Whole World on Fire

Organizations, Knowledge, and Nuclear Weapons Devastation

Lynn Eden

CORNELL UNIVERSITY PRESS ■ *Ithaca and London*

Conclusion: Routine Surprises

W*hole World on Fire* is a history of the knowledge and knowledge-laden organizational routines made—and not made—about nuclear weapons damage. Four interwoven processes compose this history. First, military organizations concerned with nuclear war planning and target intelligence focused on increasing knowledge about nuclear blast damage. Second, these same organizations paid little attention to predicting nuclear fire damage. Third, this disparity in knowledge was self-reinforcing: The contrast between the military's ability to predict blast damage and its inability to predict fire damage confirmed for many the inherent unpredictability of fire damage. The failure by U.S. civil defense-funded researchers to predict fire damage further confirmed its unpredictability. Finally, despite this disparity in attention and knowledge, the seeds of organizational innovation lay within, and a significant capacity to predict fire damage was developed, although not adopted at the time it was proposed.

Organizational attention to blast damage predated by decades the invention of nuclear weapons. The U.S. air force (then the Army Air Corps) developed doctrine in the 1920s and '30s that emphasized the destruction of an enemy's ability and will to fight by the precision bombing of specified targets. Such operations would be conducted with high-explosive, not incendiary, bombs. This initial setting of ends and means—which entailed the building of an arsenal of high-explosive bombs, planes configured to carry such bombs, and training centered on handling such bombs—structured for many years to come the problems that airmen and others would work to solve. In other words, they worked within an organizational "blast damage frame" that shaped how problems were represented and the constraints and requirements placed on possible solutions. During World War II, mathematicians, structural engineers, and operations analysts working within the blast damage frame greatly increased organizational capabilities to predict such damage. This new knowledge drew on the existing knowledge base of professionals and researchers and also significantly added to their expertise.

The invention of the atomic bomb and the extraordinary blast and fire damage wreaked on Hiroshima and Nagasaki in 1945 did not disrupt the pre-atomic dynamics. After the bomb, organizational goals remained concentrated on the destruction of specific targets, and government analysts, many of whom had served in the war, continued to understand blast damage as more certain, hence more predictable, than fire damage. In short, analysts "saw" atomic weapons as "blast weapons."[1] In the early post-war period, analysts in air target intelligence worked with university-based research engineers and others to develop systematic predictions of atomic blast damage. By 1954, the government had codified damage prediction in the knowledge-laden organizational routine known as the VN system. By the late 1950s, analysts were able to take into account the much higher yields of the hydrogen bomb in the VNTK system.

The continuities in goals, personnel, and problem focus did not preclude significant innovation. If the broad problem of predicting blast damage carried over from the war, the specifics did not. The new scale of war planning and the new physics of atomic and hydrogen bombs presented great challenges: making a compendium of targets, carrying out experiments in nuclear weapons tests, devising new routines to predict nuclear blast damage. Meeting these challenges required bureaucratic creativity, a significant mustering of resources, large-scale coordination, and the creativity and hard work of individuals. Some of the people involved were little known, some famous. This is, however, not an actor-centered story; always these individuals worked within organizations that provided the context for their activities.

In the same period, a second, inverse, historical process unfolded. Within the military, incendiary bomb damage prediction received no attention before World War II, relatively little during the war, and virtually none after. U.S. air doctrine before the war had led to an arsenal with almost no incendiaries, but as World War II proceeded, some incendiary bombs were developed and incorporated into selected operations. A "fire damage frame" developed alongside the blast damage frame, but it involved many fewer resources and people. Fire protection engineers, a different professional group from those who predicted blast damage, were mobilized to predict fire damage. Working mainly with the British and drawing on knowledge based in fire insurance and fire fighting, they were able to develop effective practical predictions of incendiary bomb damage. However, this knowledge was not as important to the U.S. war effort, nor was it codified.

After the war, the disparity became even greater. The same government organization, Physical Vulnerability in Air Force Intelligence, that developed predictions of atomic blast damage funded only a single study to predict fire damage. Despite the practical predictions made during the war, the study, by fire protection engineer John Wolverton, failed to generalize wartime findings. Fire protection engineers were not university-based; they were not well connected to physicists deeply knowledgeable about nuclear weapons effects; they did not have access to computers; and they did not use numerical computational techniques to model the fire environment.

Thus, organizational determination to predict blast damage during World War II led to the organizational capacity to do so, which provided the basis for building more capacity after the war. In contrast, lack of attention to the prediction of fire damage led to the allocation of fewer resources. Even less attention after the war resulted in a total incapacity to predict nuclear fire damage: no recognized experts, no manuals, no knowledge-laden organizational routines.

Third, the disparity in knowledge about blast damage and fire damage was self-reinforcing. The increasing contrast between the ability to predict nuclear blast damage and the inability to predict nuclear fire damage contributed to the sense that fire damage was inherently unpredictable. Analysts read the past and present state of knowledge as a confirmation of the state of nature.

In addition, the disparity in attention and resources that led to the understanding that blast damage could be predicted and that fire damage could not was further reinforced by the very resources that *were* devoted to fire damage. This seeming contradiction is explicable. With a few exceptions, fire studies were funded primarily by U.S. civil defense. Unlike nuclear war planning, civil defense planning was not an operational undertaking. If nuclear war could not be won, it could be "fought": Forces could be launched with devastating consequences. However, civil defense—evacuation, protection, and recovery from nuclear war with the Soviet Union—simply could not be done. Instead, the purpose of research was to provide the *appearance* of research; in sociologist Lee Clarke's terms, to provide the appearance of "solving problems for which there are no solutions."[2]

Because civil defense research on fire was not part of a highly focused operational effort, the research was diffuse. Civil defense funded many projects on many topics by many contractors; there was no penalty when methods and findings were inconsistent. Further, fire researchers—generally trained as chemical engineers, fire protection engineers, or foresters—developed their ideas primarily from their knowledge of urban building fires and forest fires. As a result, they had neither the intuition nor the disciplinary tools to model the large regional atmospheric flows that drive nuclear mass fires. Although the fire research community was distant from the war planning process, the inconclusive results of their research reinforced war planners' sense that fire damage could not be predicted.

