CHAPTER 1

THE STAKES

News Release (sometime in the future)
Reliability Crisis Leads to National Conference in the Capitol

The National Infrastructure Reliability Conference opened this morning in Washington, D.C. The conference was convened by the president in response to the growing number of high-profile failures among the nation's electricity grids, air traffic control operations, telecommunications systems, financial exchanges, and interstate water supplies. "We have to get to the bottom of why these critical systems are not meeting the reliability standards our citizens have a right to demand and expect," the president said in his opening remarks. "We have to reduce an unprecedented and unacceptable vulnerability."

The president's address also focused on the recent catastrophic failure of the western grid, which knocked out electricity for over thirty million households, disrupted ports and related transportation along the West Coast, and

for more than a week interrupted major financial services and cell phone activity west of the Mississippi. The economic losses are estimated to be in the tens of billions.

The president addressed over eight hundred conference delegates. They included the CEOs and top executives of the largest corporations, which own over 85 percent of our nation's critical infrastructures. Also attending were leading engineers, economists, lawyers, and consultants who design these systems. Following the president's remarks, conference delegates listened to the keynote speech by the director of the National Academy of Engineering.

"We've invested billions of dollars, over the past two decades, designing and building the world's most advanced technical systems to provide essential services," she told delegates. "We've developed the most sophisticated risk assessment methods and designed the latest in safety systems. We've restructured our organizations to run these systems—streamlining them to operate with maximum efficiency to adapt to the changing requirements of new technology.

"Yet for all the investment," the director continued, "we have not realized proportionate improvements in the security and reliability of electricity, transportation, telecommunications, and financial services. Our national academy panel of experts now finds the nation vulnerable to more failures in our interdependent systems. The public is less confident today of infrastructure safety and dependability than they have ever been, and with good reason.

"We have lost something very important," the director concluded. "There were organizations in the last century that had far better records in running large, unforgiving technical systems. We need to recapture what they knew thirty years ago about running our complex infrastructures. The consequences are enormous if we don't."

A spirited debate followed.

"It's doubtful we've lost any real skills or information over the past quarter century," countered one of the country's leading electrical engineers. "Our expert systems and engineering models cover far more of the operation of complex technical systems than the design principles of the past or the 'gut' feelings of operators ever did." One Nobel laureate economist added, "Remember, many of the organizations of the past that ran these technical systems were rigid bureaucracies. They lacked the flexibility and incentives to adapt to changing infrastructure requirements."

The debate was especially heated in the afternoon conference panel, "Where to Now?" A prominent historian of science and technology gave a presentation on "high-reliability organizations." "Decades ago, a small group of researchers claimed to identify a set of organizations that made a special commitment to defy the odds and manage highly hazardous technical systems in air transport, nuclear power, and electric power with reliability as their highest objective. And these organizations did so with extremely high effectiveness," he argued. "Much of the research on these organizations was done in old-fashioned case studies quite difficult to locate now through our parallel cyberscans," added the historian. "But these studies, which focused more on organizational and management issues than on technology, could have considerable value in relation to present reliability problems."

"What is it they could possibly tell us?" questioned a well-known engineer from the floor. "What design principles can we distill from another era and a world so removed from our present technology and optimization methods?" The historian answered by describing the principal features of high reliability organizations.

"These organizations treated high reliability as a management challenge as much as a challenge in formal design. They didn't entirely trust the diagrams for their technical systems, nor the formal models and specifications describing how they should work. A great deal of attention was paid by members of these organizations to preparing for the unexpected—system conditions and events outside those specified in formal designs. Given this skepticism, a much larger role was played by operators and supporting personnel, who supplemented formal training with informal, experience-based perspectives on how these systems actually worked."

"This operator-and-discretion-based orientation would scarcely be possible today," interjected a panel member, "given our widespread self-referential expert control systems and flex-job practices throughout the economy. Worse yet, these older organizations clearly valued reliability of operations over optimization of performance. Their optimizing methodologies were primitive, and on many fronts these organizations didn't even try to optimize resource use or output, deeming it a potential threat to reliability."

An economist on the panel added that high reliability operations existed in conditions that artificially insulated them from competitive market pressures.

"Even those high reliability organizations exposed to market pressures were protected by regulatory frameworks that forced all of their competitors to make similar reliability investments and bear similar costs."

