



RETHINKING RISK IN NATIONAL SECURITY

**Lessons of the Financial Crisis
for Risk Management**

Michael J. Mazarr



Rethinking Risk in National Security

Rethinking Risk in National Security

**Lessons of the Financial Crisis for Risk
Management**

Michael J. Mazarr

palgrave
macmillan

Michael J. Mazarr
RAND Corporation
Arlington, VA, USA

ISBN 978-1-349-94887-1 ISBN 978-1-349-91843-0 (eBook)
DOI 10.1007/978-1-349-91843-0

© The Editor(s) (if applicable) and The Author(s) 2016

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Palgrave Macmillan imprint is published by Springer Nature
The registered company is Nature America Inc. New York

Contents

<i>List of Tables and Boxes</i>	vii
<i>Acknowledgments</i>	ix
Part I Background	
1 Risk, Judgment, and Uncertainty	3
2 Defining Risk	19
3 Approaches to Risk in National Security	35
Part II Lessons of the Crisis—The Character of Risk	
4 Risk and Uncertainty	51
5 Risk Is What We Make of It	65
6 Indifferent to Consequences	77
7 The Swans to Worry About Are Gray	89
8 Risk Becomes Personalized	101
9 What You Don't Know Can Destroy You: Ignorance and Correlated Risk	113
10 Risk, Incentives, and Culture	125
Part III Toward Improved Risk Practices	
11 The Role of Risk in Strategy	143
12 Outcome Assessment of the Emerging US National Security Strategy	157
13 Principles of Effective Risk Management	175
14 Managing Uncertainty	189
<i>Notes</i>	207
<i>Bibliography</i>	233
<i>Index</i>	241

List of Tables and Boxes

Tables

14.1	Managing risk vs. managing uncertainty	192
------	--	-----

Boxes

11.1	Categories of outcome risks	154
13.1	Revised approach to categorizing risk	186
14.1	Principles for managing uncertainty	195

Acknowledgments

This study emerged from research undertaken several years ago on the potential for “strategic insolvency” inherent in the US national security posture. I would like to thank Alex Lennon, the esteemed editor of *The Washington Quarterly*, for his support in getting that analysis to a wider audience. Without him, this book would not have been possible.

I owe most direct thanks for this work to Dr. Marin Strmecki of the Smith Richardson Foundation. Marin supervises one of the most rigorous and influential international relations grant-making programs in the foundation world. He saw the potential for a study on risk, encouraged me to pursue the issues in this volume, and shepherded the process by which this proposal was selected for a grant. His support has been deeply appreciated on a number of occasions over the last several years. Also at Smith Richardson, Dr. Nadia Schadlow offered critical help in executing the grant.

I would like to offer sincere thanks to Peter Bergen at the New America Foundation. Peter is renowned for his work as a journalist and scholar of South Asia, terrorism, and related issues; he is also a generous collaborator and supportive colleague. He offered his program at New America Foundation as a virtual home for this project and was supportive throughout.

The staff at New America was incredibly helpful in everything from finances to grant administration to support for meetings.

Several individuals gave generously their time and expertise in the research for this book. I owe a special debt to my good friend Stuart Rabin, who shared his own expertise and helped to arrange a number of meetings for the project. Paul Slovic has been wonderfully encouraging throughout the process. Joel Mentor, Chris Severson, Troy Thomas, and Bradley Ziff all offered assistance during the research. Several esteemed scholars consented to extensive informal dialogues, either in person or on the phone, with the promise that they would not be named—I thank each of them. And a number of policy experts and public servants attended a roundtable in Washington, D.C. that helped to refine the concepts. Needless to say none of these individuals are responsible for any of the analysis or claims in the book.

More than anyone, I owe a sincere thanks to my wife Jennifer and sons Alex and Theo for putting up with an often distracted father, especially as I worked to finish the final draft of the book.

This work reflects only my personal views.