All three historical processes together exhibit what social scientists call "path dependence," in which the effects of learning, high research costs, interdependence, and self-reinforcing expectations reinforce choices already made. For blast damage, expert knowledge was encoded into routines that continually built more organizational capacity to predict blast damage. For fire damage, expert knowledge was not translated into organizational routines, and predictive capacity was not built. Blast damage prediction, once woven into the routines used in nuclear war planning, did not prevent change, but it meant that the initial investment would sharply constrain any future change. Perhaps most significant, participants made sense of the increasing divergence of organizational capability to predict blast damage and incapability to predict fire damage by understanding it to demonstrate that fire damage could not, in fact, be predicted. Thus, the choice made to solve problems of blast damage prediction but not fire damage prediction seemed to be based not on prior choices but grounded in nature itself. This became, in turn, a self-fulfilling prophecy[3]

These three interrelated historical processes could, in theory, provide a full account of why predictions of nuclear blast damage were developed, why predictions of nuclear fire damage were not, and why the disparity has persisted for over a half century. Indeed, had this account been written at any time from the mid-1950s to the early 1980s, this would look like a classic story of organizational inertia, or "lock-in," in which choices once made are inflexible and cannot be reversed. Frequently cited examples are the lock-in of inferior choices: the adoption of the QWERTY keyboard (named after the top left row of letters) over the apparently more efficient Dvorak keyboard (named after its inventor) and of the VHS videotape standard over the reputedly better Betamax standard. A different kind of lock-in example is the decision to launch the *Challenger* space shuttle. In sociologist Diane Vaughan's words, "Socially organized and history-dependent, it is unlikely that the decision they reached could have been otherwise."[4]

Yet, as we saw—and this is the fourth historical strand—by the early 1980s, physicist Harold Brode and several collaborators had begun research to predict fire damage for use in nuclear war planning. Long before, Brode had modeled the nuclear blast environment and contributed to the understanding of a variety of nuclear weapons effects. He had an entirely different perspective from the fire research community. Brode did not extrapolate from large conventional urban or forest fires but developed a regional model of the nuclear fire environment created by the heat rising from many simultaneous ignitions. Consistent with the physics of the fire environment, he then developed algorithms to predict fire damage that incorporated both the form and content of existing blast damage algorithms. In short, Brode worked within a "fire-blast damage frame," and enlarged the blast VNTK system to the fire-blast VNTK system.[5] This was a significant innovation that in many cases predicted far greater damage: On average, damage from fire reached two to five times farther than damage from blast alone. At the same time, the procedure was made palatable and "useable" by preserving the VNTK system and by conforming to expectations that predictions took the form of damage to specific structures. The effort was well-funded by the

Defense Nuclear Agency, the government organization then responsible for research on nuclear weapons effects. Over about a dozen years, the research yielded increasingly powerful predictions and looked likely to become incorporated into war planning. However, the effort was halted in 1992, though evidently revived in different form later.

Clearly, the organizations involved in nuclear war planning had the potential for generating significant innovation. Given the deep entrenchment of routines to predict blast damage, the development of a method to predict fire damage was certainly not inevitable; indeed, it seems unlikely. Without the determination of a single researcher who was well connected and bureaucratically savvy, nuclear fire damage would have continued to look utterly unpredictable.

How does this innovation square with the institutional persistence expected in path-dependent processes? Although much of the history above is self-reinforcing, in two respects it does not conform to important scholarly understandings about path dependence. Perhaps not coincidentally, these are the earliest and most recent aspects of the history, where issues of change are particularly salient.

Scholars have argued that the origins of path-dependent processes lie in random events. For economist Brian Arthur, history itself is equated with "random historical sequence," "historical chance," and "the small elements outside our economic model that we must treat as random."[6] For sociologist Jack Goldstone, historical "outcomes are related stochastically to initial conditions."'[7]

History as "random"—this is a remarkably ahistorical approach to history. As we saw, the origins of U.S. precision bombing doctrine and its emphasis on blast weapons were structured by the past; for example, a longstanding political aversion to "promiscuous bombing" and a service tradition of precision marksmanship. If the development of U.S. precision bombing doctrine does not seem inevitable—as what history does?—neither does it seem random. No wonder that historians have opted for weaker notions of path dependence, claiming only that "what has happened at an earlier point in time will affect the possible outcomes of a sequence of events occurring at a later point in time," an idea that denies stochastic origins.[8]

Regarding change, scholars are divided on how "locked-in" are path-dependent processes. In the strong version, path-dependent processes "set into motion institutional patterns or event chains that have deterministic properties." Like inertia, "once processes are set into motion . . . [they] tend to stay in motion" and to "reproduce a particular institutional pattern over time."[9]

Political scientist Kathleen Thelen critiques this strong version of path dependence as "both too contingent and too deterministic . . . too contingent in . . . the initial choice" in not emphasizing, among other things, antecedent conditions and "too deterministic in that once the initial choice is made . . . the argument becomes mechanical."[10]

Political scientist Paul Pierson argues for a less deterministic version: "Path dependent analyses need not imply that a particular alternative is permanently locked in following the move onto a self-reinforcing path. . . . Asserting that the social landscape can be permanently frozen hardly is credible, and that is not the claim. Change continues, but it is bounded change—until something erodes or swamps the mechanisms of reproduction that generate continuity"[11]

Self-reinforcing processes and bounded change characterize the innovations made over decades in the VNTK system to predict blast damage. However, erosion or swamping—presumably some form of external pressure or shock—is not the mechanism at work in the recent past. The external environment mattered—whether in political mobilization around nuclear winter issues or the end of the Cold War. Those involved in predicting nuclear weapons damage used these events as political resources to bolster their own claims. But the key to potential change was not the external environment but an internal incubator of innovation: the government agency responsible for research on nuclear weapons effects, the Defense Nuclear Agency. At the same time, far-reaching innovation was anything but guaranteed.

