Risk Analysis, Vol. 34, No. 7, 2014 DOI: 10.1111/risa.12241

Perspective

Improving Risk Management: From Lame Excuses to Principled Practice

Elisabeth Paté-Cornell[1] and Louis Anthony Cox Jr.[2,3,*]

The three classic pillars of risk analysis are *risk assessment* (how big is the risk and how sure can we be?), *risk management* (what shall we do about it?), and *risk communication* (what shall we say about it, to whom, when, and how?). We propose two complements as important parts of these three bases: *risk attribution* (who or what addressable conditions actually caused an accident or loss?) and *learning from experience* about risk reduction (what works, and how well?). Failures in complex systems usually evoke blame, often with insufficient attention to root causes of failure, including some aspects of the situation, design decisions, or social norms and culture. Focusing on blame, however, can inhibit effective learning, instead eliciting excuses to deflect attention and perceived culpability. Productive understanding of what went wrong, and how to do better, thus requires moving past recrimination and excuses. This article identifies common blame-shifting "lame excuses" for poor risk management. These generally contribute little to effective improvements and may leave real risks and preventable causes unaddressed. We propose principles from risk and decision sciences and organizational design to improve results. These start with organizational leadership. More specifically, they include: deliberate testing and learning—especially from near-misses and accident precursors; careful causal analysis of accidents; risk quantification; candid expression of uncertainties about costs and benefits of risk-reduction options; optimization of tradeoffs between gathering additional information and immediate action; promotion of safety culture; and mindful allocation of people, responsibilities, and resources to reduce risks. We propose that these principles provide sound foundations for improving successful risk management.

KEY WORDS: Causation; excuses; high-reliability organizations; learning; organizational design; risk assessment; risk attribution; risk management

1. INTRODUCTION: WHY DO CATASTROPHES HAPPEN? BAD LUCK IS RARELY THE WHOLE ANSWER

Even excellent risk management decisions usually do not reduce risk to zero. Failures in large, complex engineered systems can still occur under good risk management and in spite of flawless informed decisions under extreme conditions. However, failures of large systems and operations, such as the *Challenger* space shuttle disaster in 1986, the core breaches at the Fukushima Daiichi reactors in 2011, or the fire at the BP Mobile Drilling Unit Deepwater Horizon, also in 2011, are often rooted in flawed decision making at high levels of the organization, with disregard or poor use of available information and of the actual effects of the incentive structure.(1) Human, organizational, and management factors that predispose to human errors and operation failures can and should be addressed explicitly in a systems analysis to support risk management decisions and

[1]Stanford University, Stanford, CA, USA.
[2]NextHealth Technologies, Cox Associates, Denver, CO, USA.
[3]University of Colorado, Denver, CO, USA.
*Address correspondence to Louis Anthony Cox, NextHealth Technologies, Cox Associates, 503 Franklin Street, Denver, CO 80218, USA; tel: +1-303-388-1778; fax: +1-303-388-0609; Tony.Cox@NextHealthTechnologies.com.

reduce accident risks.[2] These decisions in the design, manufacturing, construction, and operation of complex systems affect their longevity and failure risk as well as their daily performance and productivity. Too often, however, catastrophic failures lead to a tight focus, after the fact, on assigning blame, and to expensive litigation over who knew (or should have known) what and when, and who should have done things differently. People design decision rules or operating practices, attract blame, and tend to be replaced.[3] Although hindsight bias may add asperity to the prosecution, the defense frequently finds that questionable (or ridiculous) arguments have been advanced for why things were done as they were, often in an effort to deflect blame from those who actually made the basic decisions.

Post hoc assignment of blame is prominent in our culture and in our justice system. It provides material for daily news and discussion: political spin on who is to blame for flooding and power outages in the wake of Hurricane Sandy; decisions about whether scientists should be jailed in Italy for opining, following small tremors, that earthquake risks in L'Aquila were modest if not negligible right before a lethal one struck;[4] and debate and recriminations over why Ambassador Stevens's calls for help did not prompt additional actions to save his life in Benghazi. More generally, the question is why precursors and near-misses were systematically ignored or misread, as they were, for instance, during drilling of the Macondo well.[5,6]

Psychologists, however, have convincingly documented that both prospective risk assessment and retrospective blame assignment are often flawed by various heuristics and biases.[7] In assessing risky prospects and projects, several classic fallacies have been identified in the literature. People systematically tend to overestimate benefits and underestimate costs (the planning fallacy).[4] Optimistic bias leads some to accept risky prospects that they might reject if their expectations and perceptions were more accurate (illusion of control and overconfidence). In retrospect, they may believe that whatever happened was inevitable (hindsight bias). They may prefer to continue with systems and assumptions in which large investments have already been made, rather than acknowledging, in light of new experience, that they are flawed and should be updated (sunk-cost bias). They may also blame bad luck and other people for undesired outcomes while attributing success to their own skill and efforts (self-serving bias), and altogether distort their understanding of what happened and why. Yet, identifying mistakes and their fundamental causes after a failure or a near-miss is key to learning effectively about what went wrong and how to do better in the future.[8] Therefore, to learn from costly experience how to improve risk management, it is essential to do realistic post-mortems, and not to let the opportunities for learning dissipate in a cloud of evasion and misdirection. Accordingly, this article focuses on a set of well-worn "lame excuses" often advanced to justify the decisions and behaviors that preceded catastrophic failures of complex systems. It then proposes some principles to improve risk assessment and management by cutting through these excuses to identify needed changes in design, operations, and culture.