Part I

Background

1

Risk, Judgment, and Uncertainty

In October 2008, under the fearsome shadow of the most serious economic crisis since the 1930s, a man who had unwittingly done much to bring it about—former Federal Reserve Chairman Alan Greenspan—testified before Congress. Facing a barrage of heated questions, Greenspan made a remarkable confession. He admitted that his worldview had been wrong.

Greenspan had discovered a miscalculation in his ideology, he confessed—a “flaw in the model that I perceived [to be] the critical functioning structure that defines how the world works.” That flaw was the assumption that markets and firms could be rationally self-policing, in part through the effective control of risk. “In recent decades,” Greenspan testified, “a vast risk management and pricing system has evolved combining the best insights of mathematicians and finance experts supported by major advances in computer and communications technology.” This “modern risk management paradigm held sway for decades,” he explained. But it would have to be rethought. “The whole intellectual edifice ... collapsed in the summer of last year.”¹

Greenspan’s point appeared self-evident by the time he testified, but it would hardly have seemed that way two years before. The financial crisis had laid bare profound underlying dangers in the ways in which major financial institutions dealt with risk. New York Fed chief Timothy Geithner, speaking just as the crisis broke, explained that the crisis had “exposed a range of weaknesses in risk management practices within financial institutions.”² It was “obvious,” one *Financial Times* columnist argued, “that there has been a massive failure of risk management across most of Wall Street.”³ A scholarly paper assessing the causes of the crisis later referred to the “nearly unanimous view amongst the regulators that lapses in risk management played a critical role in exacerbating the crisis.”⁴

Yet at the same time that flaws in procedural risk management were being exposed in the financial sector, the same practices were becoming commonplace in areas well beyond finance—most notably, for the purposes of this

analysis, in national security. This essential paradox is the inspiration for this study: The national security enterprise is relying in increasingly important ways on a tool whose limitations and perils have become increasingly evident.

My argument is not that risk management itself is bankrupt, even in its more quantitative approaches. I am sympathetic to the arguments of Nassim Nicholas Taleb, for example, about the limits of modeling under uncertainty, but I also appreciate the proven value of quantitative models, for assessing risk as well as other purposes, even in a context as protean as the market. More broadly, taking risk into account is an essential component of strategy. Many firms have employed risk management techniques to great advantage.

This study, in other words, is not intended as a frontal attack on risk management. Instead, by deriving common patterns from the experiences of a number of firms and agencies in the financial crisis, it examines ways in which risk efforts can be misused and abused. In particular, it is the story of how even extensive risk procedures can be brought low by human factors such as overconfidence, herding, groupthink, institutional culture, and malign incentives.

The core argument of the study is not that risk management is useless. Instead the study makes three more discrete arguments designed to enhance its application in national security—but lessons which might be of equal interest to decision-makers in business and even intelligence, whose warning function shares much in common with risk management.

The first conclusion is that, in order to do its job effectively, a risk process must have a clearly defined purpose in strategy. When the concept of risk becomes fragmented to the point of obscurity, it cannot contribute in meaningful ways to effective strategic choice.

Second, the role of risk management must match the kinds of decisions being made. Too often before the financial crisis (and even today), quantitative risk models were used to generate supposedly reliable, objective forecasts of situations that reflected deep uncertainty. Models can be accurate and entirely appropriate to assess certain issues—short-term anomalies in specific markets, for example. When used as a substitute for strategic judgment under uncertainty, however, risk management invites disaster.

I am, in particular, interested in the highest-level decisions that enterprises can make: big bets on which decision-makers will always have too little information, which involve intensely nonlinear dynamics and contested values, and much else. I will term such choices “complex strategic judgments.” This term reflects a critical distinction at the heart of this study: It is not most directly relevant to highly specific risk assessments of incredibly particular, and sometimes reasonably deterministic, issues—the risk assessment of the fuel system of the National Aeronautics and Space Administration’s (NASA) newest rocket, for example. I am interested in how thinking about risk supports transformative strategic decisions. Evidence from the financial crisis

points to the potential value of such a focus: It is precisely because risk management has become so complex, professionalized, quantified, model-based, and arguably disjointed that many risk processes have become disconnected from the most important choices made by senior leaders.

