



An Introduction to the Crisis

COPYRIGHTED MATERIAL



Healthy Skepticism for Risk Management

*It is far better to grasp the universe as it really is than to persist in delusion,
however satisfying and reassuring.*

—CARL SAGAN

Everything's fine today, that is our illusion.

—VOLTAIRE

Any new and rapidly growing trend in management methods should be considered with healthy skepticism, especially when that method is meant to help direct and protect major investments and inform key public policy. It is time to apply this skepticism to the “risk management” methods meant to assess and then mitigate major risks of all sorts. Many of these methods are fairly new and are growing in popularity. Some are well-established and highly regarded. Some take a very soft, qualitative approach and others are rigorously quantitative. But for all of these methods, we have to ask the same, basic questions:

- Do any of these risk management methods work?
- Would anyone in the organization even know if they didn't work?
- If they didn't work, what would be the consequences?

4 CHAPTER 1 HEALTHY SKEPTICISM FOR RISK MANAGEMENT

For most organizations, the answers to these questions are all bad news. Natural, geopolitical, and financial disasters in the first few years of the 21st century have, perhaps only temporarily, created a new awareness of risk among the public, businesses, and lawmakers. This has spurred the development of several risk management methods, in both financial and non-financial sectors. Unfortunately, when these methods are measured rigorously, they don't appear to work. Most of the new non-financial methods are not based on any previous theories of risk analysis and there is no real, scientific evidence that they result in a measurable reduction in risk or improvement in decisions. Where scientific data does exist, the data shows that most methods fail to account for known sources of error in the analysis of risk or, worse yet, add error of their own. Even in the financial sector and other areas that use the most sophisticated, quantitative methods, there is a growing realization that certain types of systematic errors have undermined the validity of their analysis for years.

The answer to the second question (whether anyone would know that the risk management system has failed) is also *no*; most managers would not know what they need to look for to evaluate a risk management method and, more likely than not, can be fooled by a kind of “placebo effect”¹ and groupthink about the method. Even under the best circumstances, where the effectiveness of the risk management method itself was tracked closely and measured objectively, adequate evidence may not be available for some time. A more typical circumstance, however, is that the risk management method itself has no performance measures at all, even in the most diligent, metrics-oriented organizations. This widespread inability to make the sometimes-subtle differentiation between methods that work and methods that don't work means that ineffectual methods are likely to spread. Ineffectual methods may even be touted as “best practices” and, like a dangerous virus with a long incubation period, are passed from company to company with no early indicators of ill effects until it's too late.

COMMON MODE FAILURE

Finally, to answer the question about the consequences of unsound risk management methods, I'll use an example from a historic air-travel disaster to explain a concept called *common mode failure* (a concept from one of the more scientific approaches to risk analysis). In July 1989, I was the

commander of the Army Reserve unit in Sioux City, Iowa. It was the first day of our two-week annual training and I had already left for Fort McCoy, Wisconsin, with a small group of support staff (the “advance party”). The convoy of the rest of the unit was going to leave that afternoon, about five hours behind us. But just before the main body was ready to leave for annual training, the unit was deployed for a major local emergency.

United Airlines flight 232 to Philadelphia was being redirected to the small Sioux City airport because of serious mechanical difficulties. It crashed, killing 111 passengers and crew. Fortunately, the large number of emergency workers available and the heroic airmanship of the crew helped make it possible to save 185 onboard. Most of my unit spent the first day of our annual training collecting the dead from the tarmac and the nearby cornfields.

During the flight, the DC-10’s tail-mounted engine failed catastrophically, causing the fast-spinning turbine blades to fly out like shrapnel in all directions. The debris from the turbine managed to cut the lines to *all three* redundant hydraulic systems, making the aircraft nearly uncontrollable. Although the crew was able to guide the aircraft in the direction of the airport by varying thrust to the two remaining wing-mounted engines, the lack of tail control made a normal landing impossible.

Aviation officials would refer to this as a “one-in-a-billion” event² and the media repeated this claim. But since mathematical misconceptions are common, if someone tells you that something that just occurred had merely a one-in-a-billion chance of occurrence, you should consider the possibility that they calculated the odds incorrectly.

The type of event that caused the crash is called a *common mode failure*, because a single event caused the failure of multiple components in a system. If they had failed independently of each other, the failure of all three would be extremely unlikely. But because all three hydraulic systems had lines near the tail engine, a single event could damage all of them. The common mode failure wiped out the benefits of redundancy.