In sum, path dependence powerfully describes much of the self-reinforcing history of nuclear weapons damage prediction. But it does less well as a description of origins or of more recent innovation. Wherever one cuts into the history, whether the earliest origins of bomb damage prediction before and during World War II or the post-World War II development of nuclear damage prediction, random processes are not a good characterization. And, more recently, the history looks neither locked-in along a predetermined path nor jolted off that path by external change.

The Science of Destruction Is Social

I want to return to a theme that has run throughout the book: how deeply social is the enterprise to understand nuclear weapons effects and predict weapons damage. Of course, nothing could be more powerful in its facticity than the effects of nuclear weapons. Yet understanding such effects is a social enterprise in at least four respects.

First, as we have seen, nature is read from inside institutions. These include institutions of science and engineering, which bring to bear distinctive intellectual orientations and tools, and organizations with goals and frames that shape how problems are represented and solutions conceived.

Two examples from the text suffice to show how organizations embed the social in the understanding of nuclear weapons effects. "Target hardness" is not simply a characterization of the strength of structures. The phrase is meaningless without a specification of social purpose, which is indicated by the type and degree of damage sought. Take an example of a heavy industrial structure containing steel or chemical production. If the goal is to collapse some supporting walls to preclude use of the building until major repairs are made ("moderate" blast damage), then the structure is rated as less hard, or more vulnerable, than if the goal is to turn the structure into matchsticks ("severe" blast damage). Many structures and installations have two such ratings (or more if "light" damage is also included). If fire damage is considered, the structure will be rated as more vulnerable than if fire is ignored. Thus, how hard a target is considered to be depends not only on its physical construction but on the level of damage sought and the kind of damage that "counts."

Similarly, the predictability of mass fire and resulting fire damage depends not only on whether variation in weather conditions significantly affects the probability and range of mass fire, but on the criteria, or social requirements, for prediction. Weather itself illustrates the social aspects of prediction. For purposes of packing my suitcase, I know that in the summer, I should pack for much warmer weather in Washington, D.C., than in San Francisco, and I know that it is extremely unlikely that I will need an umbrella in Los Angeles. For these purposes, the weather can be predicted with high probability. On the other hand, if I want to know if I can wear suede shoes in January in San Francisco, I probably will think that the weather cannot be predicted very well. Thus, it is only meaningful to say that something can or cannot be predicted against some understanding of purpose. Regarding the predictability of fire damage, the requirements for solution were critical. If nuclear war planners had accepted a solution in which area damage was predicted (e.g., the area corresponding to the perimeter of mass fire at Hiroshima), a robust prediction of mass fire, and resulting damage, could have been developed in the first decades after World War II. But because war planners thought it necessary to characterize the vulnerability of specific structures to fire damage, in a way consistent with how blast damage was predicted, a method to do so was not developed until the 1980s.

Second, the institutions dedicated to planning and predicting nuclear destruction are sites of great sociability. We have seen the rivalry and exasperation of wartime service; the camaraderie of office life; the hard work and sly strategies involved in mobilizing bureaucratic support; and the drama of high-level briefings in the protected room in the Pentagon known as the "tank." This sociability is seen in the jokes and slang used at the time (e.g., "If we meet our damage goals . . . then we don't give a hoot if we incinerate everything in the area");[12] in unpublished manuscripts and private publications such as Jerry Strope's "Autobiography of a Nerd" and Frank Shelton's *Reflections of a Nuclear Weaponeer;* and in published articles and books. (We should not forget that this social world was a highly compartmentalized and secret one—but secrecy involves its own social interactions. Henry Nash, for example, who worked in target intelligence in the early 1950s to identify government control targets, did not speak to the analysts who identified atomic power targets. Richard Grassy, for many years the head civilian in the group that analyzed the physical vulnerability of structures, did not discuss his work with his family.)[13]

Third, through omission, abstraction, classification, disembodiment, a focus on physical forces, specialized vocabularies, and whole systems of knowledge, nuclear war planners engage in a social construction of the asocial.[14] The very social world in which nuclear war planners live, both at work and outside, is entirely omitted from the environment they make plans to destroy. As if to anticipate the effect for which they plan—the utter effacement of human society—the environment they consider is abstracted from and devoid of the buzz and hum of human activity. The world of nuclear weapons damage is generally an unpeopled one of physical objects—structures, installations, and equipment.

In this world, buildings are structures that house war-making activities, not people per se; structures are targets; targets are categorized in census-like classifications, identified by numerical designators, located by latitude and longitude, and keyed to specially constructed maps. Specific buildings, installations, and equipment are classified as structural types; structural types are rated by how they respond to the physical forces, such as duration and drag, that act on them; physical forces are studied and described, as are the very large physical environments created by the blast, fire, and other effects of nuclear weapons. The effects of nuclear weapons on plants, animals, and people are studied in terms of material and physiological response. Nuclear weapons are allocated for maximum efficiency in a language of cost, requirements, and transportation logistics (weapons, for example, are "deconflicted" so that a warhead detonating does not destroy another aimed at the same or a nearby location). Each aspect is part of an edifice of disciplinary understanding, empirical study, and particular bodies of knowledge created to solve specialized problems. Much of this knowledge has been put in the form of computer codes and embedded in organizational routines. Such abstraction is inevitable in war planning and, indeed, in any planning, but it is not the less striking for it.