Third and finally, procedural risk management—models and processes designed to offer warning of accruing risk—is no match for human factors. The crisis makes abundantly clear that cognitive and social factors ranging from simple overconfidence to the personalization of risk to risk-obsessed corporate cultures consistently overrode the findings of risk processes. Risk management, I conclude, is not a challenge of process—it is a challenge of leadership, analytical rigor, and institutional culture.

What the financial crisis uncovered, as much as anything else, is that organizations do not so much face a challenge of designing ideal risk management procedures. Much more fundamentally, their health and success depends on something much broader: creating a culture that integrates consequence management into strategy, in part (I will argue) by adopting principles for managing uncertainty. The following chapters lay out these arguments and apply them to a field that has lately become widely committed to the use of risk to inform strategic judgment: national security.

The rise of “risk” in national security

The current national security context is crowded with references to risk. Many defense documents now include sections on the issue—especially the Quadrennial Defense Reviews, whose risk sections were specifically mandated in law. The Department of Homeland Security prepares a Strategic National Risk Assessment and tutors its leaders on risk management fundamentals. Program risk is a common feature of procedures at NASA, the Department of Energy, and many other agencies. Senior officials routinely make reference to risk in testimony, speeches, and public statements. The term “risk” crops up constantly in discussion of current issues and defense policy: The United States is “taking risk” with a certain decision; Russian or Chinese actions pose “risks”; the US defense posture reflects significant “risk” relative to the defense budget and capabilities of the force; additional investments would help to “buy back risk.”

Considerations of risk are infused in all manner of public and classified planning documents, and senior military and civilian leaders increasingly refer to the importance of dealing with risk in defense planning. There are literally dozens of different risk management processes and frameworks in place in the national security enterprise, from intensely specific and discrete program-specific efforts to programs that attempt to measure risk across the whole defense enterprise.⁵ Beyond the United States, moreover, a number of countries have consciously integrated risk management into their defense planning processes.⁶

At the same time, national security leaders are increasingly referring to “uncertainty” to describe the context for defense planning. Senior Army leaders have described uncertain and unpredictable futures as the “biggest threat” to their service: Without knowing what wars to anticipate, they could get many fundamental choices wrong.⁷ Former Chairman of the Joint Chiefs General Martin Dempsey repeatedly claimed that the strategic environment was “as uncertain as I have seen in 40 years of service.”⁸

This growing use of these concepts is largely a function of the dominant reality of US national security strategy: At a time of fiscal austerity and a full plate of pressing security challenges, the managers of the US national security enterprise are facing increasing difficulty reconciling ends and means, even as the international context seems to be growing more unstable than at any time in the last two decades. At such a time of volatility when the United States—the acknowledged engineer of the global system and the source of its most important security guarantees—is becoming less willing and able to play its traditional role, national security strategists are looking to concepts of risk to help them manage a seemingly diverging gap between ends and means. Increasingly, this gap is being conceived as risk.

One of the most bracing recent statements on the issue came from the official review panel for the 2014 Quadrennial Defense Review (QDR), which focused on declining capabilities as a source of risk. They concluded that “the trend line is clear: The delta between threats and capabilities is rapidly growing. Given the uncertain global threat environment, the erosion of certain American advantages, and projected budget levels, we are prepared to say that unless recommendations of the kind we make in this report are adopted, the armed services will in the near future be at high risk of not being able to fully execute the national defense strategy.”⁹

This emerging challenge to US national security strategy is often presented as a fundamental problem of risk. Leading defense planning documents increasingly boast sections on risk, and frameworks for its evaluation. But this fact merely brings us back to the essential paradox that is energizing this study: The national security enterprise is making increasing use of a concept and approach that proved incapable of evading disaster in the financial sector.