Now consider that the cracks in the turbine blades would have been detected except for what the National Transportation Safety Board (NTSB) called “inadequate consideration given to human factors” in the turbine blade inspection process. Is human error more likely than one in a billion? Absolutely; in a way, that was an *even more common* common mode failure in the system.

6 CHAPTER 1 HEALTHY SKEPTICISM FOR RISK MANAGEMENT

But the common mode failure hierarchy could be taken even further. Suppose that the risk management method itself was fundamentally flawed. If that were the case, then perhaps problems in design and inspection procedures would be very hard to discover and much more likely to materialize. Now suppose that the risk management methods not just in one airline but in most organizations in most industries were flawed. The effects of disasters like Katrina and the financial crisis of 2008/9 could be inadequately planned for simply because the methods used to assess the risk were misguided. Ineffective risk management methods that somehow manage to become standard spread this vulnerability to everything they touch.

The ultimate common mode failure would be a failure of risk management itself. A weak risk management approach is effectively the biggest risk in the organization.

If the initial assessment of risk is not based on meaningful measures, the risk mitigation methods—even if they could have worked—are bound to address the wrong problems. If risk assessment is a failure, then the best case is that the risk management effort is simply a waste of time and money because decisions are ultimately unimproved. In the worst case, the erroneous conclusions lead the organization down a more dangerous path that it would probably not have otherwise taken.

The financial crisis occurring while I wrote this book was another example of a common mode failure that traces its way back to the failure of risk management of firms like AIG, Lehman Brothers, Bear Stearns, and the federal agencies appointed to oversee them. Previously loose credit practices and overly leveraged positions combined with an economic downturn to create a cascade of loan defaults, tightening credit among institutions, and further economic downturns. If that weren't bad enough, poor risk management methods are used in government and business to make decisions that not only guide risk decisions involving billions—or trillions—of dollars, but are also used to affect decisions that impact human health and safety.

What happened is history. But here are just a few more examples of major, risky decisions currently made with questionable risk assessment

methods, some of which we will discuss in more detail later. Any of these, and many more, could reveal themselves only after a major disaster in a business, government program, or even your personal life:

- The approval and prioritization of investments and project portfolios in major U.S. companies
- The evaluation of major security threats for business and government
- The decision to launch the space shuttle
- The approval of government programs worth many billions of dollars
- The determination of when additional maintenance is required for old bridges
- The evaluation of patient risks in health care
- The identification of supply chain risks due to pandemic viruses
- The decision to outsource pharmaceutical production to China

Clearly, getting any of these risks wrong would lead to major problems—as has already happened in some cases. The individual method used may have been sold as “formal and structured” and perhaps it was even claimed to be “proven.” Surveys of organizations even show a significant percentage of managers who will say the risk management program was “successful” (more on this to come). Perhaps success was claimed for the reason that it helped to “build consensus,” “communicate risks,” or “change the culture.”

Since the methods used did not actually measure these risks in a mathematically and scientifically sound manner, management doesn’t even have the basis for determining whether a method works. Surveys about the adoption and success of risk management initiatives are almost always self-assessments by the surveyed organizations. They are not independent, objective measures of success in reducing risks. If the process doesn’t correctly assess and mitigate risks, then what is the value of building consensus about it, communicating it, or changing the culture about it? Even if harmony were achieved, perhaps communicating and building consensus on the wrong solution will merely ensure that one makes the big mistakes faster and more efficiently.

Fortunately, the cost to fix the problem is almost always a fraction of a percent of the size of what is being risked. For example, a more realistic evaluation of risks in a large IT portfolio worth over a hundred million

8 CHAPTER 1 HEALTHY SKEPTICISM FOR RISK MANAGEMENT

dollars would not have to cost more than half a million—probably a lot less. Unfortunately, the adoption of a more rigorous and scientific management of risk is still not widespread. And for major risks such as those in the previous list, that is a big problem for corporate profits, the economy, public safety, national security, and you.

WHAT COUNTS AS RISK MANAGEMENT

There are numerous topics in the broad category of *risk management* but it is often used in a much narrower sense than it should be. When the term is used too narrowly, it is either because *risk* is used too narrowly, *management* is used too narrowly, or both.

