his account of the 1996 crash of ValuJet Flight 592 in the Florida Everglades. The accident occurred when incorrectly stowed oxygen canisters caught fire in the airplane's cargo hold, creating an inferno that engulfed the main cabin. Investigators and the media largely blamed this on airline ground personnel, who, they argued, failed to follow correct procedures and falsified paperwork when loading the canisters (NTSB 1996). By Langewiesche's account, however, the underlying cause is more reasonably attributed to ambiguities in the procedures themselves, and the decision to place all the blame on the ground personnel speaks to the aforementioned contradictions between engineering practices and their public portrayal. He writes:

> The falsification [that ValuJet's ground personnel] committed was part of a larger deception—the creation of an entire pretend reality that includes unworkable chains of command, unlearnable training programs, unreadable manuals, and the fiction of regulations, checks, and controls. Such pretend realities extend even into the most self-consciously progressive large organizations, . . . The systems work in principle, and usually in practice as well, but the two may have little to do with each other. Paperwork floats free of the ground and obscures the murky workplaces where, in the confusion of real life, . . . accidents are born. (Langewiesche 1998b, 97)

The same dynamic can affect regulators themselves, with the authority born of couching their work in the language of quantitative certainty sometimes fostering an image of incompetence when their assertions fall short. As previously discussed, for instance, regulatory failure is a prominent theme of coverage around the 737-MAX (e.g., Mihm 2019; Robison and Newkirk 2019), and of the official investigation into Deepwater Horizon (NCBP 2011, 56–57).

Again, it is important to tread carefully here. Attributions of negligent behavior are not always unwarranted. There have undoubtedly been many accidents caused by organizations and individuals who should have been more diligent in designing, building, regulating, and operating catastrophic technologies. These actors are not beyond reproach, therefore, and should in no way be immune to charges of culpable shortcomings. Nevertheless, it is also true that charges of wrongdoing sometimes say less about real failings—organizational or individual—than about the way that technical work is portrayed and perceived. As we have seen, civil aviation's remarkable successes, as well as its very occasional failures, are largely born outside the formal rules, and, contrary to the implicit promises of white-boxing, some accidents are not easily attributed to wrongdoing or negligence (especially those that can be described as "rational" or "normal").

When blame is misplaced, moreover, it can be harmful in ways that go beyond injustice, not least because organizations often respond to charges of wrongdoing in ways that are themselves perverted by the same misconstruals of technical practice. Perhaps most straightforwardly, for instance, it is common for experienced personnel to lose their jobs in the wake of accidents—the symbolic sacrifice of high-placed careers having become a ritual of modern technological disaster investigations—and, when unjustified, this has costs that organizations underestimate due to white-boxing's elision of qualitative judgment. It is also common for organizations to respond to misplaced blame in ways that exacerbate white-boxing's contradictions and tensions (which, ironically, foster the misplaced blame). Blind to the less rule-governed dimensions of technical practice, for example, organizations frequently respond to perceived failings by expanding their formal oversight requirements, deepening their performances of objectivity by drafting even stricter, less practicable rules and requirements. (The Deepwater Horizon report, for example, advocates for the creation of "strict [safety] policies requiring rigorous analysis and proof" [NCBP 2011, 126].) It is also common for organizations to reform their practices to make them better match their idealized portrayal of those practices, further constraining experts' discretion by demanding that unrealistic rules and suboptimal requirements are actually observed (NAS 1998; Francis 1993; Sismondo 2010, 143; Rothstein et al. 2006).

12.4 MITIGATED CONSEQUENCES

NO HARM, NO FOUL?

The significance of the burdens and contradictions outlined here varies, but it is generally muted in the context of civil aviation. The direct costs of making jetliners accountable undoubtedly create inefficiencies and have other complex consequences for the industry, as does the tendency to undervalue subjective judgments (and the experts capable of making them). At the end of the day, however, neither phenomenon has prevented the industry from making jetliners ultrareliable.

Much the same could be said about the tendency to misapportion the blame for accidents and respond to that blame in counterproductive ways. Civil aviation has seen injustices and inefficiencies because of this, and—given its propensity to progressively exacerbate white-boxing's contradictions—this dynamic could lead to greater problems in the future. As with the issues described in this chapter, however, the evidence suggests that it has not yet had a huge effect on the industry's safety (in part, no doubt, because its effects are greatly mitigated by the fact that jetliners crash so infrequently).

Probably most worthy of concern is white-boxing's tendency to hide the role that economic incentives play in underpinning the industry's safety. Insofar as the costs of the 737-MAX's failings are mitigated by declining competition between airframers, for instance, this is a meaningful concern. It will take time for this relationship to become fully apparent, however, and it is probably too early to speak definitively about fundamental changes to the industry's incentives.

Civil aviation's practices might not be transparent, in other words, and they might not be perfect, but this should not eclipse its undeniably impressive service record. By any reasonable definition, jetliners are safe to fly and have been for decades. Jetliners are not like other catastrophic technologies, however, and, as chapter 13 will explain, the real harm that arises from white-boxing their safety is that it conceals this difference.

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