This seeming irony—the fact that the US national security community may be placing growing faith in a potentially unreliable tool—provides the basic motivation for this study. The central research question is whether the experience with risk management in the 2007–2008 financial crisis holds specific lessons for the use of risk to inform national security strategy decisions. The resulting analysis is designed to be useful to national security professionals, but it should also be of interest to senior decision-makers in business or other fields who regularly confront the concept of risk.

Part of the problem is that the term “risk” has come to mean too many things, and to be used for too many purposes. In its essence, risk involves

something that can go wrong in relation to a value or objective of an organization.¹⁰ It is part of a constant and dynamic series of balances—between risk and opportunity, risk and reward—that must be struck in the process of managing complex enterprises. Only by assessing, comparing, and managing risk can an enterprise effectively address its goals and interests with a full picture of the possible consequences of its actions.

And yet in service of these reasonable goals, the concept of risk has been stretched to the breaking point. It now encompasses everything from dangers in the strategic environment to gaps between means and ends to the role of domestic politics. It has come, in some cases, to substitute for strategy altogether. “If we were to read 10 different articles or books about risk,” two writers have concluded, “we should not be surprised to see risk described in 10 different ways.”¹¹ A concept that means several different things to different people, whose essence changes in the eye of the beholder, can end up meaning nothing at all. The literature on risk, the scholar John Adams explains, has become “vast, sprawling and ill-disciplined.”¹²

More pointedly, these cases suggest that there is a critical gap between procedural approaches to risk and the real underlying causes of the crisis. It turned out that the most elaborate and complex procedures—even, perhaps especially, those grounded in quantitative approaches using data sets and algorithms—could not stand in the way of skewed incentives, cognitive biases, groupthink, and a dozen other human factors that led companies to take excessive risk.

The experience of the financial crisis should therefore invite us to rethink what we mean by risk. In a seminal essay, the scholar Jack Dowie even argued that risk had become “an obstacle to improved decision and policy making.” The “multiple and ambiguous usages” of the term, Dowie argued, “persistently jeopardize the separation of the tasks of identifying and evaluating relevant evidence on the one hand, and eliciting and processing necessary value judgments on the other.” The idea of risk “is simply not needed” to make strategic judgments, he contends, recommending that we eliminate it altogether and replace its various functions with more classic terms and stages of strategy.¹³

I have sympathy for Dowie’s perspective. The concept of risk has often been more misleading than helpful, and all (or nearly all) the issues to which it refers can be more profitably handled by different elements of a strategy process. The gap between means and ends, for example, is a problem of sufficiency or feasibility, and should be addressed in a basic analysis of the degree of resources available.

Yet the terminology of risk has become firmly embedded in corporate and national security practice. Organized properly, moreover, to support complex strategic judgments, a risk process can help force an institution to take seriously an element of strategy that many would rather avoid: the consequences of strategic choices. (Used in more pointed and discrete ways, risk management

in a more objective and quantitative sense can also inform individual choices.) Rather than simply forsaking the term, then, this study will propose a revised approach and framework that aim to clarify and narrow its scope. It will suggest that the most important question to ask, when conceiving of complex strategic judgments, is *what sort of conversations an organization is trying to generate with its risk process*. It will argue for the use of the term to focus on outcomes, and wrap that emphasis inside a larger and more encompassing process that I will term “managing uncertainty for competitive advantage.”

A word on methods

In order to evaluate these issues, I chose a methodology of qualitative case studies. My goal was to understand how institutions attempted to use risk, how risk procedures interacted with financial calamities, and why elaborate risk procedures failed. The best source of information for such issues was the stories of decision-making groups attempting to manage risk, either in finance or national security. This analysis, therefore, reflects both an effort to engage the literature on risk management and a study of the experience of specific firms in the recent crisis.

In particular, this study relies on an assessment of the accumulating literature on planning and decision-making processes in key financial institutions before the crisis—firms such as Merrill Lynch, Bear Sterns, American International Group (AIG), and Goldman Sachs, among others. It offers a comparative discussion of two earlier crises in risk management, at the hedge fund Long-Term Capital Management and the infamous energy trader Enron. For the most part I have relied on the extensive secondary literature on such cases, though I also conducted a number of dialogues with experts in the risk industry.