"We must recognize today," continued the panel economist, "that no modern, global supply-chain network can afford such excess capacity and inefficiency. Nor can we afford to return to those stultifying regulatory environments that depress technological innovation and destroy organizational flexibility."

"Quaint, and even noble, though these organizations may have been," another panel member summed up, "we can hardly expect engineers and economists to turn their backs on the contemporary world of formal design-based systems and reprogrammable organizations."

The conference continues tomorrow with the unveiling of the System Contingency Analyzer and Response Optimization Tool (SCAROT). This parallel-processing control system produces performance forecasts based on simultaneous analysis of hundreds of thousands of interdependent components across critical infrastructures. Every five seconds it issues intervention commands to better optimize intersystem performance. High-Reliability Products, a global LLC and one of the world's fastest growing multinationals, markets it.

THE PRECEDING NEWS ACCOUNT may seem fanciful, but we argue that something like it could very well occur, if we do not take more seriously the growing challenge of managing our critical systems. This book is about why such a conference has not been necessary—yet.

The interdependence of society's large technical infrastructures for electricity, telecommunications, financial services, and the like is well established (Dillon and Wright 2005; Conference on Critical Infrastructure Protection 2007). So too are demands for heightened performance of these complex networked systems, accompanied by technological and economic approaches to optimizing their operations. Increasingly common are technological strategies to protect against external threats and internal failures that might cascade through these systems (Pool 1997; Evan and Manion 2002).

Today is the high season of engineering and the historical hour of the economist in the world of critical infrastructures. This book lays out the case for an alternative focus—what we term "high reliability management."[1] We believe that many current approaches to design pose a danger to our critical

services as significant as the prospect of natural disaster or human attack. In fact, some design and technological changes promising added security realize the opposite by hobbling management and resources that could better protect us. We demonstrate the importance of these managerial skills to our current and future safety and security. The skills of the people who manage our critical infrastructures are misunderstood and neglected, even by some in the organizations charged with the creation, operation, or regulation of critical infrastructures.

Our argument is founded on research conducted over many years to understand the challenges and competencies of high reliability organizations (HROs). To this we add ongoing research on electrical grid reliability. This research has been undertaken over a six-year period, including extensive interviews with managers, dispatchers, engineers, and regulators within private utilities and power generators, from organizations such as the California Independent System Operator (CAISO), the California Public Utilities Commission (CPUC), and the California Energy Commission (CEC). Our analysis and argument are also informed by countless hours of direct control room observation.[2]

RELIABILITY HAS BECOME a worldwide watchword of citizens, clients, leaders, and executives. For some, it means the constancy of service; for others, the safety of core activities and processes (LaPorte 1996). Increasingly, it means both *anticipation* and *resilience*, the ability of organizations to plan for shocks as well as to absorb and rebound from them in order to provide services safely and continuously. But putting these together in a single strategy is a formidable challenge.

The study of reliability in critical systems is the study of strategic balances that must be struck between efficiency and reliability, between learning by trial and error and the prevention of risky mistakes, and between maximizing our anticipation of shocks and maximizing our resilience to recover after them. *High Reliability Management* is about how these balances are achieved and sustained.

From our research into nuclear power plants, electricity grid control rooms, water system control rooms, and air traffic control centers, we sketch the specifics of the skill set and managerial approaches that promote reliability in

settings where high reliability has long been the *sine qua non* not only of operational success but of organizational survival. We know that managers and executives, as well as system designers in varied settings, can learn useful lessons from the case material. Further, we show that in the absence of this understanding, many critical infrastructures are vulnerable to the very threats design and technology are meant to prevent. In demonstrating this, we tell the story of "reliability professionals"—a special group of professionals whose commitment and dedication make the difference on a daily basis between catastrophic failure of services we all depend on for life and livelihood and the routine functioning we have come to expect. Physicians call their life-threatening errors "never events," and reliability professionals are just as keen to avoid similar failures.

We direct our analysis to society's critical infrastructures because of their overwhelming importance. "Critical infrastructures" are core technical capabilities, along with the organizations that provide them, that enable the provision of a wide variety of social activities, goods, and services. Infrastructures in the domains of electricity, water resources, communications, transportation, and financial services are by their very nature multipurpose. They are critical in that they are necessary elements for more specific secondary systems and activities. If they fail, a wide variety of social and economic capacities are affected, with considerable economic consequences. The financial damage due to the August 2003 blackout in the northeastern United States alone is estimated to have been more than US$6 billion (de Bruijne 2006, 1).