2. "IT'S NOT OUR FAULT": SOME COMMON EXCUSES FOR BAD RISK MANAGEMENT

Many arguments advanced to deflect the blame for conspicuous failures are based on claims of unpredictability: "it was a black swan," or extremely low probability and not reasonably foreseeable; "it was a perfect storm," or extremely rare conjunction of conditions; and so forth.[9] More generally they attempt to put the blame elsewhere, claiming that other people, nature, or even supernatural influences (bad luck or an act of God) were responsible. Excuses in this category are seldom justified. The failure might well be rooted, for example, in flawed procedures as might be expected, in case of poor monitoring, from the theory of principal-agent problems.[10]

- "It was not our job: we were not paid for it," or "it was not our responsibility to report the deterioration of the structure or design errors that threaten its integrity."

This kind of reasoning can be attributed to poor safety culture and the detachment of the individuals from the proper functioning of an organization or a system. Misplaced faith in the *status quo* encourages us to accept how things are and how work gets done until something breaks badly enough to force us to recognize that we should have made changes sooner.[11]

- "It was an act of God," implying that natural forces were involved that are uncontrollable

[4]This was the case, for example, of the U.S. space shuttle, when the risk of mission failure was originally estimated in the order of 1/5,000.

and therefore the blame cannot be put on any human being.

Of course, this is a fallacy if human choices create the exposure to the risk and determine a system's capacity to survive naturally occurring stresses. Some of the initial claims regarding the accident at the Fukushima Daiichi nuclear power plant belong in that category. Earthquakes of magnitude greater than 8 (and the large tidal waves that came with them) had occurred several times in recorded history,[12] but the siting of the reactor and the initial choice of a 5.7 m tsunami design criterion were human decisions.

- "It was a black swan," meaning that it was both unprecedented and unimaginable before the fact.

This one has become a favorite excuse after Taleb's 2007 book comparing unimaginable (or, at least, unanticipated) events to the discovery of black swans in the 17th century by Dutch sailors who, until then, had seen only white ones.[13] Intended by the author to explain financial crises that only a few had seen coming, it has become a much-used descriptor for events such as the attack on the United States on 9/11/2001. Yet, an attack on the World Trade Center had occurred in 1993 and FBI agents had detected suspicious flying training in the preceding months.[5] With preparation, vigilance, and effort, much more is reasonably foreseeable than might be expected.[14]

Deliberate exercises in applying "prospective hindsight" (i.e., assuming that a failure will occur in the future, and envisioning scenarios of how it could happen) and probabilistic analysis using systems analysis, event trees, fault trees, and simulation can be used to overcome common psychological biases that profoundly limit our foresight.[14] These include anchoring and availability biases,[7,15] confirmation bias, group-think,[16] *status quo* bias, or endowment effects.[11,14] In the case of 9/11, for example, a similar terrorist attack had been mounted against an Air France flight in 1994 but thwarted by French security forces before it reached Paris. Yet, the experience had faded in the *status quo* of 2001.

- "It was a perfect storm," i.e., it required such a rare conjunction of unlikely events that it seemed that it could be safely ignored.

This phrase became popular after the publication of a book and the release of a movie describing a convergence of several storms in the northern Atlantic in 1991, in which a ship sank and the crew perished. Such conjunctions of unusual conditions, although individually rare, are collectively quite common;[9] but their probabilities are often underestimated because dependencies are unrecognized or misunderstood. In technical systems, these conjunctions sometimes come from common causes of failure such as external loads (e.g., extreme winds) or human errors that affect the whole system (e.g., flawed maintenance of all engines of an aircraft). They happen regularly in the economic sector and in supply chains, for instance, if difficult market situations and lean inventories are compounded by a natural catastrophe (independent events in this case). They are even more likely to occur in the financial industry or other tightly coupled organizations where the failure of one institution may have devastating effects on related ones. This may occur, for instance, because the failure factors are economically and statistically dependent (risk "contagion"), or because psychological reactions to one event are likely to cause further failures (bank runs).

- "It never did that before" (e.g., "the system had not blown up yet in previous near-misses, so, we thought that we were fine").

Ignoring near-misses and considering them successes is a classic reaction. Indeed, responding effectively to prevent near-misses from developing into full-blown catastrophes may reflect competence in managing hazardous situations and be a justifiable source of pride and confidence.[6] Yet, interpreting near-misses as justification for complacency can make one miss potentially valuable lessons. This was the case, for example, of tire blowouts on the SST Concorde, which had happened 57 times before a piece of rubber from one of them punctured the fuel tank and caused the death of everyone on board in July 2000.[7,17]

[5]The 9/11 commission report (2004) calls "failure of imagination" the fact that these attacks were not anticipated given past experience and new signals.

[6]This was the case of the U.S. Airways flight 1549 out of JFK, which, in January 2009, safely landed on the Hudson when it had been crippled minutes earlier by a bird strike and could not reach a close airport.

[7]This was also the case of the Deepwater Horizon platform, which was destroyed by explosions and fire in 2011 after close calls that should have alerted the operators and altered their course of action.[5] Yet, neither the operators nor the regulators saw reason to intervene earlier since no accident had happened.

- "Those in charge did not know, and could not reasonably have known, what was about to happen."

Excusable ignorance is a common plea in organizations that fail to pass messages, especially bad news, from the front lines to decisionmakers higher in the hierarchy. Indeed, "plausible deniability" is sometimes sought as a way to deflect responsibilities from top decisionmakers. Clearly, many signals are gathered every day and organizations need some filters to function effectively.[18] Incentives sometimes are such that an agent can rationally take shortcuts to meet a resource constraint rather than bringing the problem to the attention of the principal.[19] Ineffective elicitation and use of the information known to the members of an organization is both common and costly when the objective should be to align the goals of the employees and those of the organization. Therefore, changing incentives and procedures of deliberations and decisions can do much to elicit, exploit, and reward information that might otherwise remain hidden.[14]

- "We did not know that things had changed so much."