My basic method was to accumulate a set of hypotheses of what might have gone wrong with risk and then test them against evidence from the cases, looking for the issues where all or nearly all reflected the same issues. Those lessons are presented in Chapters 4 through 10, and they focus largely on the role of human factors in obstructing effective procedural risk management. Those lessons derive from specific events, behavior, and analysis; a more inferential finding is the importance of a specific role for risk in making strategic decisions, a case I make here and in Chapter 11.

This approach comes with all the potential limitations of qualitative case studies. The findings could be idiosyncratic; one must be careful about generalizing too readily from a small handful of cases. It can be difficult to obtain reliable information about what actually went on in specific institutions. The findings of any series of case studies are bound to be suggestive rather than determinative.

Nonetheless I found this methodology worthwhile for a number of reasons. The differences from one case to another make generalization difficult—but

they are also precisely what motivates such an approach, because applying quantitative methods to large data sets from fundamentally incomparable cases would have little value. In a context of institutions that differ significantly from one another, a researcher can have no confidence in the ability to build a truly representative set. Broad themes can be identified and tested among the cases to look for common trends.

The result of this analysis, to be sure, is suggestive and qualitative. This study does not reflect detailed results from a data set. However much I searched for common patterns in a series of case studies, the results necessarily reflect my own interpretation of the evidence. Others will draw different lessons about risk management from the financial crisis (and many have, including some I consulted as part of this research).

In the course of the research I have sought to deal with three methodological challenges. The first is that the findings might be overdetermined, or trivial: “Human factors” will always have influence on institutional behavior, for example. And yet the assumption of institutional risk management is just the opposite—that effective risk procedures can correct for perceptual bias and group dynamics. My hypothesis goes beyond the mere presence of human factors to suggest that they make procedural risk management, as commonly practiced today, bound to fail.

A second methodological challenge stems from the nature of the context. There may simply be too many factors at work to isolate the unique effect of any one, or small number, of them. “Human factors,” for example, remains an inevitably ambiguous concept. The specific character and role of such factors as wishful thinking and herding could vary significantly from case to case, and play different roles in a buzzing crowd of variables affecting behavior. This is true of any complex case study, however, and the goal here is not to precisely isolate some quantifiable effect of any given variable, but rather to find consistent patterns and relationships that can help guide our thinking about the nature of risk management.

Third and finally, with any case study research there is a risk of overgeneralizing from unrepresentative cases. If we were to examine the three or five cases out of a hundred in which a given causal relationship emerged, the findings would be highly misleading. I have tried to deal with this potential risk in a number of ways: by surveying a wide range of companies and national security cases; by examining the wider risk literature for themes that emerge from these cases; and through a series of discussions—notably with a number of senior national security officials in the fall of 2014, at an October 2014 roundtable in Washington, DC, and a January 2015 series of interviews in the financial sector—to test the general applicability of the lessons.

In sum, these findings are designed as a spur to continued dialogue and reflection on the role of risk in strategy. There are few if any conclusive results here. But the patterns that emerge appear to be consistent and significant

enough—pointing, as they do, to significant institutional dangers in the management of risk—that they ought to be of interest to senior officials in national security and corporate contexts.

Risk, uncertainty, and judgment

As suggested above, this study examines the role of risk in a very particular class of decisions. Its focus is reflected in two key issues, or distinctions. I am interested in large-scale strategic choices that I will term complex strategic judgments. As a result, this is a study of risk management in *non-deterministic environments*, and it is—in a closely related sense—a study of risk judgment under *deep and comprehensive uncertainty*. This is an important distinction and serves to limit the reach and applicability of my findings, because not all risk analysis takes place under such conditions.