When the social is brought into nuclear war planning via considerations of civil defense, the effect is comic, an unintentional parody of the planning process that emphasizes the inability of planners to incorporate the social aspects of destruction. In civil defense planning, human society is neither omitted nor obliterated. Rather, it endures in a fantasyland of normalcy where a few buildings and mannequins represent a "typical American community"; the family car provides some protection "against the radiation, heat, and blast of a nuclear bomb"; "fireproof housekeeping" is efficacious; prompt rescue and recovery operations are possible; and families in protective shelters consume peanut butter and play charades. Even acknowledgment of the limits of civil defense planning parody it, as in the statements made in congressional testimony in the early 1960s that "we are not trying to maintain the present standard of living under thermonuclear attack," or that while "nuclear attack on the United States could be very serious, it need not be catastrophic."

Fourth, the meaning of "conventionalization," a term coined by Hans Morgenthau to characterize thinking about nuclear weapons as though they were conventional weapons, should be rethought. Robert Jervis argues that conventionalization is psychologically attractive because it denies what is disturbing about nuclear weapons and intellectually attractive because it allows analysts to use familiar strategic concepts.[15] The concept is lodged in the individual psyche. However, conventionalization should not be understood primarily as a fallacy of thought. It is much more powerfully understood as a social phenomenon residing in organizational capabilities and knowledge-laden routines in which nuclear weapons are treated as though they were conventional weapons—as though they cause damage to specific structures and do not lay waste to vast areas, as though they cause damage by blast and not by mass fire or other less "conventional" means, as though their destructive power should be measured as so many tons of dynamite, and so on. Conventionalization lies in the problems organizations seek to solve and in their routines, not in individuals' minds. The implications for change are significant. The point is not to change thinking or attitudes or psychological acuity but to change the problem-focus of organizations in building knowledge and routines.

Organization-Made Disasters

These ideas—organizational frames, path dependence, the deep embeddedness of the social in our readings of nature, and the social construction of the asocial—illuminate the particularities of the history represented in *Whole World on Fire.* I want now to step outside this history to explore briefly some broader implications.

In the introduction I said that the partial prediction of nuclear weapons damage is a case of poorly understood or unanticipated physical processes whose resulting representation of the physical world in documents, technologies, and routines is inaccurate or incomplete. I mentioned some other examples: the shipbuilder's lack of understanding of how brittle the steel was in the *Titanic;* Grumman's lack of understanding of the severity of potholes on New York City streets and their effect on newly designed buses; and the engineers[5] and architects' lack of anticipation and understanding of the effects of burning jet fuel inside the towers of the World Trade Center.

These are also examples of problem solving consistent with the best contemporary standards of professional practice: There was no suppression of evidence or disregard of a well-understood body of knowledge. (Grumman may not have been engaged in best practice, but the company did not suppress evidence or lie, and it quickly took responsibility for its design errors.)[16] To say this is to adumbrate four categories of organizational approaches to problem solving (see Table C.i).

In the first category, physical processes are well understood by organizational problem solvers, and organizational actions are consistent with best practice. Consequently, resulting technologies and knowledge-laden routines are reliable and safe. This is the world each of us hopes to live in all the time, a world in which ships float, buildings are structurally sound, elevators are safe, airplanes fly as expected, and mushrooms and meat in the supermarket are untainted.

Second, there is the "dark" side of organizations, the world of corporate wrongdoing and crime, in which physical processes are well understood, but organizational approaches to problem solving fly in the face of contemporary best practice: Organizational actors neglect, suppress, or lie about evidence and the state of knowledge. In these cases, the pursuit of craven organizational interests in profit or, rarely, organizational or national pride, cause people in positions of responsibility to ignore physical processes that they could, and should, understand. Examples include the cigarette industry after about i960, when the connection between lung cancer and cigarette smoking had been clearly established and the industry both denied the state of knowledge and suppressed and twisted evidence; the suppression of evidence of danger in the Corvair automobile that rolled over; the decision not to change the design of the Ford Pinto gas tank that exploded; and the continued manufacture and sale of Firestone tires that shredded. Other examples include French officials who did not prevent contamination of their national blood supply by the HIV virus; builders in India who did not construct buildings to code, which led to disastrous collapses in a large earthquake in early 2001; and British manufacturers dumping in Europe thousands of tons of feed suspected of causing mad cow disease after the feed had been banned in the United Kingdom. According to one editorial, the dumping was "morally unforgivable even if legal."[17]

Third is the "dumb and dark" side of organizations in which problem solvers poorly understand relevant physical processes, and their approaches

TABLE C.1.
ORGANIZATIONAL APPROACHES TO PROBLEM SOLVING

Category	Physical processes well understood?	Organizational actions consistent with best practice?	Potential results	Examples
1. Ideal	Yes	Yes	Technologies and knowledge-laden routines are reliable and safe	Ships float; buildings stand up; food is safe
2. "Dark" side	Yes	No	Dire consequences as organizational actors neglect, suppress, or lie about the state of knowledge, usually for profit	Cigarette industry after 1960; Ford Pinto gas tanks; Firestone tires
3. "Dumb and dark" side	No	No	Dire consequences, but was there negligence in organizational actors' interpretation of evidence?	Mad cow disease; Cerro Grande fire; *Challenger* explosion
4. "Ignorant but upright"	No	Yes	Dire or potentially dire consequences due to poorly understood physical processes reflected in knowledge-laden routines	*Titanic*; effects of burning jet fuel in World Trade Center; partial prediction of nuclear weapons damage; but disaster averted at Citicorp Center

to problem solving are apparently inconsistent with contemporary best practice. These cases often raise difficult questions about what could and should have been known, whether evidence was interpreted poorly, and if action was taken without proper precaution. Was there negligence, in other words, and if so, to what degree, and why? Examples include the very profitable use of sheep remains in cattle feed in Britain in the 1970s in the face of some early warnings of risk, which resulted in what we now know as mad cow disease; and the conduct of clinical trials for an asthma study at Johns Hopkins medical school that resulted in the death of a participant. The study "failed to obtain published literature about the known association between hexamethonium [the drug used in the study] and lung toxicity," which was "readily available," and violated federal regulations that required, among other things, the convening of face-to-face meetings of medical review boards overseeing such studies.[18]