If you start looking for definitions of *risk*, you will find many wordings that add up to the same thing, and a few versions that are fundamentally different. For now, I'll skirt some of the deeper philosophical issues about what it means (yes, there are some, but that will come later) and I'll avoid some of the definitions that seem to be unique to specialized uses. Chapter 5 is devoted to why the definition I am going to propose is preferable to various mutually-exclusive alternatives that each have proponents who assume their's is the "one true" definition.

For now, I'll focus on a definition that, although it contradicts some definitions, best represents the one used by well-established, mathematical treatments of the term (e.g. actuarial science), as well as any English dictionary or even how the lay-public uses the term (see the box below).

DEFINITION OF RISK

Long definition: The probability and magnitude of a loss, disaster, or other undesirable event

Shorter (equivalent) definition: Something bad could happen

The second definition is more to the point, but the first definition gives us an indication of how to quantify a risk. First, we can state a probability that

the undesirable event will occur. Also, we need to measure the magnitude of the loss from this event in terms of financial losses, lives lost, and so on.

The undesirable event could be just about anything, including natural disasters, a major product recall, the default of a major debtor, hackers releasing sensitive customer data, political instability around a foreign office, workplace accidents resulting in injuries, or a pandemic flu virus disrupting supply chains. It could also mean personal misfortunes, such as a car accident on the way to work, loss of a job, a heart attack, and so on. Almost anything that could go wrong is a risk.

Since risk *management* generally applies to a management process in an organization, I'll focus a bit less on personal risks. Of course, my chance of having a heart attack is an important personal risk to assess and I certainly try to manage that risk. But when I'm talking about the failure of risk management—as the title of this book indicates—I'm not really focusing on whether individuals couldn't do a better job of managing personal risks like losing weight to avoid heart attacks (certainly, most should). I'm talking about major organizations that have adopted what is ostensibly some sort of formal risk management approach that they use to make critical business and public policy decisions.

Now, let us discuss the second half of the phrase *risk management*. Again, as with *risk*, I find multiple, wordy definitions for *management*, but here is one that seems to represent and combine many good sources:

DEFINITION OF MANAGEMENT

Long definition: The planning, organization, coordination, control, and direction of resources toward defined objective(s)

Shorter, folksier definition: Using what you have to get what you need

There are a couple of qualifications that, while they should be extremely obvious, are worth mentioning when we put *risk* and *management* together. Of course, when an executive wants to manage risks, he or she actually

10 CHAPTER 1 HEALTHY SKEPTICISM FOR RISK MANAGEMENT

wishes to reduce it or at least not unduly increase it in pursuit of better opportunities. And since the current amount of risk and its sources are not immediately apparent, an important part of reducing or minimizing risks is figuring out where the risks are. Also, risk management must accept that risk is inherent in business and risk reduction is practical only up to a point. Like any other management program, risk management has to make effective use of limited resources. Putting all of that together, here is a definition (again, not too different in spirit from the myriad definitions found in other sources):

**DEFINITION
OF RISK
MANAGEMENT**

Long definition: The identification, assessment, and prioritization of risks followed by coordinated and economical application of resources to minimize, monitor, and control the probability and/or impact of unfortunate events

Shorter definition: Being smart about taking chances

Risk management methods come in many forms, but the ultimate goal is to minimize risk in some area of the firm relative to the opportunities being sought, given resource constraints. Some of the names of these efforts have become terms of art in virtually all of business. A popular, and laudable, trend is to put the word *enterprise* in front of *risk management* to indicate that it is a comprehensive approach to risk for the firm. *Enterprise risk management (ERM)* is one of the headings under which many of the trends in risk management appear. I'll call ERM a type of risk management *program*, because this is often the banner under which risk management is known. I will also distinguish programs from actual methods since ERM could be implemented with entirely different methods, either soft or quantitative.

The following are just a few examples of various management programs to manage different kinds of risks (*Note:* Some of these can be components of others and the same program can contain a variety of different methods):

- Enterprise risk management (ERM)
- Portfolio management or project portfolio management (PPM)
- Disaster recovery and business continuity planning (DR/BCP)
- Project risk management (PRM)
- Governance risk and compliance (GRC)
- Emergency/crisis management processes

Risk management includes analysis and mitigation of risks related to physical security, product liability, information security, various forms of insurance, investment volatility, regulatory compliance, actions of competitors, workplace safety, getting vendors or customers to share risks, political risks in foreign governments, business recovery from natural catastrophes, or any other uncertainty that could result in a significant loss.