Status quo bias lulls into assuming that things will remain what they are: the environment will not change and our system will remain what it is. This is seldom true. A general failure to monitor the situation and its evolution (including, for instance, markets, competitors, and employees' performance) and to disregard or misinterpret signals that a new one is looming is a natural tendency.[14] This change in business environment is at the core of enterprise risk management and critical in these times of globalization and quick emergence of new technologies.

- "It was permitted and convenient, so we did it."

This was the case, for instance, of the design of the Piper Alpha platform, where, against common sense, the control rooms had been located above the production modules[20] for the convenience of operators who could easily go from one to another. Yet, an explosion in the production modules could (and eventually did) destroy possibilities of controlling the situation.

- "We took every reasonable precaution and followed standard operating procedure."

This would be a convincing excuse if the precautions and standard operating procedures were effectively applied to the situations for which they were intended. Yet, as potential catastrophes begin to unfold, the system, the environment, or the situation may change in ways that make the standard procedures and precautions inadequate. Blind rule-following may be disastrous if the assumptions behind the rules no longer hold. Deterioration, for example, affects automobile safety as well as that of airplanes, especially when maintenance on schedule fails to address some obvious parts such as the fuselage of an aircraft[8] and when there is not sufficient latitude for maintenance on demand. A quick shift from standard to crisis operation mode and creative improvisation to respond to the new and unforeseen situation may then be essential to avert a disaster.[3]

- "Everybody does it."

The everybody-does-it-defense, commonly used by teenagers, implicitly assumes that it is no one's obligation to examine current *status quo* practices and their implications, especially in a changing environment[11] and that imitating others justifies one's actions. Heedless imitation and herd-following behavior can, of course, multiply the consequences of failure if tight technological couplings make imitation easy and reflection more difficult. This is the case, for instance, of computer reactions to financial market situations, where automatic trading platforms and rules amplify the effects of initially minor price fluctuations. Similarly, destructive memes of harmful behaviors (whether teenage hyperventilation, recreational drug use, or copycat crimes) can spread rapidly through social media. Imitative learning is a powerful force of social coherence, but thoughtless imitation, as well as following orders without questioning their ethics and consequences, can destabilize systems, spread destructive habits, and amplify risks.

- "All signals indicated that we were doing fine."

Good test results may be falsely reassuring if the tests were performed in the wrong environment, or the sample size was too small. Overconfidence in biomedical circles has become so prevalent that

[8]This was the case, for example, of the Aloha Airlines in 1988 in Hawaii where part of the airplane fuselage broke apart in an explosive decompression and peeled off exposing the inside of the plane.[21]

commentators are starting to worry that science is failing us.[22] For lack of operating experience, engineering and physical models as well as expert opinions may be needed to complete the picture. Besides, "we used only the best/most credible/reliable results" may reflect a self-serving selection bias.[9]

- "But everyone agreed that it seemed a good idea at the time!"

A common reaction to failure is that everyone agreed beforehand to the course of action. "Our best experts approved, our reviewers and stakeholders (Congress, clients, funders, public, etc.) loved our analysis, our models were detailed and coherent, the results were perfectly clear it all seemed to make sense." In reality, such consensus may reflect "groupthink" and mutual influence among the players.[16] That clients receive analytical results that fit their interest may be the result of the incentives or the information that they have given to the analysts. These results are thus directly affected by motivated reasoning, confirmation and interpretation bias, and premature closure.

In practice, full validation of a risk analysis model in the traditional statistical sense may be impossible for new systems or systems that have changed because the statistics at the system level are not yet available. Yet, these models can be justified based on what is known of the functions and dependencies of the various components, and of the external factors (loads, environment) that affect their robustness.

- "It was an operator error."

Blaming the personnel in charge at the lower level is sometimes convenient to an organization. For instance, more than 60% of recent airline accidents have been blamed on pilot errors.[23] Yet, in some cases, a system design may be the accident root cause and must be corrected to allow for some pilot mistakes before they cause a crash. In other cases, the pilots may not have been sufficiently trained, for instance, to understand the signals that they get from electronic monitors or to operate in absence of these signals. Similarly, in the medical field, some accidents are directly caused by residents who do not have sufficient experience; but the true responsibility may lie with the supervisors if they are 15 minutes away from the operating room when they should be accessible in two minutes. More generally, managers sometimes blame the operators for their own failures when the leadership did not provide the training, the supervision, the information, or the incentives required to face critical situations.

[9]The same is true of data selection, for instance, as mentioned earlier, the choice of the Fukushima reactor designers to ignore all tsunami data older than 1,000 years.

3. SOME FOUNDATIONS FOR BETTER RISK MANAGEMENT

Valuable lessons about risk reduction can be derived from these accidents and from the excuses that are invoked after the fact. This section proposes some constructive foundations for improved risk management practices. They are selected from the management science and risk management literatures and reinforced by the cases described earlier. As witnessed by the extensive literatures on improving organizational decision making[14] and building high-reliability organizations,[24] we are far from the first to suggest principles and practices for overcoming the foregoing limitations. But we believe that the following recommendations, which emphasize causal understanding, quantitative methods, and deliberate design for learning, can add value to earlier improvement efforts.

- Understand the causes of the hazard, then its potential effects.