Another example is the accidental burning by the U.S. Park Service of 48,000 acres (75 square miles) in the Cerro Grande fire near Los Alamos National Laboratory in May 2000. Eighteen thousand residents were evacuated, hundreds of homes were destroyed or damaged, and total damage was estimated at about $1 billion. In a contingency never planned for, the nuclear weapons laboratory itself was threatened, and forty laboratory structures were destroyed.[19] This gigantic wildland fire resulted from a deliberately set, or prescribed, fire going out of control. First, the prescribed fire burned beyond its boundaries and then, to contain it, a backfire was introduced that, in conjunction with the wind and seasonal conditions, was disastrous.[20] A National Park Service board of inquiry found that "questionable judgment was exercised" but that there were "no violations of policy."[21] It is clear that Park Service fire managers did not understand the risks, and this was due partly to procedural problems in risk assessment. The National Weather Service did not predict winds in their three- to five-day forecast due to constantly changing conditions, and Park Service personnel evidently took that to mean high winds were not expected. Indeed, the Park Service official in charge "said that if he had better information on the wind ... he would not have introduced fire . . . into the burn area."[22] Yet, a General Accounting Office (GAO) study said, "This time of year typically brings high winds, [further,] the area was in the midst of a 3-year drought. . . . Also, during the 2-week period before the fire was started . . . four prescribed fires got out of control in that region.[23] Another procedural problem was that the fire complexity ratings for prescribed fires had been mistranscribed on the web site used by National Park Service fire managers, resulting in a significant underestimate of the difficulties that could be encountered. Given evident incompetence and inadequate procedures, the GAO recommended that prescribed burn plans "need to be 'peer-reviewed[5] by independent, knowledgeable individuals."[24]

Another example is the U.S. space shuttle program's understanding of the behavior of the O-rings that sealed in the hot propellant gases in the shuttle's booster rockets—the failure of which resulted in the explosion of the *Challenger* on January 28, 1986. The engineers did not understand the mechanisms of sealing in cold weather (although they thought they did), nor did they clearly see the correlation of cold temperature and erosion of the O-rings by hot gases. But were they negligent in not understanding these complicated processes? On the one hand, they believed they were following best practice, and in many respects they were. On the other hand, in part due to design compromises (and all projects have design compromises), the difficulty of understanding the sealing mechanisms, and the scale and complexity of the whole enterprise that caused them not to know what they did not know, they unwittingly departed from best practice in a process Diane Vaughan terms the "normalization of deviance." It appears that the catastrophic failure of the *Columbia* space shuttle on February 1, 2003, reflects a similar "incremental descent into poor judgment."[25]

In these examples, organizational history and goals contributed to incompetent problem solving for physical processes that were not well understood. Explicit or implicit pressure from the top to proceed seems likely. In addition, in the short term, precaution can be mind-bogglingly expensive (though not as costly as the failure that may result).[26]

Best Practice, Mostly Bad Outcomes

Finally, let us turn to the fourth category in which poorly understood physical processes, embodied in knowledge-laden technologies and routines, combine with best contemporary practice to produce dire or potentially dire consequences. In addition to the partial prediction of nuclear weapons damage and cases mentioned above—the *Titanic*^ steel, the potholes on New York City streets, and the effects of burning jet fuel in the World Trade Center—there are other striking examples. These include ignorance about the spread of childbed, or puerperal, fever in the eighteenth and nineteenth centuries—"the most serious, deadly, and terrifying of all the complications of childbirth and the most common cause of maternal deaths" in this period;[27] not understood dynamic loads on suspension bridges in the early twentieth century that resulted in the sudden collapse of the Tacoma Narrows Bridge on November 7,1940; and the failure to calculate certain forces on New York's Citicorp Center and Boston's John Hancock Tower in the 1970s, which could have resulted in catastrophic collapses.

These examples—and there are many others—are especially troubling because they involve no willful misinterpretation of evidence or obvious deviation from best practice. This suggests that incentives to engage in best practice or punitive measures to inhibit such actions will not be effective since competent, even preeminent, practitioners are already doing the best they can and acting with integrity

This does not mean that nothing can be done. These examples represent a wide range of outcomes. In these cases, the understanding that paves the way to solution can occur after persistent failure (childbed fever) or a single failure *(Titanic,* Tacoma Narrows Bridge, World Trade Center). Dire consequences can also be averted (Citicorp Center, John Hancock Tower) or may remain unrecognized (underestimates of nuclear weapons damage). Thus, these cases have implications both for understanding dire consequences and for preventing disaster.

Let us begin with childbed fever. In the eighteenth and nineteenth centuries, lying-in hospitals for women in childbirth became widespread in Europe. These hospitals provided rest and nutrition for women, professional delivery by midwives and doctors, and training facilities for midwives and medical students. One problem associated with these hospitals was the very high rate of childbed fever. Within a few days after delivery, affected mothers began to suffer from terrible shivers and fevers, excruciating abdominal pain, and, often, death. Sometimes entire maternity wards would suffer epidemics in which "nearly every patient died."[28]

Drawing on current medical knowledge, doctors tried to understand the problem. French doctors who did postmortems on the affected women observed a milky white substance covering the intestines and omentum and theorized that breast milk had metastasized to the abdominal cavity. An English doctor thought it was due to the "putridity" of the indoor atmosphere in which, deprived of an essential ingredient, the air became "vitiated." Others thought childbed fever was due to miasmas—odorless materials in the air emanating from vegetable decomposition—or to "mental depression, malnutrition, or its opposite," gluttony.[29]

Outbreaks of childbed fever were sometimes associated with particular midwives and doctors, but until the early 1850s, no one thought that midwives and doctors themselves might play a role. However, as part of their training, medical students performed postmortem examinations and then routinely went from examining cadavers to delivering women—without washing their hands. And they went from mother to mother delivering babies—again without washing their hands.[30] Their knowledge-laden childbirth routines did not include hand washing for the same reason that we do not routinely stand on our heads before taking tests: They could see no causal connection. It was not until the last quarter of the nineteenth century, after the development of the germ theory and its incorporation into hospital practices of antisepsis and sterilization—thanks largely to the physician Joseph Lister—that incidents of childbed fever, caused primarily by streptococcal bacteria, declined dramatically.[31] Until the germ theory was developed, the theoretical knowledge base from which to derive a solution was beyond not only practitioners in hospitals but all contemporaries.