ANECDOTE: THE RISK OF OUTSOURCING DRUG MANUFACTURING

At a conference organized by the Consumer Health Products Association (a pharmaceutical industry association), I witnessed a chemical engineer describing a new risk management process he had developed for his firm. The risk analysis method was meant to assess an important and emerging risk in this field.

To control costs, this large pharmaceutical manufacturer was more frequently outsourcing certain batch processes to China. Virtually all of this manufacturer's competition was doing the same. But while the costs were significantly lower, they had a concern that batches from China might have additional quality control issues over and above those of batches manufactured here in the United States. These concerns were entirely justified.

The conference was in October 2007, and earlier that year there had already been several widely publicized product safety incidents with goods produced in China. In June, there was a toxin found in toothpaste and lead found in toys produced in China. Then there was tainted pet food that killed as many as 4,000 pets. There was even the disturbing case of "Aqua Dots," the children's craft-beads that stuck together to make different designs. The coating of these beads could metabolize in the stomach to produce gamma hydroxy butyrate—the chemical used in date-rape drugs.

Except for me, almost all of the audience were chemists, chemical engineers, and industrial engineers. They were previously listening to extremely technical sessions on sheer stress of particles in various processing equipment, yield curves, and mechanical details of drug packaging. There was no shortage of scientific thinkers and, from what I could tell, no timidity about mathematical models.

Yet, when the presenter was explaining the details of his company's new method for analyzing the risk of batches outsourced to China, I saw none of the hard science and skeptical peer-review that seemed common in the other sessions. He was describing a method based on a subjective "weighted score."³ In it, several "risk indicators" were each scored on a scale of 1 to 5. For example, if the manufacturer already produces a similar, but not identical, drug, it might get a low risk score of 2 on the indicator called "proven technical proficiency." If it was inspected by and got a positive evaluation from the Chinese health agency, but was not yet inspected by the Food and Drug Administration, then it might get a 4 on the "formal inspections" indicator. If the components of the drug required certain special safety controls that would be harder to outsource, then it might score as a higher risk in other areas. Each of these scores was based on the judgments of a team assembled to make these evaluations.

Then these scores were each multiplied by a weight of somewhere between 0.1 and 1.0 and then all of the weighted scores were totaled. The total of the weighted score might be 17.5 for one outsourcing strategy, 21.2 for another, and so on. The team that chose the scores also chose the weights and, again, it was based only on subjective judgments. The team further separated the resulting scores into various stratifications of risk that would, apparently, have some bearing on the decision to use a particular China-based source for a drug. For example, risk scores of over 20 might mean "Extremely high risk: Find an alternative"; 10 to 19 might mean "High risk: Proceed only with increased quality assurance," and so on.

When the engineer had finished describing the approach, I noticed that several heads in the room turned to me expecting some response. Earlier that day, I had given the keynote address describing, among other things, how risk can be quantified in a mathematically and scientifically meaningful way. Perhaps some were implementing something similar in their firms and were curious to see whether I would endorse it, but I suspect it was more likely they were expecting a criticism.

I neither endorsed nor rejected the approach outright. To be perfectly fair, neither position could yet be positively justified at that point without knowing a few more details (although there is a good chance it shared the flaws of many weighted scores, which I discuss later). I simply asked, “How do you know it works?” This is the most important question we could ask about a risk analysis and risk management approach. Once I knew the answer to that question, then I could legitimately take a position.

There was a long pause. It was obvious that they hadn’t even considered how to answer that question. So I thought it would be helpful (if a bit leading) to prompt them with another question: “Would you call this approach *scientific*?” After another pause, I asked, “Do you see how an actuary or statistician might not call this a *risk analysis*?” At this point, I sensed the questions were more like brow-beating than being helpful.

I then suggested to the presenter that the engineers in this field could be as scientific in their approach to this problem as they are in any other aspect of their profession. I pointed out that, for one, there was no need to start from scratch. If they were developing a new process for pharmaceutical manufacture, I’m sure they would examine existing research in the area. Likewise, there is quite a lot of literature in the general area of assessing risks in a mathematically and scientifically sound manner. It would be helpful to know that they don’t have to reinvent any of the fundamental concepts when it comes to measuring risks.

