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

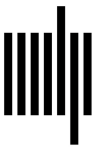

The MIT Press

10 SAFETY COSTS: THE STRUCTURAL FOUNDATIONS OF ULTRARELIABLE DESIGN

> Of the major incentives to improve safety, by far the most compelling is that of economics.
>
> —Jerome Lederer

10.1 AN ORGANIZATIONAL ANOMALY

GOOD BEHAVIOR

If "the story of Concorde was to demonstrate that the age of irrational decision-making was not yet past," as a magazine article put it (Gillman 1977), then the lacuna created by the airplane's retirement—the conspicuous absence of supersonic transport by the third decade of the twenty-first century—arguably speaks to an age of technological maturity, where airframers manage rational accidents by curbing their design ambitions.[1] The preceding chapters of this book have argued that the civil aviation industry transcends the uncertainties of its tests and models, in part, by committing to a common design paradigm and keeping that paradigm stable.

In helping resolve the epistemological problem of civil aviation's extreme reliability, however, the relative stability of its airframe designs poses an organizational problem. As we will see, stability for the airframers is expensive. It means forgoing, or long delaying, tempting innovations like composite materials that promise competitive advantages in a challenging marketplace. And while it is intuitive to see the industry's commitment to design stability as

something driven by a concern for safety, organizations are generally thought to be poor at consistently prioritizing safety over profits.[2]

SAFETY SECOND

The finding that airframers are willing to consistently subordinate economic interests to safety is difficult to reconcile with the literature around organizations and technology. Outside of civil aviation, scholars routinely find that private companies are susceptible to moral hazards arising from incentives to maximize short-term economic gain (e.g., Power 1997), and organizations responsible for catastrophic technological systems are no exception to this rule.

Consider, just by way of illustration, the 2010 Deepwater Horizon disaster. The offshore drilling platform was exceptional even by the lofty standards of drilling platforms. Built by Hyundai at a cost of $350 million, the huge, semi-submersible rig was of a different class to most of its peers, holding records for the depths that it drilled (NCBP 2011: xiii; Transocean 2010). It met its catastrophic end in the Gulf of Mexico on April 20, 2010, while completing an exploratory well for the oil company BP. A blowout caused an explosion that killed 11 of its 126 workers and engulfed it in unquenchable flames. The stricken platform burned fiercely for two days before disappearing ignominiously beneath the waves. In its wake, it left a gushing oil leak at the seabed, 5,000 feet (1,500 meters) below, which took three months and billions of dollars to cap. To date, this environmentally devastating spill remains the largest ever in US waters, having contaminated 1,100 miles of coastline and 68,000 square miles of water (NCBP 2011).

As the scale of the disaster became evident, the White House convened a commission to investigate its causes. The commission's report, published the following year, attributed the accident largely to a corporate culture that encouraged workers to maximize revenue by cutting corners (NCBP 2011). It identified nine separate management decisions—each pertaining to the platform's design or operations—that increased the risk of a blowout, and it highlighted that seven of the nine unambiguously had saved BP money (NCBP 2011, 125). The specific contribution of each decision was ambiguous, the report concluded, but collectively they pointed to a common underling cause: a failure to consistently prioritize safety over short-term profits. Time and again, it found that "[d]ecisionmaking processes . . . did not adequately

ensure that personnel fully considered the risks created by time- and money-saving decisions" (NCBP 2011, 125).

The findings of the Deepwater Commission might be shocking, but they should not be surprising. Indeed, any serious engagement with the literature around technological disasters—from academic ethnographies to formal accident reports—testifies to the regrettable normalcy of organizations resisting expenditures on marginal risk reduction (e.g., Vaughan 1999; Perrow 1999, 2015; LaPorte and Consolini 1991; Perin 2005; Hopkins 2010; Reason 1997; Silbey 2009; NCBP 2011; Sagan 1993). So abundant is this evidence that many social scientists consider it almost axiomatic that the safety of sociotechnical systems suffers when it clashes with profits, especially over time.

The difficulty of prioritizing safety over profit becomes more intuitive if we look at organizational incentive structures. Such structures and their relationship to risk behavior are always complex, as is the literature around them (see, e.g., Vaughan 1999; Gephart et al. 2009), but two generalizable observations are worth highlighting. The first is that organizations, and the individuals staffing them, invariably operate with regard to time frames that are ill suited to managing the kinds of high-consequence and low-probability risks that characterize catastrophic technologies. This is simply to say that—since management positions are only held for a few years at a time, and corporate strategies must cater to investors and quarterly earnings reports—it is difficult to create accountability structures that prioritize the prevention of very rare accidents, no matter how consequential those accidents might be. The second is that when technological accidents occur, the organizations held responsible for them rarely bear the full costs of those accidents. Insurance plays a role in this, as does the fact that courts and legislatures, conscious of jobs and shareholders, almost never penalize large corporations in ways that would undermine their viability.

