Chapter 1

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

What is your single biggest risk? How do you know? These are critical questions for any organization regardless of industry, size, structure, environment, political pressures, or changes in technology. Any attempt to manage risk in these organizations should involve answering these questions.

We need to ask hard questions about new and rapidly growing trends in management methods, especially when those methods are meant to help direct and protect major investments and inform key public policy. The application of healthy skepticism to risk management methods was long past due when I wrote the first edition of this book more than a decade ago.

The first edition of this book came out on the tail end of the Great Recession in 2008 and 2009. Since then, several major events have

resulted in extraordinary losses both financially and in terms of human health and safety. Here are just a few:

- Deepwater Horizon offshore oil spill (2010)
- Fukushima Daiichi nuclear disaster (2011)
- Flint Michigan water system contamination (starting 2012)
- Samsung Galaxy Note 7 battery failures (2016)
- Multiple large data breaches (Equifax, Anthem, Target, etc.)
- Amtrak derailments/collisions (2018)

Events such as these and other natural, geopolitical, technological, and financial disasters in the beginning of the twenty-first century periodically accelerate (maybe only temporarily) interest in risk management among the public, businesses, and lawmakers. This continues to spur the development of several risk management methods.

The methods to determine risks vary greatly among organizations. Some of these methods—used to assess and mitigate risks of all sorts and sizes—are recent additions in the history of risk management and are growing in popularity. Some are well-established and highly regarded. Some take a very soft, qualitative approach and others are rigorously quantitative. If some of these are better, if some are fundamentally flawed, then we should want to know.

Actually, there is very convincing evidence about the effectiveness of different methods and this evidence is not just anecdotal. As we will see in this book, this evidence is based on detailed measurements in large controlled experiments. Some points about what works are even based on mathematical proofs. This will all be reviewed in much detail but, for now, I will skip ahead to the conclusion. Unfortunately, it is not good news.

I will make the case that most of the widely used methods are not based on any proven theories of risk analysis, and there is no real, scientific evidence that they result in a measurable improvement in decisions to manage risks. Where scientific data does exist, the data show that many of these methods fail to account for known sources of error in the analysis of risk or, worse yet, *add error of their own*.

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 "analysis placebo effect" (more to come on that).[1] 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-difficult differentiation between methods that work and methods that don't work means that ineffectual methods are likely to spread. Once certain methods are adopted, institutional inertia cements them in place with the assistance of standards and vendors that refer to them as "best practices." Sometimes they are even codified into law. Like a dangerous virus with a long incubation period, methods are passed from company to company with no early indicators of ill effects until it's too late.

The consequences of flawed but widely adopted methods are inevitably severe for organizations making critical decisions. Decisions regarding not only the financial security of a business but also the entire economy and even human lives are supported in large part by our assessment and management of risks. The reader may already start to see the answer to the first question at the beginning of this chapter, "What is your biggest risk?"

A "Common Mode Failure"

The year 2017 was remarkable for safety in commercial air travel. There was not a single fatality worldwide from an accident. Air travel had already been the safest form of travel for decades. Even so, luck had some part to play in the 2017 record, but that luck would not last. That same year, a new variation of the Boeing 737 MAX series passenger aircraft was introduced: the 737 MAX 8. Within twelve months of the initial roll out, well over one hundred MAX 8s were in service.

In 2018 and 2019, two crashes with the MAX 8, totaling 339 fatalities, showed that a particular category of failure was still very possible in air travel. Although the details of the two 737 crashes were still emerging as this book was written, it appears that it is an example of a

common mode failure. In other words, the two crashes may be linked to the same cause. This is a term familiar to systems risk analysis in some areas of engineering, where several failures can have a common cause. This would be like a weak link in a chain, but where the weak link was part of multiple chains.

I had an indirect connection to another common mode failure in air travel forty years before this book came out. 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 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 rest of my 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 the 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" event[2] and the media repeated this claim. But because mathematical misconceptions are much more common than one in a billion, if someone tells you that something that had just occurred had merely a one-in-a-billion chance of occurrence, you should consider the possibility that they calculated the odds incorrectly.