"Acts of God" such as earthquakes, tsunamis, or hurricanes often have a history, and their mechanisms and recurrence times can be at least partly understood. There is no more valuable tool for reducing risk than correctly analyzing and understanding causes.[20,25] This requires identifying the factors affecting the performance of people and systems, and their technical characteristics, as well as the environment in which they operate. It is essential, in particular, to understand who can control what, and how incentive and information structures affect agents' decisions, response times, and error rates.[2] In the oil industry, for instance, rewarding production alone will likely discourage interrupting operations when immediate maintenance is needed. A general nonchalant attitude towards safety can be corrected by training, changing incentives, and making top managers aware of the true costs of risk and of opportunities to manage them more effectively. Because risks are often invisible until they have struck, it is an easy and common practice to dismiss them or to leave

them to actuaries and insurers to price and manage their financial side. The human costs, however, cannot be truly redressed after the fact.

Risk analysis can make the cumulative impacts of risks on a company and its employees, customers, and partners more vivid. It can quantify, with emphasis on uncertainties, avoidable possible losses as well as potential changes in insurance premiums and cost of capital, and can highlight cost-effective opportunities for enterprise risk management. Other risk factors are simply facts that can only be accounted for. The realities of shrinking global markets may not be changeable, but some level of diversification, innovation, and decoupling meant to protect a system from cascading failures may go a long way towards reducing risks.

More complex cases are those in which the risk is caused in large part by the activities or the threats of intelligent adversaries such as drug gangs or insurgents. The key to analyzing the risk that they present is to understand who they are, their intent, their capabilities, and the types of response that one is willing to implement given the possibilities of deterrence but also escalating conflicts.[9] The issue here is thus to address not only the symptoms (e.g., immediate threats of attacks) but also the basic problems, although sometimes, as in saving a patient, treating the symptoms may have to come first.[10]

- Characterize risk magnitudes and uncertainties.

Once a hazard—a possible source of risk—is identified, the next step is to try to figure out what one is dealing with and how large the risk might be.[26] Are probabilities and magnitudes of potential losses large enough, compared to the costs of reducing them, to warrant urgent attention, or are they small enough that waiting to address them implies little possible loss? This is where science reporters often fail as risk communicators, by publishing articles exclaiming that exposures have been "linked" to adverse effects, but without noting the absolute sizes of the risks involved or, often, even failing to check whether the claimed "links" are causal.[27] If the risk is uncertain, can this uncertainty be clarified for a cost that is less than the value of information obtained by additional investigation, and the benefits of improved decisions that it would make possible?[11] If so, acting quickly out of concern about uncertain risks may be less prudent than first collecting better information.

This quantification of the risk and of the associated uncertainties may be a difficult task depending on the nature and the relevance of the available evidence. Quantitative risk assessment (QRA) or probabilistic risk analysis (PRA), developed originally for engineered systems, involve all available information that can help to answer practical risk management and uncertainty reduction questions. These methods are based both on systems analysis and on probability, including essential functions, feedback loops and dependencies caused by external events, and common causes of failure.[28] For a structural system, these external events can be earthquakes or flooding that affect simultaneously several subsystems and components. In these cases, the risk analysis is based on an assessment of the probability distributions of loads and capacities, and computation of the chances that the former exceeds the latter. When needed, the results should include, and if needed display separately, the effects of aleatory as well as epistemic uncertainties[12] to accurately characterize the limitations of the analytical results.

Realistic PRAs, including those for failures of complex technological systems, must also include human and organizational factors. This analysis can be achieved, starting from a functional and probabilistic analysis of system failures, then considering the potential decisions and actions of the actors directly involved (errors as well as competent moves), and linking these to the environment created by the management.[2] This requires examining in details the procedures, the structure, and the culture of the organization, including the information passed along to operators, the resource constraints, and the incentive system.

[10] That was the dilemma in managing the financial crisis of 2008, when governments were facing several options: some favored the injection of stimulus capital first then the regulation of banking reserves and others the reverse.

[11] Of course, additional information may also increase the uncertainties about a risk and justify increased safety measures.

[12] Aleatory uncertainties are caused by randomness. They remain even when the probability of an event is known with certainty. Epistemic uncertainty refers to imperfect basic knowledge: it is the uncertainty about the probability of an event, for example, when several hypotheses are possible, when experts disagree, etc. Separating the two types of uncertainty in the display of the results as a family of risk curves is especially useful either when the analysis applies to several systems and/or over several time periods, or when the decisionmaker is "ambiguity averse" and epistemic uncertainties affect his/her preference function.[29–31]

These risk analyses, imperfect as they are, can be invaluable tools in identifying risks that were not considered or were underestimated before,[13] and in setting priorities among safety measures.

- Identify what can be done to reduce risk, and with what costs and benefits. Candidly show uncertainties about the answers.

For risks that are worth acting on now, the next step is to identify the risk mitigation alternatives and challenges in an implementing them, and to assess how much difference they would make. This is an essential step in rational ("System 2") thinking, which is often entirely missing from emotional and intuitive (e.g., outrage-driven) "System 1" responses to risk.[32] Our emotions, often based on recent events that have been widely advertised, may tell us that a situation is unacceptable and urge us to adopt a change to address the problem ("Ban that product!"). Indeed, the "precautionary principle" to implement such bans systematically when there remain uncertainties has been adopted by some governments.[33] Yet, reasonable (if imperfect) calculations of how much difference alternative interventions would actually make are needed to guide risk management actions to achieve desired results.