The upshot of this, as Perrow (1994, 217) observes, is that the costs of accidents are rarely high enough to be more determinative of organizational behavior than wider economic incentives. The 2007–2008 financial crisis probably exemplifies this most clearly. Autopsies of the crisis and its causes routinely highlight the fact that banks rewarded traders—and investors rewarded banks—on the basis of short-term profits rather than long-term stability, and the price that these actors paid for the disaster was wholly eclipsed by the money that most of them accrued in making it possible (e.g., Lowenstein

2011). And although the financial crisis might be exemplary in this regard, it is far from unique. Studies of more classically "technological" accidents routinely come to similar conclusions. BP famously paid record penalties for the Deepwater Horizon incident, for instance, but those penalties still failed to cover the full cost of it. BP remained an extremely viable company (Uhlmann 2020), one that in 2018 reported an annual profit of over $49 billion.

The primacy that scholars ascribe to economic incentives becomes more intuitive if we try to appreciate the subtle ways that incentives can shape behaviors. Most accounts of economic pressures influencing organizational safety behaviors frame the process in terms suggesting moral compromise: deliberate corner-cutting, born of greed or irresponsibility. But such narratives are often misleading.

There is no shortage of greed or irresponsibility in the world, of course, but organizational decision-making is usually better understood in relation to the conditions that structure its reasoning. Consider, for instance, that the division of labor in bureaucracies invariably separates financial expertise from safety expertise, thereby creating circumstances where even well-intentioned economic decisions can have underappreciated safety ramifications. Consider also that the benefits of reliability engineering do not scale very intuitively with their costs. Achieving ultrahigh reliability requires organizations to spend large amounts of money to offset extraordinarily improbable failure modes: one-in-a-billion events that might never occur in the entire lifetime of the system. Such expenditures do not fit neatly into the kinds of economic calculus taught at business schools, and the divisions within bureaucracies help occlude their necessity.

"Safety" is an ambiguous condition, moreover, and, homilies notwithstanding, it can never be something that organizations prioritize absolutely. Even in optimal conditions, it is always being negotiated against rival considerations like cost and functionality. (The safest possible jetliner would be one that rarely left the ground and charged millions of dollars per flight, but few would consider this ideal.) No organization has unlimited resources to spend on risk reduction, in other words, and all must balance the safety of their operations against the ever-present risk of failing to compete (economically or otherwise): what Reason (1997) calls the tension between production and protection (see also Lazonick and O'Sullivan 2000; Froud et al. 2006; Perrow 2015; Bazerman and Watkins 2004, 5–7). Such considerations

require complex and ambiguous judgments, and—when seen in light of the division of expertise, and the counterintuitive costs of reliability—it is easy to imagine how incentives might shape these judgments without presupposing much wrongdoing by the actors involved.

This is all to say that any organization that consistently prioritizes extreme, long-term safety over nearer-term profit—as civil airframers appear to do—deserves scrutiny. The organizations that make jetliners are clearly conscientious; they care about safety and employ many intelligent, earnest, and principled people. Any exposure to the industry speaks to this. But such considerations have limited explanatory value. It is reasonable to assume that most organizations responsible for catastrophic technologies care about safety and employ well-meaning people, but, as we have seen, most still fail to consistently prioritize safety.

There are few reasons to imagine that airframers should be exceptional in this regard. The economic pressures on them are intense. The market in which they operate is competitive and unforgiving; their customers have real options (or at least they have in the past—but this may be changing, as we will see in chapter 11); and the development costs of a new jetliner are such that they practically bet the company's future on the commercial success of new designs (Newhouse 1982; Mowery and Rosenberg 1981). (Lockheed, for example, was forced to exit the commercial jetliner market after its L-1011 Tristar failed to compete effectively with the DC-10.) And close accounts of airframers' circumstances, practices, and cultures, rarely suggest that they float free of the pressures that characterize other organizations (e.g., Schiavo 1997; Perrow 1999, 163–169; Fraher 2014; Barlay 1990; Oberstar and Mica 2008; Heimer 1980). Indeed, experts commonly voice concerns about the potentially corrupting effects of economic pressures on civil aviation's design choices (e.g., *Federal Register* 1998, 68642).

This is why the civil aviation industry's sustained commitment to design stability raises important questions with far-reaching implications for our understanding of organizations more broadly. Resolving these questions will take us on a somewhat circuitous—arguably even circular—journey. The excursion is worthwhile, however, because it sheds light on how the epistemology of ultrahigh reliability shapes the incentive structures around catastrophic technologies, and the role that regulators play in policing those technologies.

Let us start, therefore, by taking some time to further establish that the industry does indeed appear to be acting against its short-term economic incentives in pursuit of safety.

10.2 CHOOSING STABILITY

ECONOMIC SACRIFICE

The problem of why airframers pursue design stability would essentially disappear if stability itself was economically advantageous for them, and there is a credible argument that this is sometimes the case. Certainly, it is not always true that delaying or resisting innovation costs airframers money. Innovation is expensive, after all, and scholars often find that technological designs are kept stable by the costs of altering them (e.g., Musso 2009; Kemp, Schot, and Hoogma 1998; Freeman and Perez 1988; Mokyr 1990). This might be especially true in catastrophic technological contexts like civil aviation, where systems have extreme technical and regulatory requirements. Forgoing innovation in these contexts avoids the considerable development and testing costs that design changes incur, and it allows the industry to take advantage of time-honed efficiencies in manufacturing and servicing.