The ship *Titanic,* which sank in 1912 when it hit an iceberg in the North Atlantic, is similar in one regard: Knowledge of the time was inadequate to prevent catastrophe or to directly address it after. The proximate cause of disaster was the flooding of the forward five compartments in the ship, which as every moviegoer now knows, had been designed to withstand flooding in the first four. The design itself, which set new marks for safety, cannot be faulted by contemporary standards. For many years the prevailing theory held that the iceberg had torn a large continuous gash in the side of the ship. But when the ship was found at the bottom of the ocean many years later, it turned out that the iceberg had not forcefully punctured the side; rather, the pressure of the iceberg had caused the ship's inch-thick steel plates to buckle and to open in several thin discontinuous slits. The plates had buckled because they were brittle in cold water.[32]

Contemporaries understood that brittle metal was a problem in shipbuilding. For that reason, the *Titanic's* steel plates were not made by the Bessemer process, which produced brittle steel, particularly at low temperatures (due to its high nitrogen content). Indeed, according to an authoritative study, the steel used in the *Titanic* was "probably the best. . . ship plate available in the period of 1909 to 1911." The only other manufacturing method available was the open-hearth process, which was most commonly done in acid-lined tubs. The acid-lined tubs produced steel with a high sulfur content and other chemicals that, as it turned out, also embrittled steel and produced a hull "not suited for service at low temperatures."[33]

Given the lengths to which the *Titanic's* builders went to design and build a safe ship, it seems highly unlikely that they were aware of the effects of their manufacturing methods. Whether steel makers and metallurgists did not understand how the acid-lined tubs interacted with the steel being produced, or could not analyze the steel content and/or the embrittling effects of certain chemicals, it seems likely that no one understood the vulnerability of the steel produced for the *Titanic.* If the required knowledge regarding content or effects of steel content was not beyond the theoretical knowledge base of the time (and I do not know whether it was or was not), in all likelihood it was beyond the knowledge base available to the organizations and practitioners involved in steel production.

Three other cases illustrate failure to understand physical processes within a context of high professional standards. In the Tacoma Narrows Bridge collapse in 1940 and in the serious design errors in the Citicorp Center and the John Hancock Tower in the 1970s, the required understandings of physical processes were well within the knowledge base of contemporaries and, hence, were much more amenable to solution. In these cases, the failures lay in the problems engineers sought to solve, and the problems they did not.

At the time it was built, the Tacoma Narrows Bridge near Seattle was the third longest suspension bridge in the world (after the Golden Gate Bridge in San Francisco and the George Washington Bridge in New York). From the time it opened in 1940, the bridge, known as "Galloping Gertie," undulated in the wind. Flexible suspension bridges were not unusual in this period and were frequently stiffened after construction. Engineers observed the bridge and began to take steps to reduce the sway, but no one expected a catastrophic failure. A few months after it opened, in a light wind in early November, the bridge not only swayed but began to twist, the sides of the roadway seesawing. Within a short time, the bridge tore itself apart and collapsed.[34] Fortunately, the bridge was closed to traffic that day, and no one was killed.

The bridge had been designed by an eminent engineer, Leon Moisseiff, who worked within well-established suspension bridge design principles of the period. Modern engineers had developed what appeared to be robust design algorithms that calculated wind forces on bridges as static, or steady, forces rather than as dynamic forces. Using this method, they worked within a "design climate" in which they focused on principles of structural simplicity and aesthetics to produce "ever longer, slenderer, and lighter suspension bridges."[35] The methods they used had been successful in the George Washington Bridge, built in the 1920s, and in later suspension bridges. However, the algorithms poorly represented the forces on the bridges, although bridge engineers were unaware of it at the time. Engineers modeled the wind forces pushing sideways on the roadway, but they did not take into account the forces that could lift the road and drag it down, much like an airplane wing. It was these forces that would cause the Tacoma Narrows Bridge to twist and collapse.[36]

Clearly, the effects of dynamic forces on bridges were beyond bridge designers' understanding at the time. But the new field of aerodynamics, used in the design of airplanes in the 1930s, provided precisely the dynamic analysis that was needed for suspension bridge building. Indeed, engineer W. Watters Pagon published a series of eight articles on aerodynamics in the 1930s—the first titled "What Aerodynamics Can Teach the Civil Engineer." However, according to author-engineer Henry Petroski, "the whole series seems largely to have been ignored by the bridge builders," in large part because "bridge building was becoming so highly specialized that there was the 'danger of losing contact with the other branches of engineering and with allied sciences/"[37] In this case, a fully developed knowledge base was available, but bridge builders did not make use of it.