The question, again, is to ensure that these assessments are as objective as humanly possible. Separation of facts and values may sometimes require that an analyst waste no time working for someone who will disregard fact-based results, or who insists on constraining or influencing them based on values and preconceptions, for instance, by forcing some inputs. For example, the Environmental Protection Agency (EPA) required its experts to express their uncertainty about "lives saved per microgram per cubic meter" of reduction in fine particulate matter by using Weibull distributions, which are constrained to show a 100% probability of positive life savings (no matter what the data say). In that case, analysts might have insisted on being given the flexibility to use other distributions that could also assign positive probability to zero (or negative) values if that is what the data indicate.[34] An illustration of the "risk of no risk analysis" is, again, the choice of a surprisingly low tsunami design criterion at the Fukushima Daiichi nuclear reactor, despite a recorded history of such events over more than a thousand years as mentioned earlier. Insisting that risk management be guided by risk analyses is particularly critical for new nuclear reactors, whose design criteria must meet the characteristics of each site and the local hazards of external loads at a time where 68 new nuclear power plants are under construction across the world.

- Assess the urgency and value of information: Is collecting (or waiting for) additional information before acting more costly than it is worth?

The value of gathering new information depends on the possibility that it will permit better (e.g., higher expected utility) decision making. It therefore depends on the uncertainties faced by the decision-maker, as well as his or her risk attitude.[36] When deciding the urgency of action and evaluating whether to wait for additional information, a risk manager should consider:

(1) Is the system stable? If not, how quickly is it deteriorating?
(2) Are the benefits of gathering or waiting for additional information, which might improve the decision, expected to outweigh the costs?
(3) What does one know (and can expect) of new technologies that may allow elimination of the risk altogether, for instance, by replacing a hazardous technology at an acceptable cost?

An example of the first consideration—deterioration—is the speed at which one might expect, for instance, deterioration of the climate with and without proposed interventions, with an assessment of its likely impacts (both beneficial and harmful) on human population in different parts of the globe. Examples of the second consideration—value of information and optimal stopping—include the choice of whether to perform additional medical tests before an operation, whether to engage in more invasive medical tests on a routine basis, or whether to delay a repair in a car or a chemical factory to ensure that the potential risk reduction benefits justify the costs of an immediate fix. An example of the third type—risk reduction by the substitution of a new technology—might be the decision to live with the consequence of coal burning to generate electricity, after closing nuclear plants and before solar energy becomes truly economical, understanding the pace and the costs of such new

[13]This was the case of the importance of the auxiliary feedwater systems in the safety of nuclear power plants, as emphasized in the first PRAs performed for these systems.[35]

development and the actual potential for future risk reduction.[14]

- Anticipate, monitor, and prepare for rare and no-so-rare events.

Not-so-rare events can generally be analyzed relatively easily because there is a base of experience, either with the system itself (e.g., classic earth dams) or with its components and subsystems (e.g., an aircraft that has been in service for decades, so that there is substantial operating experience with its subsystems and components). Rare events that result from the conjunction of known components ("perfect storms") with or without dependencies may be a bit more difficult to analyze if either the probabilities or the dependencies among them are difficult to establish.

Rare or unknown events for which there is little or no information as to whether or not they can actually occur are especially difficult to manage sensibly. Starting with the most difficult case, genuine "black swans" that one knows nothing about and cannot reasonably anticipate, the best strategy may be to monitor for signals of unusual occurrences (e.g., of new diseases) and to put in place a "resilient" structure of organizational connections, financial reserves, and access to human intelligence and knowledge that allows for quick, creative local responses.[9,38] For instance, new types of flu occur on average every two years. A system managed by the World Health Organization permits monitoring and sharing of information across countries and identification of virus types. Although imperfect, that system allows relatively quick response to new strains of flu viruses such as H1N1. But the slow response to the spread of AIDS illustrates the difficulty of identifying and responding to a new type of pathogen.[15] Managing the more straightforward case of "perfect storms" is easier in that it involves "anticipating the unexpected" but imaginable,[39] and observing conjunctions of dangerous events such as the convergence of storms, loads on a system, or economic problems.

- Deliberately test and learn.

An avoidable pitfall in organizational risk management is to neglect to deliberately acknowledge and test key assumptions, to learn from experience, and to capture data and lessons for future reference as opportunities arise.[14] The world is full of "natural experiments"—unplanned but potentially highly informative shocks to systems—which can be used to test and refine critical assumptions underlying risk assessment and risk management—if we remember to do so. For example, if air pollution in Beijing during a winter inversion soars to dozens of times higher concentrations than are permitted in the United States, but mortality rates do not increase correspondingly, the common risk assessment assumption that life expectancies decrease in direct proportion to pollution concentrations[40,41] should be revisited in light of the new data. Nor, outside the domain of human health, is it always necessary to wait for natural experiments. Intelligence and security professionals know that deliberately testing their systems (e.g., by "red teaming," which grew more popular after 9/11) and trying to bypass or disable safeguards is a key to active identification and elimination or mitigation of exploitable vulnerabilities.

- Learn from near-misses and identify accident precursors

Many accidents have been preceded by close-calls, for instance, when only one event did not occur in a known accident sequence. That these have not turned into a disaster has sometimes been viewed as evidence that the system needs no correction. Pro-active risk management, of course, is the best way to avoid disasters. Yet, industries and regulators seem to believe at times that they should not intervene because the system has worked and no disaster has occurred—even if only by chance. The experts who claimed after a small tremor that all was safe in L'Aquila (Italy), where a large earthquake then occurred shortly after, were relying on a recent occurrence of a false alert and failed to communicate to the public the fact that small shocks can also be precursors of large ones. In the case of the 2011 accident at the Macondo well, the regulators as well as the three companies involved did not intervene when they knew that some worrisome near-misses had occurred (presuming that they were doing well enough), and decided to ignore precursors and test results[5] presumably for a variety of immediate benefits.

- Establish and maintain a culture of safety.