These savings can be significant, and they undoubtedly incentivize stability in many circumstances, but radical innovation is not always a poor investment, and it is easy to underestimate the pressures on airframers to innovate faster and more ambitiously. Among the many design changes that civil airframers abstained from, or long-delayed introducing, are some that seemingly promised to make their products meaningfully more competitive. Arguably, the clearest examples of these—let us call them "uncompetitive deferments"—involve innovations that promised to increase jetliners' efficiency.

To understand the pressures on airframers to increase the efficiency of their jetliners, and thus the costs of delaying or forgoing efficiency-maximizing innovations, it helps to consider the incentive structures that act on their primary customers: the airlines.

Airlines find it notoriously difficult to make money. The fact that they are highly regulated, together with the related fact that many are valued national assets, have historically limited their ability to consolidate, leading to chronic oversupply on many routes. The resulting competition—often distorted further by subsidies and, in recent decades, intensified by more

transparent internet pricing—has squeezed revenues to a point where many airlines almost perennially hover on the edge of bankruptcy (Vasigh, Flemming, and Humphreys 2015; Gritta, Adrangi, and Davalos 2006). Over the period 1978–2010, the average net annual profit for the world's airlines was less than zero US dollars (–$0.04 billion, to be precise) (Cronrath 2017, 3). "If a capitalist had been present at Kitty Hawk back in the early 1900s, he should have shot Orville Wright," Warren Buffet told *The Telegraph* in 2002, adding, "He would have saved his progeny money" (Vasigh et al. 2015, 264).

In this environment of tight margins and tenuous profits, airlines are strongly incentivized to make hard-nosed economic calculations about airframe purchases. There are many factors at work in such calculations, but operating costs are the most prominent, and even small economies can make an enormous difference to an airline's bottom line. This is especially true in relation to fuel consumption, fuel being a major operating expense for most airlines and its price being their primary cost variable. (In 2012, for example, fuel alone was estimated to account for 33 percent of total airline operating costs [IATA 2013].)

The importance of fuel economy to airlines creates powerful incentives for competing airframers to adopt innovations that promise to improve the fuel efficiency of their jetliners. A design change that even marginally decreases a jetliner's fuel requirements would have to drive up its base price and maintenance costs by a lot before it became uncompetitive with less expensive but less efficient offerings from other manufacturers.

There are many ways of increasing the fuel efficiency of an airframe,[3] but the most straightforward of these is to make it lighter. The relationship between weight and fuel consumption is widely underappreciated. Heavier jetliners require more fuel to move—this is straightforward enough. But it is easy to forget that any additional fuel adds its own weight, which then requires more fuel, which adds more weight, which, at some point, requires structural reinforcements that add even more weight, which requires more fuel, and so on.

This irony of aeronautical engineering means that even small changes to the mass of an airframe can have disproportionate implications for an airline's operating expenses. Take, for instance, Southwest Airlines, which MIT researchers used as the basis for a 2013 study of the industry's fuel costs (Jensen and Yutko 2014). A smaller airline than many, Southwest operated about 1.6 million flights in 2013. According to the MIT model, if every passenger

on each of those flights carried an extra cell phone, the combined weight of those phones would have cost the airline $1.2 million in fuel over the course of the year—a number would increase to $21.6 million if the passengers carried laptops instead (Jensen and Yutko 2014). (Understanding this relationship helps explain a lot about the behavior of airlines, from the growing trend of charging for baggage, to the decision by some carriers to stop offering free bottles of water. It is partly why American Airlines replaced pilots' emergency binders with iPads—a measure that it anticipated would save it $1.2 million per year—and why All Nippon Airways started requesting that passengers visit the bathroom before boarding [Brown 2009; Jensen and Yutko 2014]).

Such figures offer an important perspective on civil aviation's adoption of innovations like fly-by-wire and advanced composite materials. Both these adaptations promised to make airframes significantly more efficient, primarily by reducing weight. And, once embraced, they lived up to this promise. The fly-by-wire A320 was claimed to be 40 percent more fuel efficient than the B727, with which it was designed to compete (Beatson 1989). The composite B787, meanwhile, was designed to consume a fifth less fuel per passenger mile than the B767 it was intended to replace (Marsh 2009, 16–17; Waltz 2006). These exact figures are contestable, and in each case there were multiple factors that contributed to the airframes' efficiency gains, but both aircraft were undoubtedly more efficient, and fly-by-wire and composites played important roles in these achievements. Even though both innovations were costly to implement, therefore, they promised significant efficiencies for customers that had averaged negative returns over most of their history.[4]

For airframers eager to offer competitive products, the efficiencies offered by fly-by-wire and composites—both clearly viable innovations that had long been employed in other aviation contexts—must have loomed large against abstract and invisible misgivings about reliability. It would have been economically rational for executives to push the earlier adoption of both in efforts to sell more jetliners than their rivals. At least in respect to these design choices, therefore, it is fair to say that keeping airframe architectures stable—changing them very slowly and long delaying innovations until they had been validated exhaustively in other spheres—came at a cost. It is difficult to interpret this behavior except in terms of a consistent ability to prioritize long-term safety objectives over shorter-term economic incentives.