Critical infrastructures have unusual properties that make them challenging to manage reliably. They are generally networked systems (de Bruijne 2006) with significant spatial dispersion, consisting of multiple organizations and varied interdependencies. Many of them, such as electric and transportation grids, have grown and developed by accretion, meaning they frequently consist of elements of varied age, design, and performance characteristics. The elements are difficult to characterize in a single model or consistent set of engineering or economic principles. To understand these systems and their reliability we must understand the special challenges they pose to management. In describing the challenges of high reliability management, we intend simultaneously to identify the practices, describe the professionalism, sketch out a

new research field, and draw implications of our argument for more effective policy and management.

There are no shortcuts to high reliability, though the temptation to try to find them is strong. Our own concluding recommendations, if simply distilled into design principles or management recipes, will assuredly make these systems more brittle than current design strategies have already rendered them. Failing to understand and appreciate the practices, professionalism, and research findings associated with high reliability management means a continuation of what we believe are great risks associated with many policies, technical designs, and business fads current today.

OUR ANALYSIS is a cautionary tale cast around three propositions:

1. There is an important divergence between dominant approaches to designing systems and the process of managing them. The difference is primarily in cognitive orientation—ways of framing, knowing, and responding—among different groups of professionals.
   - In particular, a distinct set of reliability professionals can be found among operators and middle managers of complex technical systems—individuals who are more than usually competent and motivated to have things "turn out right" in the operation of complex systems. They balance learning by trial and error with the prevention of high-risk mistakes; they anticipate to avoid shocks but maintain an ability for quick response to them when they happen.
   - This balancing typically requires skills to work effectively in a special cognitive frame of reference between the general principles and deductive orientations of designers and the case-by-case, experience-based preoccupation of field operators. These concerns drive reliability professionals to manage in terms of the patterns they recognize at the systemwide level and the action scenarios they formulate for the local level.
2. High reliability management is focused less on safeguarding single-factor or invariant performance than on maintaining a set of key organizational processes through adjustments within upper and lower limits acceptable for management (what we call "bandwidths"). The boundaries of these

bandwidths can be reliably changed only in proportion to improvements in the special knowledge base of the system's reliability professionals.

3. Despite the vulnerabilities they generate, centralization and interdependency among the component parts of a complex technical system can actually be significant managerial resources for reliability. Notwithstanding recent preoccupations among designers to do away with these properties, they provide options with which reliability professionals can make key adjustments and preserve balances needed for resilience and anticipation in a complex technical system.

As we illustrate, each proposition represents a neglected perspective in current policy and technology approaches to the operations of large technical systems. In fact, high reliability management in many respects is about the management of errors associated with both technology and policy.

In California's electricity restructuring, for example, the stated theories of economists, regulators, and legislators that electricity markets would quickly evolve and attract new players in electricity generation as well as keep wholesale prices down by means of market forces did not prove correct. In reality, managers of the grid confronted quite the opposite. They faced an unstable cycle of design errors leading to underperformance leading to ever-more frantic efforts at redesign.[3]

The architects of electricity restructuring were quick to say that their design was not really tried. For example, the retail market for electricity was not deregulated the way the wholesale market was. But smart and talented economists, policymakers, and regulators over a sustained period of time gave deregulated energy markets their best shot in California, and there were still major unforeseen and undesirable consequences they could neither predict nor control. Very clever reformers failed to forecast the first-order, let alone second-order, consequences of their policies, as we document in the book. Here, as in other cases, engineers, regulators, and economists were sorely undereducated when it came to the management requirements for highly reliable performance. We hope through this book to fill that education gap by making it clear just what the management requirements for high reliability performance entail.

We argue that no strategy of policy, technology, or markets can ensure reliability on its own without a strong management base. The importance of

that base is seen everywhere in the California example. Throughout the electricity crisis and the long aftermath induced by restructuring, the lights by and large have stayed on in the state. Why? Quite simply, the California Independent System Operator (CAISO) and the distribution utilities, including Pacific Gas & Electric (PG&E) and Southern California Edison (SCE), have taken reliability seriously, when others did not. How so?