[14] In the case of coal burning, this would require an understanding of the true health and safety effects of it, as well as the costs, timing, and risk reduction effects of carbon sequestration.[37]

[15] The HIV retrovirus had been present in the human population for decades before it was clearly identified and researched.[42]

It is possible to deliberately create and maintain a safety culture that reduces accident risks and losses. This requires acting beyond the classic ritual statements of "safety first." A safety culture starts at the head of an organization, with a true commitment to recognize and properly manage unavoidable tradeoffs, and by training those who are closest to operations to make appropriate decisions when needed. Therefore, the deliberate design and development of highly-reliable organizations typically emphasize adopting a vigilant, risk-aware mind set, and instilling the following five principles throughout the organization: preoccupation with failure at all levels and by all hands; reluctance to jump to conclusions or simplify interpretations of data and anomalies; sensitivity to operations at all levels; commitment to resilience; and deference to relevant expertise, rather than to authority.(25)

A key part of a safety culture thus involves the incentives provided by the management. The structure and the procedures of organizations such as an oil company reflect an attitude at the top of the corporation that permeates all levels of the organization. Concretely, the incentives, constraints, and directions explicitly communicated to employees shape their decisions, especially when they have little time to react or little information to evaluate their decisions. This was one of many problems at the Fukushima Daiichi nuclear power plant, where operators had to wait for hours before deciding on their own to flood a crippled reactor. It was also true on the Deepwater Horizon platform, where ignoring negative pressure tests results contributed to the already high risks of an accident.(5) Economic incentives that encourage motivated reasoning may thus distort risk-taking and risk-management decisions. As pointed out earlier, organizations that reward exclusively a production level and *de facto* penalize those who slow down production, put at risk not only their employees but also possibly, the general public both from a safety and a financial point of view.

- Put the right people in the right place with the right knowledge, incentives, and resources.

Training and learning are two of the most important requirements for effective risk management. Risk analysis can clarify the effectiveness and performance of risk management decisions and their importance in affecting outcomes. The results, in turn, allow assessing where additional resources and training, as well as changes in incentives and responsibilities, are most likely to pay off in reduced risks and improved performance. Having examined what drives the operators of a complex system (e.g., the conductors of high-speed trains), one can also review management procedures, structure, and culture for fitness to meet the needs of both regular operations and responses to crisis. In normal operations, disciplined rule-following can protect us against the temptations, heuristics, and biases that undermine so much human decision making. These range from succumbing to short-run impulses that we may come to regret such as hyperbolic discounting,[16](27,43) to letting fears, doubts, and desire control decisive actions that, upon reflection, no one favors. On the other hand, when reality turns to crisis or emergency situations, narrow rule-following can lead to blinkered vision and to abdication of the responsibility, creativity, and active cooperation needed for adaptive responses to the unexpected.(3) A key challenge in many organizations is to know when to shift from normal procedures to emergency response, which implies that crisis signals have been observed and transmitted in time for quick effective response.

In a context where operators can face unexpected delays and problems, it is essential to provide people with reasonable amounts of resources and deadlines, and to be willing to make adjustments. Otherwise, agents might satisfy the managers by cutting corners in ways that they may not even imagine until and unless they see consequent failures.(10) Therefore, when managers set these constraints they have to ask themselves what are their "shadow price," that is, by how much would one reduce the failure risk if one relaxed that constraint by one unit (one more day?); or on the contrary, whether one can tighten these constraints at a net benefit.

- Clearly define leadership and responsibilities.

Key to the effectiveness of managers is their leadership in providing role models, and setting the tone for the organization's performance. Leadership in a risk management context implies not only having (or deferring to) relevant knowledge and authority but also establishing clear lines of

[16]Hyperbolic discounting describes "present-biased" preferences in which the same delay in reward is valued at different rates in the present than in the future (e.g., the case where $10 now is preferred to $20 in one year, while $20 in six years is preferred to $10 in five years).

accountability and building trust from the people involved that their leaders can and will make proper and prudent decisions in difficult situations.

Who is responsible for avoiding accidents and mishaps? There are often several lines of responsibility and accountability, which should be properly defined and coordinated. The feeble defense of "responsible but not guilty" was used, for instance, by a high government official head of a health ministry in Europe in 1991, after contaminated blood infected a large number of people. The question, of course, is: What constitutes guilt on the part of a leader who fails to define proper procedures and ensure their application? Another failure of leadership can occur when a conflict of authority emerges from a two-head structure. For instance, a surgeon and an anesthesiologist who disagree when neither of them has the ultimate decision-making power can cause (and have caused) the death of a patient.(44) It may be possible to pinpoint precisely an error at the bottom of the organizational hierarchy that has led to an accident sequence.[17] But, as in the case of rogue traders, the overall question of supervision, incentives, and safety culture emanates directly from the leadership of the company and the regulators.

Leadership is thus a key ingredient of a solid system of risk management decision making in which the decisionmaker hears the message on time, understands it (and the uncertainties involved if any), and is able and ready to act when needed. The decisionmakers must be willing to know the truth, to make difficult choices and tradeoffs of what is tolerable and what is not, and to decide when it is time to shift from regular operations to crisis management with the ability to make quick, well-informed decisions. When the risk is borne by a group of people, this requires an acceptable collective decision process, able to balance the interests and the safety of different groups, and the overall costs and benefits.

- Sharing knowledge and experience across organizations.

Not all risk management responsibilities can or should be defined within a specific organization. Distributed control of risks, shared among multiple organizations or individuals, also creates a need for legal and institutional frameworks to clearly define and enforce rights and duties. Clarifying whose responsibility it is to avoid risks that arise from joint decisions[18] can reduce the average costs of risk management. To provide a rational basis for coordinating liability and incentives to reduce the costs of risk externalities and jointly caused risks in a society of interdependent agents, one might adopt several possible principles. From the economic analysis of law,(45) the Learned Hand formula states that parties should take additional care if and only if the expected marginal benefit of doing so exceeds the expected marginal cost.(46) Similarly, the "cheapest cost avoider" principle states that the party who can most cheaply avoid a jointly created risk should do so.