Whether this ability reflects any real restraint on behalf of the airframers is another question, however, as it is probably more intuitive to imagine design stability being imposed on airframers by regulators. Airframers are kept from innovating too aggressively, we might suppose, by the FAA and its elaborate code of certification requirements. After all, policing the industry is ostensibly one of aviation regulators' primary functions.

As we will see, however, this too is problematic.

10.3 REGULATING STABILITY

AN INTUITIVE NOTION

The idea of aviation regulators policing the stability of jetliner architectures is undoubtedly appealing. Preceding chapters of this book outlined in detail why formal rules and metrics, with their many ambiguities and interpretive flexibilities, do not allow experts to monitor and enforce ultrahigh reliabilities directly. (As we have seen, a jetliner could meet every standard, pass every test, and still be unsafe to fly.) Nevertheless, it might still be possible for regulators to monitor and enforce practices like design stability. Innovation scholars sometimes speak of "technology-forcing" regulations, designed to incentivize actors to take design risks in pursuit of newer and better technologies (e.g., Gerard and Lave 2005; Lee et al. 2010). In the context of catastrophic technologies, by contrast, it arguably makes sense to think of "technology-halting" regulations, intended to promote design stability in pursuit of safety.

There is evidence to support this idea. Type certification is ostensibly framed around quantitative reliability targets as a way of avoiding innovation-constraining design stipulations, but it clearly rewards design stability. If the regulator finds a proposed system to have insignificant differences compared to a previously certified design, for instance, then the rules largely (or wholly) allow it to retain its original certification. If it deems a system to be unusually innovative, by contrast, the rules mandate extra scrutiny, often requiring the formulation of an "issue paper" to identify concerns and establish additional assessment protocols (FAA 1982, appx 1; NTSB 2006b). The effect of such measures is to make it significantly easier and cheaper for manufacturers to gain approval for systems that are identical to, or strongly derived from, previously certified designs.

Still, however, direct regulation offers an unsatisfying explanation for the industry's design stability.

DIRECT OVERSIGHT

To understand the FAA's ability (or rather their inability) to police the stability of jetliner architectures, it helps to think about the fundamental nature of the task. Innovativeness is a fundamentally ambiguous and interpretive property of artifacts that, like reliability itself, is impossible to formalize or objectively quantify. Assessing or controlling the stability of a system would require regulators to assess its similarity to previous systems. As we have seen, however, the question of when two things are the same—be they redundant elements, or test and real-world environments—raises extremely complex relevance questions. And navigating such questions necessarily involves a great number of consequential but qualitative technical judgments, which regulators are ill equipped to make.

There are several reasons to doubt the FAA's ability to directly make the kinds of judgments required to effectively police the innovativeness of new systems. The most straightforward is simply that it lacks the necessary resources to grapple with the nuances of new systems. It is difficult to say with precision how many technical people the FAA employs to assist directly in certification activities, but as of 2006 it was 250 (NTSB 2006b, 68), and the agency has not transformed dramatically since then. Whatever the exact number, however, there is no question that it is far fewer than would be needed to actively police the innovativeness of every system in a new airframe, especially given the industry's increasing reliance on subcontractors (FAA 2008a; Bonnín Roca et al. 2017).[5] In 2019, the FAA's acting administrator told Congress that his agency would need at least 10,000 additional employees to directly perform all its certification duties (Shepardson 2019).

Even if afforded infinite resources, moreover, there is good reason to believe that the FAA would lack the technical intimacy needed to make the requisite technical judgments. A range of official analyses have testified to this deficit. As early as 1988, for example, the Office of Technology Assessment (OTA) was reporting that FAA regulators lacked the expertise to make sound technological judgments (OTA 1988), a conclusion that the US Government Accountability Office (GAO) echoed in 1993, when it found the agency to be "not sufficiently familiar with [specific systems] to provide meaningful inputs to the testing requirements or to verify compliance with

regulatory standards" (GAO 1993, 19). The reports' findings echo those of the Deepwater Commission, which similarly found the platform's regulator lacked sufficient personnel and expertise to keep pace with the industry's technological developments (NCBP 2011, 56–74). They also echo many academic studies of technological practice, which routinely find that the understanding needed to navigate complex technical judgments is born as much from close familiarity as from formal analysis or abstract study (e.g., Collins 1982, 1985, 2001, 2010; MacKenzie and Spinardi 1996). Put in STS terms, we might say that there is an inevitable imbalance—or what Spinardi (2019) calls an "asymmetry"—between the expertise of regulators and that of airframers, wherein the former lack the "tacit knowledge"[6] to make the judgments on which key questions about design stability would hinge (Downer 2009a; 2010).