This event, as may be the case with the recent 737 MAX 8 crashes, was an example of a common mode failure because a single source

caused multiple failures. If the failures of three hydraulic systems were entirely independent of each other, then the failure of all three hydraulic systems in the DC-10 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. Likewise, a single software problem may cause problems on multiple 737 crashes.

Now consider that the cracks in the turbine blades of the DC-10 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. And human error in large complex software systems like those used on the 737 MAX 8 is almost inevitable and takes significant quality control to avoid. In a way, human error was an *even-more-common* common mode failure in the system.

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, whether it is hydraulics or software, would be very hard to discover and much more likely to materialize. In effect, *a flawed risk management is the ultimate common mode failure.*

And suppose they are flawed not just in one airline but in most organizations. The effects of disasters like Katrina, the financial crisis of 2008/2009, Deepwater Horizon, Fukashima, or even the 737 MAX 8 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 the risk management process itself. A weak risk management approach is effectively the biggest risk in the organization.

The financial crisis occurring while I wrote the first edition of this book was another example of a common mode failure that traces its way back to the failure of risk management of firms such as 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. 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 also are used to affect decisions that impact on human health and safety.

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 dollars would not have to cost more than 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.

A NASA scientist once told me the way that NASA reacts to risk events. If she were driving to work, veered off the road and ran into a tree, NASA management would develop a class to teach everyone how not to run into *that specific tree*. In a way, that's how most organizations deal with risk events. They may fix that immediate cause but not address whether the original risk analysis allowed that entire category of flaws to happen in the first place.

Key Definitions: *Risk Management* and Some Related Terms

There are numerous topics in the broad term of *risk management* but this term is often used in a much narrower sense than it should be. This is because *risk* is used too narrowly, *management* is used too narrowly, or both. And we also need to discuss a few other key terms that will come up a lot and how they fit together with risk management, especially the terms *risk assessment, risk analysis,* and *decision analysis*.

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 risk 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 6 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 definition is the "one true" definition.

For now, I'll focus on a definition that, although it contradicts some uses of the term, 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.

DEFINITION OF RISK

Long definition: A potential loss, disaster, or other undesirable event measured with probabilities assigned to losses of various magnitudes

Shorter (equivalent) definition: The possibility that something bad could happen

The second definition is more to the point, but the first definition describes a way to quantify a risk. First, we determine a probability that the undesirable event will occur. Then, we need to determine 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 surrounding 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.

Because 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. I'm referring to 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, although 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 wishes to reduce it or at least make sure it is acceptable in pursuit of better opportunities. And because 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. Similar to any other management program, risk management has to make effective use of limited resources.

Of course, we must accept that risk is inherent in business and risk reduction is practical only up to a point. 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, analysis, and prioritization of risks followed by coordinated and economical application of resources to reduce, 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, I think, 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 because ERM could be implemented with entirely different methods, either soft or quantitative.

The following are just a few examples of various programs related to managing 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)
- Project portfolio management (PPM) or Project risk management (PRM)

- Portfolio management (as in financial investments)
- Disaster recovery and business continuity planning (DR/BCP)
- Governance risk and compliance (GRC)
- Emergency/crisis management processes

The types of risks managed, just to name a few, include 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.

As the previous definition indicates, risk management activities include the analysis and mitigation of risks as well as establishing the tolerance for risk and managing the resources for doing all of this. All of these components of risk management are important but the reader will notice that this book will spend a lot of time on evaluating methods of *risk analysis*. So let me offer both a long and short definition of risk analysis at this point.

DEFINITION OF RISK ANALYSIS

Long definition: The detailed examination of the components of risk, including the evaluation of the probabilities of various events and their ultimate consequences, with the ultimate goal of informing risk management efforts

Shorter definition: How you figure out what your risks are (so you can do something about it)

Note that some risk managers will make a distinction between risk analysis and *risk assessment* or may use them synonymously. If they are used separately, it is often because the identification of risk is considered separate from the analysis of those risks and

together they comprise risk assessment. Personally, I find the analysis and identification of risks to be an iterative, back-and-forth process without a clear border between them. That is, we start with some identification of risk but on analyzing them, we identify more risks. So I may use the terms *analysis* and *assessment* a bit more interchangeably.