Accidents sometimes reveal the existence of information in some parts of industry that could have saved others. Some organizations successfully permit sharing that critical information. The Institute of Nuclear Power Operations (INPO) provides a practical example of such an organization.(47) Created in the wake of the Three Mile Island accident, INPO provides a forum where industry managers can discuss existing problems behind closed doors with the support of the regulator (in the United States, the NRC). It has the role of an internal watchdog, regularly rating each power plant. These ratings, in turn, influence the insurance rate of the plants, thus promoting strong incentives for excellence in safety. What makes such an organization successful is the combination of peer pressure, of a forum for internal discussion of potential problems, blunt assessment of plant performance, and the "teeth" provided by financial incentives. What sometimes makes it difficult to generalize the model is the competition among the organizations involved and the global nature of some industries such as the oil market.

4. CONCLUSIONS

Successful risk management is usually a cooperative enterprise. Successful cooperation, in turn, requires moving past blame-casting and excuse-giving to understand the causes and conditions that contribute to catastrophes and improve the system or, conversely, that promote safety in the face of unanticipated challenges. Prominent among the addressable drivers of safety are vigilance and readiness to perceive and respond to anomalies, determination, ability to learn from experience, eagerness to continually probe and update

[17] In the case of the Piper Alpha accident in 1988, a young worker made the mistake of leaving the work of fixing a pump unfinished at the end of a day and failed to tag the pump as remaining to be fixed.

[18] An example is a consumer's decision to stir his soup with a hair dryer, together with the manufacturer's decision not to affix a label warning against such uses.

assumptions in light of new information, and capacity to adapt creatively and cooperatively when conditions change. Clear lines of duty and responsibility for risk avoidance, together with discipline and training in following well-conceived procedures and rules for routine safe operation of complex systems, are key contributors to safety cultures that work. At the same time, having the wisdom, incentives, know-how, and experience in team problem-solving required to step outside such boundaries and improvise when needed is essential for successful risk management in the face of novel threats. These are generally teachable and learnable skills.

We propose that improved practices of risk analysis, quantification, and management should be built on technical and cultural foundations, which encompass expertise both in reducing routine risks and in responding to novel ones. Such risk management practices should rely less on blame-casting and excuse-making than in the past. They will need to acknowledge that human error is not necessarily the main driver of failures in an increasingly complex and interconnected world, and that systems should be designed to withstand such errors. Unprecedented hazards, fat-tailed distributions, and risk contagion leading to cascading failures are increasingly recognized as drivers of some of the most conspicuous modern risks, from power outages to epidemics to financial failures. Improved risk management practices should thus increasingly rely on intelligent engagement with our uncertain and changing world.

They should build on the key principles we have touched upon: leadership and accountability; robust design (decoupling subsystem whenever possible); vigilant and open-minded monitoring; continual active testing of assumptions and systems; deliberate learning; optimal tradeoffs of the costs and benefits of gathering further information before acting; well-trained and disciplined habits of coordination; and ability to cooperate quickly and effectively in response to new threats. These principles have been valuable foundations for effective risk management when they were applied in the past. They should become common practice in the future.

ACKNOWLEDGMENTS

The authors thank Area Editor Warner North and two anonymous referees for useful suggestions that led to a shorter, clearer exposition.

REFERENCES

1. Paté-Cornell ME. Organizational aspects of engineering system safety: The case of offshore platforms. Science, 1990; 250:1210–1217.
2. Murphy DM, Paté-Cornell ME. The SAM framework: A systems analysis approach to modeling the effects of management on human behavior in risk analysis. Risk Analysis, 1996;6(4):501–515.
3. Harford T. Adapt: Why Success Always Starts with Failure. New York: Farra, Straus and Giroux, 2011.
4. International Seismic Safety Organization. Position Statement on Earthquake Hazard Assessment and Design Load for Seismic Safety. Arsita, Italy, 2012.
5. National Academy of Engineering. Macondo Well-Deepwater Horizon Blowout: Lessons for Improving Offshore Drilling Safety, Report to the Department of Interior. Washington, DC: National Academy Press, 2012.
6. National Academy of Engineering. Accident Precursor Analysis and Management: Reducing Technological Risk Through Diligence. Washington, DC: National Academies Press, 2004.
7. Kahneman D, Tversky A. Judgment under uncertainties: Heuristics and biases. Science, 1974; 85(4157):1124–1131.
8. Paté-Cornell ME. Accident precursors. In Cochran JJ (ed). The Wiley Encyclopedia of Operations Research and Management Science. New York: Wiley Pub, 2009.
9. Paté-Cornell ME. On black swans and perfect storms: Risk analysis and management when statistics are not enough. Risk Analysis, 2012; 32(11):1823–1833.
10. Garber RG, Paté-Cornell ME. Shortcuts in complex engineering systems: A principal-agent approach to risk management. Risk Analysis, 2012; 32(5):836–854.
11. Hammond JS, Keeney RL, Raiffa H. The Hidden Traps in Decision-Making. Cambridge: Harvard Business Review, 1998.
12. Epstein W. A Probabilistic Risk Assessment Practitioner Looks at the Great East Japan Earthquake and Tsunami. A Ninokata Laboratory White Paper. Tokyo: Tokyo Institute of Technology, 2011.
13. Taleb NN. The Black Swan: The Impact of the Highly Improbable. New York: Random House, 2007.
14. Russo JE, Schoemaker PJH. Decision Traps: Ten Barriers to Brilliant Decision-Making and How to Overcome Them. New York: Doubleday Publishers, 1990.
15. Kahneman D. Thinking Fast and Slow. New York: Farrar, Straus, and Giroux, 2011.
16. Janis IL. Groupthink: Psychological Studies of Policy Decisions and Fiascoes. Boston: Houghton Mifflin, 1984.
17. BEA, French Accident Investigation Bureau. Final report on the 2000 Concorde accident. Paris, 2002.
18. Lakats LM, Paté-Cornell ME. Organizational warnings and system safety: A probabilistic analysis. IEEE Transactions on Engineering Management, 2004; 51(2):183–196.
19. Paté-Cornell ME, with the collaboration of Guikema S, Kucik P, Caswell D, Garber R. Games, risks and analytics: Several cases involving national security and management situations. Decision Analysis, 2012; 9(2):186–203.
20. Paté-Cornell ME. Learning from the Piper Alpha accident: A post-mortem analysis of technical and organizational factors. Risk Analysis, 1993; 13(2):215–232.
21. National Transportation Safety Board. Aircraft accident report—Eastern Airlines, Flight 855, Lockheed L-1011, -N334EA, Miami International Airport, Miami, Florida, May 5, 1983; 1984.
22. Lehrer J. Trials and errors: Why science is failing us. Available at: http://www.wired.co.uk/magazine/archive/2012/02/features/trials-and-errors?page=all, Accessed January 28, 2012.