For both these reasons—a deficit of resources and of intimacy—the idea of external regulators effectively policing airframers' technical judgments to enforce appropriate levels of design stability is almost certainly unrealistic. As one aeronautical engineer succinctly put it, "[T]here is not a way for a third-party organization to assess our understanding of [complex avionics] . . . The very best method we have of discriminating between those who can, and those who can't but talk a good game, are their peers."[7]

The FAA actually concurs. This is why it has a longstanding practice of deputizing engineers from within the industry to act as its surrogates and make technical determinations on its behalf. Herein, therefore, lies a more subtle mechanism by which the regulator might be policing the stability of civil airframes: not directly, via assessments of the technology itself, but indirectly, via assessments of the personnel and organizations that design it.

As we will see, this mechanism is also insufficient for policing design stability, but it is worth exploring nevertheless.

SECOND-ORDER OVERSIGHT

The FAA's tradition of deputizing engineers to act on its behalf has roots in the earliest days of certification. "We've got certain safety factors, and we'll have our engineers check your plans with respect to them," William MacCracken, the man charged with framing the first regulations, told manufacturers, "but mainly we'll rely on you to comply voluntarily" (quoted in Komons 1978, 98). Today, this relationship is formalized in what the FAA calls the "designee program." The program employs a variety of designees

across various roles (FAA 2005b), but the group that plays the most prominent role in its assessment efforts is referred to as "Designated Engineering Representatives (DERs)" (FAA 2005a).

In many ways, DERs are the backbone of the certification process, overseeing tests, calculations, and designs to ensure that systems are compliant with regulations. In this capacity, they answer to the FAA, but most are also employees of the manufacturers (although a small number of consultant DERs work for third parties or independently). Most have about fifteen to twenty years of experience and hold key technical positions working on the systems they assess.

The responsibilities afforded to designees have expanded considerably over time. When the program began, it was intended that they would conduct well-defined tasks, allowing regulators to concentrate on larger oversight functions: framing requirements and analytical criteria, designing tests, and making final compliance determinations. As jetliners became more complex, however, the roles originally demarcated for regulators became increasingly untenable. As a result, DERs may now be authorized to assume almost all key oversight functions, to the point where it is now common for them to both frame and witness tests on the regulator's behalf (FAA 2007, 44; GAO 2004; NAS 1980).

Periodic investigations of the certification process illustrate this growth of DER responsibilities. In 1989, for instance, an internal FAA review concluded that the regulator had been forced to delegate practically all the certification work on the 747–400's new flight-management system because its staff "were not sufficiently familiar with the system to provide meaningful inputs to the testing requirements or to verify compliance with the regulatory standards" (AIAA 1989, 49). Four years later, the GAO reported that the FAA was increasingly relinquishing roles that it traditionally retained, concluding that between 90 and 95 percent of all regulatory activities were being delegated, including many "core" functions such as the framing of standards (GAO 1993, 17–22). In keeping with this shift, the ratio of designees to core regulatory personnel has changed dramatically since the 1970s. Between 1980 and 1992, for instance, the number of designees overseen by the FAA's two main branches rose 330 percent, while core regulatory personnel rose only 31 percent, bringing the overall ratio of designees to regulators from about 3-to-1 in March 1980 to 11-to-1 in 1992. (GAO 1993, 17–19). This trend has continued since the 1990s. By 2006, for example, the NTSB (2006b, 68) was

reporting that FAA Certification Offices employed 250 personnel while drawing on 4,600 designees: a ratio of over 18-to-1.

The designees offer the FAA a level of hands-on tacit knowledge and technical intimacy that its regulators, as outside observers, cannot match (NAS 1980, 7; Fanfalone 2003). Significantly, they also give its regulators access to civil aviation's social economy and reputational landscape: its rumors, hearsay, and culture. This is important because it means that even if regulators struggle to judge the significance of innovations directly, they are still likely to notice if engineers who can make those judgements have significant misgivings about a change, even amid the background noise of normal engineering dialogue and dissent. Construed in STS terms, we might say that the FAA uses the designee system to access what Collins (1981, 1985, 1988) would call the "core set" of aviation engineering: the narrow community of informed specialists who actively participate in the resolution of technical controversies. Regulators might lack the expertise required to actively participate in technical debates—what Collins and Evans (2002) call "contributory expertise"—but possess the competence necessary to understand what it means to be an expert participant, as well as the familiarity required to engage with such experts ("referred" and "interactional" expertise, in Collins and Evans's terminology). All the regulator has to do in these circumstances is assess the creditworthiness of the designees themselves.

Understood in this way, we might say that the DER program helps the FAA manage the unavoidable limitations of its knowledge by substituting an intractable technical problem for a much more tractable social one. Recognizing its limited ability to make complex engineering judgments directly, the FAA instead judges the actors who are capable of making technical judgments: a process that I have elsewhere referred to as "second-order" oversight (Downer 2010). The idea of the FAA actively policing design stability seems more plausible when construed this way. To the extent that regulators can draw on the insight of their designees, then they might be well positioned to make effective technical rulings about the significance of different design changes.

It is worth noting that external reviewers of the FAA's practices have endorsed its delegation program on these grounds (e.g., NAS 1998). It is also worth noting that many other technological domains have come to similar arrangements. It is actually very common for organizations producing high-risk technologies to play an active role in their own regulation, "if only

because they alone possess sufficient technical knowledge to do so," as Perrow (1984, 267) puts it. Nuclear regulation operates on similar principles, for instance, as does railway regulation (Perin 2005; Hutter 2001).