Two recent cases illustrate other oversights. To accommodate a church on a corner of the building site of the fifty-nine-story Citicorp Center in New York, the structural engineer William J. LeMessurier decided to support the building's steel skeleton on four massive columns placed at the center of each side rather than at the corners as was usually done; he also used an innovative system of steel braces to provide strength against the wind. But when an engineering student challenged the strength of the completed structure, LeMessurier found, to his great surprise, that the steel braces were not as strong as he had expected against winds hitting the building from the corners, called quartering winds. The New York City building code required only that the perpendicular winds pushing face-on to the structure be calculated, but LeMessurier had also calculated quartering winds in the design; in particular, the massive columns placed in the center of each side were unusually strong against them. Although the engineer's recalculations showed that the strain on the braces was greater than anticipated, it was well within the margin of safety, all other things being equal.[38]

But all other things were not equal. In his reexamination, LeMessurier also discovered that the joints that held together the building's steel girders had not been built to his original specifications. Instead, his office had approved a change recommended by the construction company that the joints be bolted instead of welded, on the grounds that welds were stronger than necessary. This was not a question of improper procedure or shoddy construction. The problem was that in designing the bolts, LeMessurier's office had not considered the sensitivity to quartering winds. This, plus another "subtle conceptual error," meant that the Citicorp building could fail catastrophically in a "sixteen-year storm"—a storm with a probability of occurring once every sixteen years.

After contemplating silence or suicide, LeMessurier explained the problem to the building's lawyers, architects, insurers, and owners. Emergency repairs were made, and disaster averted.

Somewhat similarly, in 1975 a renowned structural engineer, Bruno Thurlimann, determined that under certain wind conditions the new John Hancock Tower in Boston could fall over, not on its face but on its "narrow edge . . . as if a book standing upright on a table were to fall on its spine."[39] (This is the same building notorious for window panes falling out.) Like the Citicorp Center, the John Hancock Tower was an innovative design that met all building codes. And like the Citicorp building, the structure was scrutinized by fellow engineers. As with Citicorp, analysis revealed the problem, and it was corrected immediately.

We might be tempted to say that the structural engineers in these two cases were not adhering to best practice, since best practice would dictate that buildings be not so vulnerable to wind forces. Yet, as with the Tacoma Narrows Bridge, the engineers involved were at the top of their profession, they *defined* best practice, and their oversights were not obvious at the time. However, the understanding of the physical processes involved was well within the knowledge base of civil engineering. In both buildings, when the errors were pointed out by other engineers, those responsible immediately understood the problem. The framing of problems is clearly what caused a lack of attention to particular wind loads on these structures.

Finally, the effects of burning jet fuel inside the towers of the World Trade Center may seem like a failure to anticipate the social environment rather than the physical one. It is hard not to agree with National Security Advisor Condoleezza Rice's statement that "I don't think anybody could have predicted that these people would take an airplane and slam it into the World Trade Center."[40] Yet, Leslie Robertson, the engineer in charge of the structural design of the towers, did consider the contingency of the largest jet aircraft of the time, a Boeing 707, hitting the building, and he designed the building to withstand its impact.[41] He did not, however, design for "thousands of gallons of fuel being put inside the building," according to a prominent structural engineer, Abolhassan Astaneh-Asl.[42] Why not? According to Robertson himself, after designing for the impact of an aircraft,

> The next step would have been to think about the fuel load, and I've been searching my brain, but I don't know what happened there, whether in all our testing we thought about it. Now we know what happens—it explodes. I don't know if we considered the fire damage that would cause. Anyway, the architect, not the engineer, is the one who specifies the fire system.[43]

Aircraft impact, force, and structural response were anticipated, but potential fire damage was overlooked.

State of Knowledge, Problem Recognition, Secrecy

We see in these examples a range of determinants of understanding of physical processes. At one extreme are the spread of childbed fever and the brittleness of the *Titanic's* steel plates, which were beyond the ability of every contemporary to understand. In such cases, the background state of knowledge about the physical world defines and delimits how problems are cast.

In the other examples, the requisite knowledge base about the physical world was, at least in theory, available to contemporaries. In these cases, the key lay in how problems were represented and solutions were defined—organizational frames. For example, had the designers of suspension bridges been familiar with aerodynamics, they could have much more quickly understood the forces on the Tacoma Narrows Bridge and would have solved structural design problems differently. Had the *Challenger's* engineers been deeply grounded in statistical analysis or graphical analysis, they would have seen the danger of launching the *Challenger* shuttle in record-cold weather.[44] It is not my goal to explain how these particular organizational frames developed. Indeed, satisfying explanations have already been written. I simply want to say that these are the questions to be asked and answered.

We have seen that feedback from the environment can indicate that a severe problem exists, but contemporaries may be unable to diagnosis the problem or solve it (childbed fever). On the other hand, contemporaries may have the ability to recognize a problem before there is any direct indication that a problem exists (Citicorp).[45]

We can see that environmental feedback is *always* mediated through social expectations, whether at the organizational or societal level. For example, at the organizational and professional level, the flexibility of suspension bridges, even the galloping of Gertie, did not lead bridge designers to think there was a serious problem in their calculations of forces. They expected the swaying and thought they understood its causes. Similarly, the *Challenger* space shuttle's engineers reinterpreted the increasing erosion of the O-rings as normal and not dangerous.[46] We might think that mobilization for safety occurs when the threat to life is obvious. But even potential or actual deaths are not good predictors of the social expectations of the acceptability of failures. Why do important problems in airline safety not get addressed until after crashes bring them to public attention? Why have large numbers of deaths from routine medical mistakes persisted for so long?[47]

Understanding that physical signs must always be socially interpreted also allows us not to make the mistake of using later understandings to read back into earlier situations our own superiority and participants' apparent stupidity. As a member of the commission investigating the *Challenger* disaster, physicist Richard Feynman's famous and rhetorically effective demonstration that O-rings stiffen in ice water was, in terms of our understanding of what space shuttle engineers understood, beside the point. The *Challenger's* engineers knew that O-rings stiffen in cold temperature.[48] They had *not* understood, among other things, that the backup that they had thought would compensate for O-ring behavior in cold temperature was inadequate. Feynman and the rest of the commission did not understand the complicated organizational context that had led to these interpretations of O-ring behavior. No one did until the painstaking research and original interpretation by Diane Vaughan.