Now, obviously, if risk analysis methods were flawed, then the risk management would have to be misguided. If the initial analysis of risk is not based on meaningful measures, the risk mitigation methods are bound to address the wrong problems. If risk analysis 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. Just consider how flawed risk management may impact an organization or the public in the following situations.

- The approval and prioritization of investments and project portfolios in major US companies
- The level of protections needed for major security threats, including cybersecurity threats, for business and government
- The approval of government programs worth many billions of dollars
- The determination of when additional maintenance is required for old bridges or other infrastructure
- The evaluation of patient risks in health care
- The identification of supply chain risks due to pandemic viruses
- The decision to outsource pharmaceutical production overseas

Risks in any of these areas, and many more, could reveal themselves only after a major disaster in a business, government program, or even your personal life. Clearly, mismeasurement of these risks would lead to major problems—as has already happened in some cases.

The specific method used to assess these risks 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."

Because 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. Sometimes, management or vendors rely on surveys to assess the effectiveness of risk analysis, but they are almost always self-assessments by the surveyed organizations. They are not independent, objective measures of success in reducing risks.

I'm focusing on the analysis component of risk management because, as stated previously, risk management has to be informed in part by risk analysis. And then, how risks are mitigated is informed by the cost of those mitigations and the expected effect those mitigations will have on risks. In other words, even choosing mitigations involves another layer of risk analysis.

This, in no way, should be interpreted as a conflation of risk analysis with risk management. Yes, I will be addressing issues other than what is strictly the analysis of risk as the problem later in this book. But it should be clear that if this link is weak, then that's where the entire process fails. If risk analysis is broken, it is the first and most fundamental common mode failure of risk management.

And just as risk analysis is a subset of risk management, those are subsets of *decision analysis* in general decision-making. Risks are considered alongside opportunities when making decisions, and decision analysis is a quantitative treatment of that topic. Having risk management without being integrated into decision-making in general is like a store that sells only left-handed gloves.

What Failure Means

Now that we have defined risk management, 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 airlines and aircraft manufacturers involved in the crashes described

before were probably applying what they believed to be a prudent level of risk management. I also believe that many of the other organizations involved in other disasters I listed were not always just ignoring risk management practices. When I refer to the "failure of risk management," I do not just refer to outright negligence. Deliberately 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. No single event necessarily constitutes a failure of risk management. Nor would a lucky streak of zero disasters have indicated that the risk management was working.

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, major cyberattacks, or Fukushima 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 in which a

disaster hasn't occurred are just lucky and they may have been doing nothing substantially different from organizations in which disasters have occurred. 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 effectiveness of risk management itself is *almost* never measured:** The biggest failure of risk management is that there is usually 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 accept 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. Because 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 naive use of historical data. Even 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 they 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.

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 of 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 analysis 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 the first edition of this book in 2009 may think the risk I'm talking about is a financial risk. If I had written this just after the Fukushima Daiichi nuclear disaster of 2011 or more recent events, then 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. So, I will spend very little time on well-developed methods in actuarial science or quality control in engineering. The focus will be more on where there are numerous competing methods and the highest levels of management such as ERM.

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 in 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. Chapter 4 will show a simple "straw man" that can be the basis for developing a fully quantitative model. (This will also provide a way to imagine an alternative to current risk management methods as we go through a long and detailed criticism of them.)
- **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 book 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 build on the basic straw man model introduced in chapter 4. We will discuss 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.

Throughout this book, I will offer those who require more hands-on examples sample spreadsheets on this book's website at www.howtomeasureanything.com/riskmanagement. 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 evolving risks, new ideas and a community of other professionals interested in commenting on those.

See this book's website at www.howtomeasureanything.com/riskmanagement 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* 48, June 1991.