First, according to economists, reliability is only one attribute of electricity and can be traded off against other attributes, such as how cheap or environmentally "clean" the energy is. But that clearly has not been the case. Reliable electricity (or water or telecommunications or so on) *defines* the critical infrastructure, and we as a society are unwilling to trade off reliability against the service's other features.

Second, advocates of major redesigns in our critical infrastructures have argued that reform, through either new technologies, markets, or policies, would significantly reduce the organizational burden of coordinating complex interdependencies. Not so. Deregulated energy markets, to name but one case, have created a far greater task of coordinating interdependency among market participants. This burden has fallen upon key focal organizations such as CAISO.

Finally, the real experiment in this infrastructure reform has not been the policy or the new technology itself, but something altogether more disturbing: a scrambling and reshuffling of institutions—single integrated utilities on the one hand and entirely new organizations on the other—all on the premise that organizations can be created or dismantled at will, without undermining service reliability in the process. The broad implications of this monumental conceit are exposed in this book.

It is said of Americans that they hate "regulation" but at the same time demand all manner of safeguards to ensure their own personal safety, health, and well-being. Our critical infrastructures provide a clear example of this demand. Americans, not just Californians, have shown themselves willing and able to spend billions upon billions of dollars in the name of ensuring the reliability of critical infrastructures—think of Y2K. Such transfers of income demonstrate that critical service reliability is not like any other "attribute" we know. High reliability is not just one more quality feature in the hedonic price for electricity. In fact, it is a foundation for the operation of society, not simply an attribute

of a service, in the same way a society's distribution of assets is a foundation from which any set of prices are derived. Substantially change asset distribution, and you change the prices; substantially change service reliability of our critical infrastructures, and you change their character as assets. Critical infrastructures have become so intertwined, and we, as a society, are so dependent on their always-on reliability, that high reliability has become a background condition, much like the framework of contract law, that *enables* market transactions.

As we show, high reliability in real time is not a bargainable commodity, nor is it sensibly traded-off by individual consumers in order to cheapen the costs of service. In real time when it matters the most, reliability is not exchanged or substituted for something else; otherwise those services would not be critical. This lesson was learned dramatically and at great cost in the California electricity crisis. The state budget surplus disappeared as the governor and his administration spent it on buying high-priced "spot-market" electricity, because not enough prescheduled power was being bid into the new energy markets to keep the lights on. As this was unfolding, a sign—"Reliability Through Markets"—was unceremoniously removed from the CAISO control room. It actually should have read "Markets Through Reliability." The evaporation of the California budget surplus (over $12 billion) and the recall of the state's governor testify to the foundational role of infrastructure reliability in modern society.

WHEN WE LEAVE CALIFORNIA, the challenges to critical infrastructures, from within and outside, remain the same. It is essential that all of us understand why.

Imagine a coordinated attack by terrorists striking at major electric power transmission lines and facilities in strategic places throughout the American Midwest and Northeast. They are able to knock out nearly 70,000 megawatts of peak-load electrical capacity and throw more than fifty million people into darkness over a 240,000 kilometer area in the United States and Canada. Without electric power a variety of other critical services fail, including water supplies and hospital facilities, as well as major financial markets over the globe. Ultimately, security systems become disabled, leaving key infrastructures vulnerable to additional terrorist attacks.

By this point, you may have already guessed that many of these conditions actually existed during the Northeast blackouts of August 14, 2003. The outages were caused not by terrorists but by the failure of electric transmission systems themselves, without hostile intervention. Although power was restored quickly in some areas, other portions of major metropolitan regions were without power for over twenty-four hours, and some areas had service interruptions for several days. It could have been worse. An earlier report issued in 2002 by a task force headed by former senators Gary Hart and Warren Rudman concluded that as a consequence of a coordinated terrorist attack, because of the lack of replacement parts for aged or customized equipment, "acute [power] shortages could mandate rolling blackouts for as long as several years" (Regalado and Fields 2003, A3). On November 4, 2006, the shutting down of a high-voltage line over a river in Germany to allow a ship to pass led to a chain-reaction set of outages that plunged ten million people in six European countries into darkness.

It is not only electric grids we have to worry about. We confront information networks under assault by computer viruses and hackers, large-scale transportation systems and water supplies open to terrorist attack—even the prospect of electronic voting exposes us to all manner of fraud and undetected error. Surprisingly, it is not expanding the reach of these complex systems but rather safeguarding their reliability that has become a great preoccupation of the twenty-first century.