23. Federal Aviation Administration. Aviation Rule Making Advisory Committee Report, FAA. Washington, DC: Federal Aviation Administration, 2013.
24. Weick KE, Sutcliffe KM. Managing the Unexpected: Resilient Performance in an Age of Uncertainty. 2nd ed. New York: John Wiley & Sons, 2007.
25. Cox LA Jr. Improving Risk Analysis. New York: Springer, 2013.
26. Kaplan S, Garrick JB. On the quantitative definition of risk. Risk Analysis, 1981; 1(1):11–27.
27. Gardner D. The Science of Fear: How the Culture of Fear Manipulates Your Brain. New York: Penguin Group, 2009.
28. Paté-Cornell ME. Probabilistic risk assessment. In Cochran JJ (ed). The Wiley Encyclopedia of Operations Research and Management Science. New York: Wiley Pub, 2009.
29. Paté-Cornell ME, Davis DB. (1994). A challenge to the compound lottery axiom: A two-stage normative structure and comparison to other theories. Theory Decision, 1994; 37(3):267–309.
30. Paté-Cornell ME, Fischbeck PS. Probabilistic interpretation of command and control signals: Bayesian updating of the probability of nuclear attack. Reliability Engineering and System Safety, 1995; 47(1):27–36.
31. Paté-Cornell ME. Uncertainties in risk analysis: Six levels of treatment. Reliability Engineering and System Safety, 1996; 54:95–111.
32. Sanfey AG, Chang LJ. Multiple systems in decision-making. In Tucker WT, Ferson S, Finkel A, Long TF, Slavin D, Wright P, eds. Strategies for risk communication: evolution, evidence, experience. Annals of the New York Academy of Science, 2008; 1128:53–62.
33. European Commission. Communication from the Commission on the Precautionary Principle. Brussels, Belgium, 2000.
34. Cox LA Jr. Reassessing the human health benefits from cleaner air. Risk Analysis 2012; 32(5):816–829.
35. US Nuclear Regulatory Commission. Reactor Safety Study: An Assessment of Accidental Risks in U.S. Commercial Nuclear Power Plants. Rasmussen, NC. Washington, DC: NUREG-75/014 (WASH-1400), 1975.
36. Howard RA. Decision analysis: Perspectives on inference, decision, and experimentation. Proceedings of IEEE, 1970; 58(5):823–834.
37. Azar C et al. Carbon capture and storage from fossil fuels and biomass—Costs and potential role in stabilizing the atmosphere. Climate Change, 2006; 74:1–3.
38. Cox LA Jr. Community resilience and decision theory challenges for catastrophic events. Risk Analysis, 2012; 32(11):1919–1934.
39. Augustine NR. Augustine Laws. Washington, DC: American Institute of Aeronautics and Astronautics, 1996.
40. Pope CA 3rd, Ezzati M, Dockery DW. Fine-particulate air pollution and life expectancy in the United States. New England Journal of Medicine, 2009; 360(4):376–386.
41. Correia AW, Pope, CA, 3rd, Dockery DW, Wang Y, Ezzati M, Dominici F. Effect of air pollution control on life expectancy in the United States: An analysis of 545 U.S. counties for the period from 2000 to 2007. Epidemiology, 2013; 24(1): 23–31.
42. Gallo RC. A reflection on HIV/AIDS research after 25 years. Retrovirology, 2006; 3:72.
43. Lehrer J. How We Decide. New York: Houghton Mifflin Harcourt, 2009.
44. Paté-Cornell ME, Lakats LM, Murphy DM, Gaba DM. Anesthesia patient risk: A quantitative approach to organizational factors and risk management options. Risk Analysis, 1997; 17(4):511–523.
45. Posner R. Economic Analysis of Law, 5th ed. New York: Aspen Publishers, 1998.
46. Feldman A, Kim J. The Hand rule and United States v. Carroll Towing Co. reconsidered. American Law and Economics Review, 2005; 7(2):523–543.
47. Reilly W. Valuing safety even when the markets do not notice. In Goldwyn DL, Kalicki J (eds). Energy and Security, 2nd ed. Baltimore, MD: Johns Hopkins University Press, 2013.