While second-order oversight probably does serve a valuable function, however, it still offers an unsatisfactory explanation for the industry's enduring ability to subordinate profits to safety. This is because the FAA and its delegates are almost certainly too close to the industry that they regulate to be expected to steer it, effectively and consistently, in directions that it does not wish to travel. To invoke a slightly loaded and misleading term, we might say that the FAA and its designees are both highly vulnerable to "regulatory capture."

CAPTURE

First gaining traction in the 1970s (e.g., Peltzman 1976; Posner 1971, 1974, 1975; Stigler 1971), the concept of "regulatory capture" describes a process whereby powerful institutional actors come to control, influence, or otherwise dominate the bodies charged with regulating them. Capture "puts the gamekeeper in league with the poacher," so to speak, by creating circumstances where organizations can pursue their self-interest in ways that regulators are expected to curb on the public's behalf; or even, on occasion, circumstances where they can use regulation to further their own ends at the public's cost (Wiley 1986, 713).

Critics of the FAA frequently describe it as being captured by the airframers that it is supposed to police (e.g., Stimpson and McCabe 2008; Dana and Koniak 1999; Niles 2002; Nader and Smith 1994; Schiavo 1997; Perrow 1999; Fraher 2014; Oberstar and Mica 2008). Such assertions usually point to entwined interests and sympathies operating at varying levels of abstraction (e.g., Bó 2006). On the level of individual sympathies, for instance, critics highlight the fact that FAA personnel often remain in specific regional offices for years at a time and often become close with the people they oversee. This is by design. Fostering the kinds of interactional expertise outlined here requires that regulators develop close, long-term working relationships with their charges. But effective regulation is traditionally thought to hinge on the maintenance of emotional distance between agents of each party, and regulators in many other spheres are regularly rotated through positions and locations to mitigate capture risks.

On a different level of abstraction, critics often point out that aviation regulators and airframers share broad, structural-level interests and incentives. The FAA is a US government agency that certifies jetliners built by US corporations. To date, this primarily means Boeing, a significant national economic asset, and a major defense contractor with a powerful lobby in Washington and beyond. It does not require much cynicism to imagine that such considerations have consequences, especially insofar as they pertain to questions that might jeopardize the company's future, as the viability of a new jetliner might. (It is also worth noting in this context that the FAA was originally established to both regulate and promote the industry. It held this dual mandate until 1996, when it was amended in the wake of a high-profile accident, via legislation that changed the regulator's mission from "promoting" to "encouraging" the industry, without requiring any changes to its "organization or functions" [Mihm 2019]).

To understand these relationships and the problem that they pose, it is important to remember that incentives can shape judgments without implying moral failings. "Capture" is a loaded term. It is often portrayed as an almost conspiratorial process wherein regulators knowingly subvert clearly defined rules. Such associations often misconstrue the nature of the problem, however, and they run contrary to any real exposure to regulatory personnel, who invariably appear professional and sincere. Rather than seeing capture as fundamentally conspiratorial, it is more productive in this context to think of it as a subtle process wherein aligned sympathies and shared worldviews come to colonize irrevocably ambiguous rules and interpretations—a dynamic less analogous to a poacher and gamekeeper, we might say, than to an honest figure-skating judge who has a close relationship with the skater being judged.

This is to say that capture needs to be understood in relation to the interpretive flexibility of technical decision-making, wherein crucial regulatory questions hinge on subjective judgments about everything from the representativeness of artificial birds to the similarity of redundant hydraulic systems. There is extensive evidence that subjective judgments inevitably come to be shaped by structural interests. Scholars since Karl Marx have argued that interests tend to invisibly colonize cultures and interpretations, a finding supported by modern psychology (e.g., Kahneman 2011), by historical evidence that engineers routinely underestimate risks of their own designs

(Petroski 1992a), and by the STS literature, which has long held that interests and predispositions permeate even the most rigorous knowledge claims (e.g., Bloor 1976). To imagine that FAA regulatory determinations were an exception to this rule would be to defy generations of social research. It would also go against explicit concerns expressed by secretary of transportation Mary Peters (in testimony to Congress), and by the FAA itself (in internal memos and other semiprivate communications), about a culture of excessive "coziness" between regulators and airframers (e.g., in Oberstar and Mica 2008, 12; Mihm 2019).

None of this should be read as suggesting that regulators play *no* role in shaping new airframes and promoting reliable design. Regulators are in constant dialogue with airframers, which do not always get their way. (When Niles [2002, 384] quotes an FAA veteran as saying: "To tell the truth, the industry, they really own the FAA," that is probably an exaggeration.) Neither should it be read as suggesting that the FAA is exceptional in being captured by the industry that it regulates. Its relationship with airframers is typical of catastrophic technology regulators more broadly (see, e.g., Perrow 2015).[8]

It is important to recognize, however, that the regulator's abilities and agency are limited, and that, insofar as civil aviation exhibits an extraordinary organizational commitment to safety, regulatory oversight is not a credible explanation for that commitment. Indeed, a senior FAA official conceded as much at the Flight Safety Foundation's 1990 annual International Air Safety Seminar. "The FAA does not and cannot serve as a guarantor of aviation safety," he told the audience. "The responsibility for safe design, operation and maintenance rests primarily and ultimately with each manufacturer and each airline" (Nader and Smith 1994, 157).[9]

To return to the central theme of this chapter, this still leaves us with a problem regarding the behavior of the aviation industry. Airframers seem to be making choices that consistently prioritize long-term reliability over short-term economic incentives. This behavior runs contrary to widely held sociological expectations about the nature of organizations, and it seems unlikely that regulation could be its primary explanation. What is it, then, that makes civil aviation different?