At the same time, this preoccupation is often misguided in ways that are not fully appreciated. System designers and policymakers assume that the key to reliability lies in hardening our infrastructures so that they better resist outside attack. It is said that if we make these large technical systems more failsafe and foolproof or less tightly coupled or even more physically dispersed, we will improve their reliability (National Research Council 2002; Farrell, Lave, and Morgan 2002; Perrow 1999 [1984]). As one engineer has contended, "I try to design systems that are not only foolproof but damned foolproof, so even a damned fool can't screw them up."

It seems every week we hear of new data processing, electronic communications, and security systems that are planned or have already been put into place to increase reliability in the face of terrorist threats. More design changes

are planned for electricity grids and air traffic control systems (for example, Apt, Lave, Talukdar, Morgan, and Ilic 2004). Recent public discussions of business continuity have focused almost exclusively on design solutions to problems of businesses' protection or recovery from external threat (for example, *Financial Times* 2005). From our perspective, wrapping a patient suffering from internal bleeding in body armor of the latest style is not therapeutic.

We intend to show that the key to increased reliability for our electricity grids, water supplies, telcoms, transportation systems, and financial services, among others, lies not in the pursuit of optimal designs and fail-safe technologies for large complex systems but, rather, in their careful management. Unfortunately, there is a paradox at the core of this management. The very skills of high reliability management described in this book mask the vulnerability of that management to challenges and stresses, including those induced by misdirected designs and policy interventions. The professional attitude of key personnel and their virtuosity at working around design errors and rescuing situations in real time means that few signals are sent by these managers that conditions are worsening, until major failure actually occurs. Even higher management in their own organizations, as we will demonstrate, may not see how close to the edge the system is operating with respect to maintaining reliability.

Our intent is to signal what is happening in and to our critical infrastructures today. Rather than wait for major failures to communicate the risks, we offer a detailed look at the world of high reliability management, with careful case descriptions and close analysis of the skills at work. In the process, we offer a new method for measuring precursor conditions that serve as early-warning indicators of approaching edges to capacities for high reliability management.

THIS BOOK IS ORGANIZED into three sections. The principal focus of Part One is on the change over time in the reliability management of the California Independent System Operator (CAISO) as the transmission manager of California's high-voltage electrical grid, one of the world's most important electricity systems. To our knowledge, no other critical infrastructure control room has been examined as intensively with respect to managing for high reliability and over such a lengthy period of time (2001–2008).

We also discuss in Part One theories relevant to the challenge of high reliability, including "normal accident" theory and the theory of high reliability

organizations (HROs). We show the limitations of these theories in accounting for what we have observed, and we present our own framework for understanding high reliability management. This framework highlights the crucial role of reliability professionals—control operators, key technical department heads, and support personnel—whose special cognitive skills and flexible performance modes maintain reliable operations even in the face of widely varying and unpredictable conditions.

Part Two puts our analysis of key concepts, practices, and issues in high reliability management into a strategic perspective. Topics covered are (1) the critical balance between trial-and-error learning and failure-free performance in large technical systems; (2) strategies of managing performance fluctuations within controlled upper and lower limits and margins (bandwidths) as opposed to strategies for invariant performance; (3) the special domain of operational risks as opposed to analyzed risk, a domain in which risk *seeking* can enhance reliability; (4) the cognitive meaning of anticipation, resilience, robustness, and recovery and their operational trade-offs; and (5) the special threats to high reliability management posed by current approaches to technical design. Part Two ends with a chapter on indicators for key concepts, including the identification, measurement, and assessment of performance edges in high reliability settings.

Part Three moves the analysis to the wider context of critical infrastructures and the implications of our findings for the high reliability management of infrastructures in other social and organizational settings. National defense and homeland security are given special attention. We conclude with an examination of ways to support and protect reliability professionals, so as to ensure the provision of critical services in the future. It is a great irony that, while many economists are calling for greater efficiencies in critical service provision, and many engineers for greater capacity and effectiveness in the design of these services, the most underutilized resource we have as a society is the professionals who run these systems. They actively and consistently protect our infrastructures against disturbances, failures, and mistakes that could bring them down, including errors at policy and supervisory levels. They work, often heroically, against odds that many in the public, academia, and government can hardly appreciate.