Finally, to what extent were people other than organizational problem solvers aware of a problem and empowered to address it? As we have seen, transparency of technology and design to a wide professional community can raise awareness of problems and prevent disaster. Catastrophic failure can also make widely known what was not and can bring others into the process. This is not automatic. Under conditions of extreme secrecy and lack of democratic accountability, even disaster can be made invisible. For example, in 1979 a military biological weapons laboratory in Sverdlovsk in the Soviet Union accidentally released airborne anthrax spores that killed sixty-six people. Soviet officials lied and claimed that people had died of gastrointestinal and cutaneous anthrax due to consumption of contaminated meat and contact with diseased animals. Questions were raised over many years, but it was not until a decade and a half later that the matter was fully resolved as to the kind of anthrax and the source of the release.[49]

Nuclear Weapons Damage

These cases give us a comparative context in which to understand the half century in which the U.S. government did not predict nuclear fire damage for decades and then chose not to incorporate such predictions into knowledge-laden organizational routines. Clearly, the failure to predict fire damage lies at an extreme of persistence (not to say potential consequence), at least among known contemporary examples.

One possible explanation does not hold here: The contemporary knowledge base did not foreclose the possibility of prediction. Although most considered nuclear mass fire and resulting damage so complex as to defy prediction, an understanding of the basic physical processes involved was well within the knowledge base of physicists, and had been for many years. The applied knowledge required for damage prediction was not so ready-made as in the case of Citicorp or the John Hancock building or even the Tacoma Narrows Bridge: It was not circulating among practicing professionals. Instead, it had to be made, just as all knowledge about nuclear weapons effects had to be made. The physics of the fire environment had to be modeled. The potential variables contributing to fire damage had to be analyzed. The results had to be translated into organizational routines that were consistent with and built on past damage-predicting routines. As this book has demonstrated, the key lay in organizational frames, the approaches to problems by those in organizations, which influenced the mobilization of expertise, resources, and resulting knowledge-laden routines.

One might think that the mass fires at Hiroshima and Nagasaki would have been sufficient indication that fire damage mattered. However, the issue was not understood to be whether nuclear fire damage would sometimes occur, but whether it would occur with enough regularity that it could be robustly predicted. The answer was thought to be no. Nothing occurred after the war to shake confidence in the adequacy of this answer. But a counterfactual or two demonstrates the role that environmental feedback could have played. Had one of the atomic bombs dropped over the Nevada Test Site gone astray and accidentally burned down Las Vegas—approximately 65 miles away—or had the United States inadvertently burned down Moscow in a "limited" nuclear exchange, it seems likely that war planners would have reevaluated the necessity and feasibility of predicting nuclear fire damage. Fortunately for the world, these scenarios never occurred.

Unfortunately, as we have seen, it is often catastrophes that make known what was not known, or widely known, and that put the pressure of public accountability on internal organizational processes. Although professional standards within the nuclear weapons effects community have been high, the issue of nuclear fire damage has been nearly invisible to the public. There has been very little mention in the press or discussion by scholars of its importance or omission in war planning, and no discussion until now of how this has come about.[50] The lack of visibility resulted from both formal secrecy and opacity. The world of nuclear war planning is a secret one separated from practicing professionals and ordinary citizens.[51] To a large extent, this is a self-policing system in which those with classified knowledge pledge not to divulge it.

Still, it may not be the formal secrecy that has kept the issue from public awareness so much as opacity: Even unclassified information is not widely understood. Unlike building design, the technical issues are not familiar to a broad community of practicing professionals. With the notable exception of MIT physicist Theodore Postol, few outside the government-sponsored nuclear weapons effects community have paid attention to these issues or been available to explain them to journalists, scholars, and the wider public. The organizational processes that have determined which problems are solved and which are not are no less important, and these too have been hidden from the public.

No wonder, then, that the lack of prediction of nuclear fire damage has been so persistent. These are largely self-reinforcing organizational processes, sealed off from the public through secrecy and opacity. And since World War II, the consequences of these weapons have been in the realm of the hypothetical.

What we have seen is paradoxical: Organizations should think about what they are not thinking about—a kind of organizational walking and chewing gum at the same time. It is not that organizations should simply do worst-case analysis well. They should figure out what problems they are not trying to solve and examine how those could lead to consequences worse than the worst case being considered. The engineer LeMessurier's advice is instructive: "Any time you depart from established practice, make ten times the effort, ten times the investigations. Especially on a very large-scale project."[52] Of course, organizations do not think; people think and approach problems in ways that are structured by organizational history, capacity, and routines. As we have seen, it can be extremely difficult to change organizational approaches to problem solving. Dominant understandings, not surprisingly, dominate. Organizations that do not encourage alternatives to, or questioning of, dominant approaches to problem solving may overlook important problems. (In FBI agent Coleen Rowley's words after the September 11 attacks, this is the "don't rock the boat, don't ask a question" problem.)[53] Queries from the top of an organization, or from outside, regarding technology and the physical world may be answered in ways that simply reflect ongoing approaches to problem solving. Further, change cannot simply be mandated from the top or from outside. To be fully effective, change must be implemented at the level of knowledge-laden routines, algorithms that both represent problems and embody solutions.

The only alternative to learning from catastrophe is learning from smaller failures, near failures, and scenarios of possible failure and unforeseen consequences. Precaution regarding the unforeseen is particularly important in a world in which the full consequences of our actions will not be fully known until much later. We have experienced many unforeseen consequences of twentieth-century innovation, from the miracle mineral asbestos that has proved dangerous to human health to the miracle drugs that are steadily losing their effectiveness in promoting human health. What will the twenty-first century hold? The consequences will be great indeed if, among other things, we do not anticipate the social and ecological consequences of huge construction projects (like the Three Gorges and Narmada dams), if we do not exercise precaution in proceeding with genetically engineered organisms, and if we do not understand the effects of our actions on global warming.

In all of our interactions with the physical world, organizational integrity and intelligence is critical. Visibility of organizational actions to independent professionals and scholars is necessary. Comprehensibility of organizational processes to a wider public is essential. Democratic accountability is indispensable.