Then I pointed out that in the design of processes in drug production, once they had thoroughly reviewed the literature on a topic, no doubt they would design empirical tests of various components in the process, and measure them in a way that would satisfy the peer-reviewed journals and the FDA inspectors alike. Again, this same philosophy can apply to risk.

In fact, a much more sophisticated method is often already used to assess a different risk in the drug industry. “Stop-gate” analysis is used to determine whether a candidate for a new product should advance from formulation to animal testing, then from animal testing to human trials, until finally they decide whether to go to market. Many drug companies use proven statistical methods at each step in the stop-gate analysis. But, somehow, none of the basic concepts of stop-gate analysis were built upon to assess the risks of outsourcing production to China.

My questions to the presenter were rhetorical. I was already fairly sure that they had no objective measure for the effectiveness of this method. If

they had known to create such measures, they would probably have been inclined to create a very different approach in the first place. When it came to designing a method for assessing and managing risks, these scientists and engineers developed an approach with no more scientific rigor behind it than an ancient shaman reading goat entrails to determine where to hunt. While the lack of such rigor would be considered negligent in most of their work, it was acceptable to use a risk assessment method with no scientific backing at all.

In effect, they didn't think of this new risk in the same way as they thought of the substances and processes they use to manufacture drugs in a highly regulated industry. The chemicals they process and the vessels they use are concrete, tangible things and, to the engineers, risk might seem like an abstraction. Even the methods they use in stop-gate analysis might take on an air of concreteness simply because, by now, they have a lot of data on the problem. Perhaps, to them, the process of managing an unfamiliar risk seems like an intangible thing that doesn't lend itself to the same methods of validation that a drug manufacturing process would have to undergo for FDA approval. Applying the type of scientific reasoning and testing they use on the production of a drug to the risk analysis of producing that same drug in China is a leap they had not considered.

The presenter and the audience felt that the weighted scoring method they described was something close to "best practices" for the industry. When I asked, nobody in the room claimed to have an approach that was any more sophisticated. Most had no risk analysis at all on this problem.

Fortunately for the company that was presenting its risk management solution, it had not yet seen the worst-case scenarios that might result from unsound risk analysis. But with an entire industry approaching the outsourcing problem with either unscientific risk analysis methods or none at all, the worst case was inevitable. Just a few months after the conference, another major drug company using similarly subjective risk management methods on this problem would discover exactly how much was being risked by the outsourcing decisions (and the meager risk analysis applied to it).

Baxter International, Inc. was receiving reports of dangerous adverse reactions to its Chinese-manufactured blood-thinning drug called heparin. To its credit, by mid-January 2008, Baxter had voluntarily recalled some lots of the multidose vials of the drug. By then, the FDA was considering a

mandatory recall but had not yet done so because they believed other suppliers might not be able to meet demand for this critical drug. The FDA reasoned that this additional risk to patients requiring heparin therapy would be higher (I have no idea how much risk analysis went into *that* decision).

By February, the FDA had determined that the supply of heparin by other manufacturers was adequate and that Baxter should proceed with the recall of various types of heparin products. At the beginning of the recall in February, the FDA had linked four deaths to the Chinese-manufactured heparin and by March the number had grown to 19 deaths. By May 2008, the FDA had “clearly linked” a total of 81 deaths and 785 severe allergic reactions to the drug. Of course, chances are the various individual and class action lawsuits (just beginning as this book was written) will argue a much larger number.

The risks of outsourcing drug production to China always were high and the fact that some firms were at least attempting to develop a risk management method—regardless of its effectiveness—indicates that the industry was at least aware of the risk. The FDA is entrusted to inspect the operations of any drug manufacturer selling products in the United States, including foreign-based factories but, by March 2008, the FDA had inspected just 16 of the 566 Chinese drug manufacturers. The United States gets approximately 40% of its drugs from abroad. The scale of the problem easily justifies the very best risk analysis available.

Obviously, we can't be certain with only this information that the industry's lack of more sophisticated risk management for overseas drug manufacturing was the direct cause of the heparin incident. If the industry had used more sophisticated methods such as it already uses for stop-gate analysis, we could not be certain that some similar problem would not still have occurred. And, since the entire industry was unsophisticated in this area of risk management, there is certainly no reason to single out Baxter as particularly bad. This anecdote, by definition, is merely a single sample of the types of events that can occur and, by itself, is not sufficient to draw scientifically justified conclusions.