The answer, I suggest, lies in a factor that makes the industry exceptional in so many other ways: the unique volume at which it operates.

10.4 COSTS REVISITED

ACCIDENTS AND INCENTIVES

At the start of this chapter, I argued that most organizations are incentivized to underinvest in avoiding catastrophic technological accidents because they rarely bear the full economic costs of those accidents, and because they operate on time frames that are difficult to reconcile with extremely infrequent events. Some of these conditions apply very clearly to civil aviation. Airframers are rarely penalized directly for design-based technological failures, for example, not least because it is difficult to sue them for designs the FAA has formally certified as reliable. (For its part, the FAA itself cannot be sued for perceived shortfalls in its oversight practices [see Lagoni 2007, 247].) At the same time, however, there are good reasons to imagine that civil aviation is exposed to accidents in a way that sets it apart from other catastrophic-technological spheres.

Simply put, civil aviation—compared to other catastrophic technological spheres—might be expected to have a unique relationship to accidents for the same reason that it has a unique relationship to epistemology: its service experience. As outlined in chapter 3's discussion of the aviation paradox, civil aviation works on a radically different scale than its technological peers. With tens of thousands of jetliners operating simultaneously, they accrue service hours much faster than any other catastrophic technology. This service experience has many ramifications regarding things like recursive practice, as we have seen, but one hitherto unexplored consequence lies in the way it shapes the industry's structural incentives.

If every jetliner sold in the same numbers and operated on the same schedule as Concorde, then, practically speaking, it might be cost rational for the industry to compromise on reliability by embracing innovations that promised to make their products more competitive in the near term. In this scenario, the infrequently flying jetliners would accrue service slowly, so their designs could be much less reliable than modern airframes and decades might still pass before a manufacturer saw its first catastrophic accident (or certainly before any generalizable design shortcomings became statistically demonstrable). By that time, the individuals who designed, assessed, and approved the fallen jetliner would likely have retired or moved on in their careers. The airframer, in turn, would likely be selling a substantially

changed product, which it could plausibly claim to be free of the same shortcomings. If history is any guide, moreover, then the fines that it incurred from the accident would not be especially punitive. The same goes for any reputational damage, which would likely benefit from narratives—in newspapers, legislatures, and courts—that contextualized the accident in relation to the jetliner's preceding decades of safe service and its differences from current offerings.

Most jetliners sell in much greater numbers than Concorde, however, and operate on much more demanding schedules. This changes everything, as it means that they accrue service hours at rates that reveal any reliability shortcomings quickly. With a huge number of flights every year, it is likely that whatever can go wrong with a jetliner, will, and in a time frame that affects airframers in their present form, managers in their current positions, and products in their current incarnation.

In these circumstances, therefore, catastrophic failures have real and direct consequences, both for the airframers as organizations and for the individuals who staff them. For the airframers, these consequences can be financially devastating, not because of fines or penalties, but because the market in which they operate is unusually sensitive to public confidence, and public confidence is sensitive to the absolute frequency of accidents. As outlined previously, airframers operate in a competitive market where sales hinge on the decisions of customers (the airlines), with real choices and fickle customers of their own. In these circumstances, unreliability, or even the appearance of it, can have severe financial ramifications regardless of any formal liabilities (e.g., Cobb and Primo 2003, 5).

The conditions described here radically shape the incentives around reliability. If a jetliner is less than ultrareliable by design, then—because of the industry's high operating volume—this is likely to manifest in accidents before its airframer has stopped taking orders for it. If those accidents give that jetliner (or, almost as likely, its manufacturer) a poor reputation, passengers will start avoiding certain aircraft when booking flights, which will put pressure on the airlines to cancel orders for those aircraft and to look elsewhere for future purchases.

The reality and significance of such confidence spirals in civil aviation are amply illustrated by the industry's history, which is punctuated with once-illustrious manufacturers brought low by tarnished safety reputations. Take, for example, the two 1954 Comet crashes discussed in chapter 7. De

Havilland revised the Comet's airframe in the wake of the accidents' investigation, resolving the underlying fatigue issue. And, following this work, there were no principled reasons to suspect that the Comet would be any less reliable than its peers (military variants of the airframe would go on to serve for six decades). Commercially, however, neither airframe nor airframer would ever recover. The accidents irreparably damaged the Comet's reputation, effectively ending the pioneering British manufacturer's bid to be a major player in the emerging industry.

In many ways, De Havilland's collapse was the opening act of the jet age, and the industry took heed. Perhaps even more salutary, however, was McDonnell Douglas's experience with the DC-10 almost two decades later.