For any risk management method used in the pharmaceutical industry or any other industry, we must ask, again, “How do we know it works?” If we can't answer that question, then our most important risk management strategy should be to find a way to answer it and adopt a risk assessment and risk mitigation method that does work.

WHAT FAILURE MEANS

At the beginning of this chapter, we defined *risk* and *risk management*. Now we need to discuss what I mean by the *failure* of risk management. With some exceptions, it may not be very obvious. And that is part of the problem.

First, a couple of points about the anecdotes I just used. I believe United Airlines was probably applying what it believed to be a prudent level of risk management. I also believe the entire pharmaceutical industry and Baxter in particular were making a well-intentioned effort to manage the risks of outsourcing to China. When I refer to the “failure of risk management,” I do not just refer to outright negligence. Failing to employ the accounting controls that would have avoided Enron’s demise, for example, are not the kind of failures I examine the most in this book. I will concentrate more on the failure of sincere efforts to manage risks, as I will presume is the case with many organizations—even though we know the possible lawsuits must argue otherwise. I’m focusing on those organizations that believe they have adopted an effective risk management method and are unaware that they haven’t improved their situation one iota.

Second, I used these anecdotes in part to make a point about the limits of anecdotes when it comes to showing the failure or success of risk management. The single event of tainted blood thinner does not necessarily constitute a failure of risk management. Nor would a lucky streak of zero disasters have indicated that the risk management was working. At best, the pharmaceutical outsourcing anecdote shows one scenario of what could happen.

I think this is a departure from some approaches to the discussion of risk management. I have heard some entertaining speakers talk about various anecdotal misfortunes of companies as evidence that risk management failed. I have to admit, these stories are often fascinating, especially where the circumstances are engaging and the outcome was particularly disastrous. But I think the details of the mortgage crisis, 9/11, rogue traders, Hurricane Katrina, or Three Mile Island feed a kind of morbid curiosity more than they inform about risk management. Perhaps the stories made managers feel a little better about the fact they hadn’t (yet) made such a terrible blunder.

I will continue to use examples like this because that is part of what it takes to help people connect with the concepts. But we need a better

measure of the success or failure of risk management than single anecdotes. In most cases regarding risk management, an anecdote should be used only to *illustrate* a point, not to prove a point.

So, when I claim that risk management has failed, I'm not necessarily basing that on individual anecdotes of unfortunate things happening. It is possible, after all, that organizations where a disaster didn't occur were just lucky. They may have been doing nothing substantially different from organizations where disasters did occur. When I say that risk management has failed, it is for at least one of three reasons, all of which are independent of individual anecdotes: (1) the failure to measure and validate methods as a whole or in part; (2) the use of components that are known not to work; and (3) the lack of use of components that are known to work.

1. *Except for certain quantitative methods in certain industries, the effectiveness of risk management is almost never measured.* The biggest failure of risk management is that there is almost no experimentally verifiable evidence that the methods used improve on the assessment and mitigation of risks, especially for the softer (and much more popular) methods. If the only "evidence" is a subjective perception of success by the very managers who championed the method in the first place, then we have no reason to believe that the risk management method does not have a negative return. For a critical issue like risk management, we should require positive proof that it works—not just the lack of proof that it doesn't. Part of the success of any initiative is the measurable evidence of its success. It is a failure of risk management to know nothing of its own risks. It is also an avoidable risk that risk management, contrary to its purpose, fails to avoid.
2. *Some parts that have been measured don't work.* The experimental evidence that does exist for some aspects of risk management indicates the existence of some serious errors and biases. Since many risk management methods rely on human judgment, we should consider the research that shows how humans misperceive and systematically underestimate risks. If these problems are not identified and corrected, then they will invalidate any risk management method based even in part on human assessments. Other methods add error through arbitrary scales or the naïve use of historical data. Even

18 CHAPTER 1 HEALTHY SKEPTICISM FOR RISK MANAGEMENT

some of the most quantitatively rigorous methods fail to produce results that compare well with historical observations.

3. *Some parts that do work aren't used.* There are methods that are proven to work both in controlled laboratory settings and in the real world, but are not used in most risk management processes. These are methods that are entirely practical in the real world and, although they may be more elaborate, are easily justified for the magnitude of the decisions risk management will influence. Falling far short of what one could reasonably be expected to do is another form of failure.

In total, these failures add up to the fact that we still take unnecessary risks within risk management itself. Now it is time to measure risk management itself in a meaningful way so we can identify more precisely where risk management is broken and how to fix it.

SCOPE AND OBJECTIVES OF THIS BOOK

My objectives with this book are (1) to reach the widest possible audience among managers and analysts, (2) to give them enough information to quit using ineffective methods, and (3) to get them started on better solutions.

The first objective, reaching a wide audience, requires that I don't treat risk management myopically from the point of a given industry. There are many existing risk management texts that I consider important classics, but I see none that map the breadth of the different methods and the problems and advantages of each. There are financial risk assessment texts written specifically for financial analysts and economists. There are engineering and environmental risk texts for engineers and scientists. There are multiple risk management methods written for managers of software projects, computer security, or disaster recovery. Many of these sources seem to talk about risk management as if their methods comprised the entire subject. None seems entirely aware of the others.

The "wide audience" objective also means that I can't write just about the latest disaster. A reader picking up this book in 2009 may think the risk I'm talking about is a financial risk. If I had written this just after Katrina, risk might have meant something very different. But risk is not selective in that way and the best methods are not specific to one category of risks.

Thinking about risks means thinking about events that have not yet occurred, not just last year's news.

Finally, reaching a wide audience requires that I don't just write another esoteric text on quantitative methods for a small community of experts. Of those, there are already some excellent sources that I will not attempt to reproduce. A couple of slightly technical issues will be discussed, but only enough to introduce the important concepts.

The last two objectives, to get managers to quit using ineffectual methods and start them on a better path, are also satisfied by a "just technical enough" approach to the problem. This book won't make most managers masters of more quantitative and scientific methods of risk management. I merely want to convince them to make a radical change of direction from the methods they are most likely using now.

To accomplish these objectives, the remainder of this book is divided along the lines implied by the title:

- *Part One: An Introduction to the Crisis.* This first chapter introduced the problem and its seriousness. Chapter 2 outlines the diversity of approaches to assess and mitigate risks and discusses how managers rate their own firms in these areas. Chapter 3 examines how we should evaluate risk management methods.
- *Part Two: Why It's Broken.* After an introduction to four basic schools of thought about risk management, we will discuss the confusing differences in basic terminology among different areas of risk management. Then we will introduce several sources of fundamental errors in popular methods that remain unaddressed. We will list several fallacies that keep some from adopting better methods. Finally, this part of the text will outline some significant problems with even the most quantitative methods being used.
- *Part Three: How to Fix It.* This final part will introduce methods for addressing each of the previously discussed sources of error in risk management methods. We will talk about the basic concepts behind better methods, including how to think about probabilities and how to introduce scientific methods and measurements into risk management. Finally, we will talk about some of the issues involved in creating a culture in organizations and governments that would facilitate and incentivize better risk management.

20 CHAPTER 1 HEALTHY SKEPTICISM FOR RISK MANAGEMENT

Throughout this book, I will offer those who require more hands-on examples sample spreadsheets on this book's website at www.howtofixriskmgt.com. Those who prefer the "10,000-foot view" can still get a good idea of the issues without feeling dragged down by some technical details, whereas those who prefer to get more information can get specific example calculations. The website will also give all readers access to information on risks that evolve after this book has been published as well as a way to interact with other risk managers.

See this book's website at www.howtofixriskmgt.com for detailed examples from the book, discussion groups, and up-to-date news on risk management.

■ NOTES

1. My use of "placebo effect" requires a qualification. The placebo effect in medicine is the tendency among patients to experience both subjective and in some cases objectively observable improvements in health after receiving treatment that should be inert. This is a purely psychological effect but the improvements could be in objectively measurable ways—such as reducing blood pressure or cholesterol. However, when I refer to a placebo effect I mean that there literally is no improvement other than the subjective impression of an improvement.
2. Capt. A.C. Haynes "United 232: Coping With the 'One-in-a-Billion' Loss of All Flight Controls," *Accident Prevention* Volume 48, June 1991.
3. Some of the details of this are modified to protect the confidentiality of the firm that presented the method in this closed session, but the basic approach used was still a subjective weighted score.