The DC-10—an impressive airplane in many ways, but one that was famously designed to a tight budget—was born with a design weakness. It involved the door to the cargo hold, and, we might reasonably say, probably owed more to economically driven compromises than to the inherent epistemological ambiguities of complex systems. The cargo doors to most jetliners open inward. This makes them more secure, as they can be shaped to act like a plug when the airplane is pressurized from the inside. Inward-opening doors take up valuable cargo space, however, so the DC-10 was designed to open outward. This decision had the effect of making the door's locking mechanism safety critical, but its designers did little to treat it as such. The mechanism lacked redundancy and was designed in a way that sometimes allowed it to appear securely locked when closed improperly, and even register as such in the cockpit.

The dangers of this design became apparent in June 1972—just nine months after the DC-10 entered service—when the cargo door of American Airlines Flight 96 blew out shortly after it took off from Detroit (NTSB 1973). The blowout caused the floor of the jetliner to partially collapse into the cargo bay below, but the pilots were able to land without casualties. It was a close call, however, and per the NTSB's recommendations, McDonnell Douglas redesigned the locking mechanism and added vents to keep the cargo bay from decompressing explosively. Neither airframer nor regulator mandated changes to aircraft already in service, however, and the redesigns proved insufficient to remedy the underlying problem (AAIB 1976). This became tragically evident two years later, in March 1974, when an almost identical failure struck Turkish Airlines Flight 981, shortly after it left Paris. On this occasion, the aircraft fared less well. The floor collapsed in a way that damaged its control

cables, and it crashed into a forest with the loss of all 346 on board (AAIB 1976).

In the DC-10, we might say, McDonnell Douglas committed two meaningful transgressions. The first was to prioritize the DC-10's competitiveness over its safety by breaking with the standard airframe paradigm in a way that traded technical certainty for economic advantage. The second, more cardinal, sin was a lapse of recursive practice: it failed to fully grasp the nettle of its error, and, as a result, lost a second aircraft to the same cause. The airframer paid a steep price for these lapses. It quickly remedied the issue after Flight 981—the DC-10 has never failed in the same way since—but the damage was done. The controversy weighed heavily on public perceptions of the DC-10, which gained a reputation for unreliability. This reputation was arguably unjust—as of 2008, the DC-10's lifetime safety record was comparable to other jetliners of its generation (Boeing 2010; Endres 1998, 109)—but it shaped the coverage of later incidents,[10] and travelers took notice. Bookings on the DC-10 began to suffer, and the airlines responded. TWA put out full-page advertisements to reassure potential customers that it owned none of the aircraft. American Airlines could not make the same claim, but it ran campaigns stressing that it serviced certain routes exclusively with Boeings. Embroiled in controversy, the DC-10 tanked as a marketable product, with airlines across the world canceling their purchase options. McDonnell Douglas might have weathered the storm, but the airplane's tattered reputation became a millstone. It dropped the "DC" designation for its next jetliner—the MD-11—in an effort to distance it from its predecessor, but it failed to sell, despite being highly regarded by engineers. Saddled with debt from the airplane's development, the historied McDonnell Douglas Corporation faltered and was eventually subsumed by Boeing (Newhouse 1982; Davies and Birtles 1999; Waddington 2000; Endres 1998).

SELF-INTEREST

Herein, therefore, lies a plausible solution to the puzzling sociology of civil aviation's design choices: it has unusual incentives, which happen to align with reliability, rather than an unusual relationship *to* those incentives. In an industry that is so aware of its own history, it would be strange indeed if the fates of the Comet and DC-10—and those of their manufacturers—did not weigh heavily on the decision-making of other airframers, pulling against shorter-term incentives to maximize profit by innovating in ways

that compromised reliability. Or, at the very least, we can imagine that these fates acted as a form of "natural selection," wherein only airframers that prioritized reliability (and thus design stability) survived for long.

A succession of formal investigations into the FAA have come to an equivalent conclusion: repeatedly finding that concerns about its regulation practices—regarding capture or delegation, for example—are moot because of the industry's financial incentives keep manufacturers in line. As early as 1980, for instance, the National Academy of Sciences (NAS) was highlighting the importance of the self-interest of the manufacturer in designing a safe, reliable aircraft that would not expose them to lost sales and litigation from high profile failures (NAS 1980); a view the GAO (2004) echoed over twenty years later.

As we have seen at length, the choices airframers have made also seem to bear this out. When it comes to the reliability of airframes at least, it is difficult to fault them. The last half-century of civil aviation produced many jetliner accidents, simply as a function of its enormous volume, but it is extremely rare that jetliners have failed in ways that might be attributed to an overeagerness to innovate.

Rare does not mean unheard of, however, and somewhat inconveniently for the argument in this chapter, two Boeing 737-MAXs crashed during its drafting, embroiling the company in a crisis seemingly born of ill-managed innovation and short-term economic pressure. Chapter 11 will address that crisis, its causes, and its implications. It will then explore the industry's relationship to a property known as "crash survivability," which, I will argue, offers further evidence that economic interests, more than oversight or good intentions, drive jetliner design choices.

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