A deterministic model, or environment, is one in which the outputs are determined by the inputs, and the inputs are known—and is therefore predictable. That idea presumes both (1) a strong basis of information about a situation, and (2) the fact that causal variables are well-understood. This is different from what is commonly known as a “stochastic” environment, in which one set of inputs can produce a wide range of very different outcomes. Many mechanical devices are, in effect, deterministic systems: Put in 20 percent more power, get 20 percent more force (or speed, or whatever outcome you are looking for). A management context, on the other hand, is stochastic: The same input to different employees, or the same employee at different times, can produce wildly different results. This doesn’t mean that linear models are irrelevant to stochastic environments, or that intentional strategy is pointless in such cases—but it does mean that any thinking about how causes and variables will unfold must be done with intense care.

What I have in mind with complex strategic judgments are the high-level strategic choices which senior leaders get paid to make: whether the United States should withdraw troops from Korea, or change the composition of its Army, or invade Iraq; whether a technology firm should abandon a traditional focus on hardware and become a services company. These are issues on which there are simply too many variables, interrelationships, unknown factors, and unpredictably emergent behavior to allow an optimal solution. The result, as I will argue, is a form of deep or radical *uncertainty* that characterizes most or all truly strategic decisions facing senior leaders.

This is not to suggest that data and analysis can play no role in *informing* such judgments. They can, and indeed I will argue that they must, as one component of an effort to tame the human factors that can push strategic choices into randomly intuitive directions. Much deeper analysis of Iraq’s infrastructure, for example, before March 2003 would have made much more clear the scale of the national reconstruction that would be required after the

US intervention, and provided better perspectives on the nature of the challenge Washington was about to bite off. In the corporate world, even big strategic bets will come with helpful baskets of data: the size of potential markets, the cost of specific options, the scale of debt required for key investments.

One key distinguishing characteristic of complex strategic judgments, though, is that the best data-based analysis will never be able to *make* the choice, in the sense of providing an objective, reliable answer. There will never be enough information to be sure that the analysis has captured the necessary factors. Causalities are too fickle. Nonlinear dynamics abound. “Transmutability” means that the effect of various ongoing choices is so great that the world that will determine the effect of choices can be significantly different from the context that existed when they were made. Choices will often be determined by subjective values and considerations not subject to modeling, from politics to personalities to ethical considerations. When making big strategic choices under such conditions, the final choice is ultimately, and unavoidably, a subjective and interpretive judgment.¹⁴

A leading question for this study is how risk considerations can best contribute to such complex judgments. One of the clear, and by now widely appreciated, lessons of the crisis is that trying to force deterministic solutions onto uncertain environments is a recipe for disaster. Institutions are anxious to cope with uncertainty with formalized, often data- and algorithm-driven procedures.¹⁵ These can be perfectly useful when applied effectively. Yet when used indiscriminately, or in the wrong contexts, or when used as a substitute for strategic judgment, they can become dangerous distractions,¹⁶ because the complex problems decision-makers must tackle are often immune to solution by such techniques. All too often, the results of risk management efforts are presented as if they *were* referring to a deterministic environment: highly quantified estimates, offered in detailed stoplight charts.

For the purposes of this analysis, then, I am referring to decisions of a very particular type.¹⁷ Many strategic choices that involve risk fall on a broad spectrum ranging from more linear and predictable to more uncertain. Taken together, a non-deterministic, uncertain environment produces the need for complex strategic judgments. Such choices have a number of particular characteristics.

- Outcomes—of current trends as well as new actions or behavior—cannot be forecast from present patterns and remain highly ambiguous.
- They are necessarily based on incomplete information.
- They involve issues, problems, or actions that are inherently subjective: Their meaning varies depending on the perception of the actors involved; there is no objective value function to be assigned.
- They involve issues that are complex in the formal sense, meaning that dozens or hundreds of variables are interacting to generate emergent patterns whose outcome cannot be accurately inferred from present arrangements.

- They involve contested values.
- As a result of these factors, there is no optimization process available for complex strategic judgments. At the time they are made and even in retrospect, there will never be an objectively discoverable “right” answer.

Three terms are especially important from the preceding discussion. One is *judgment*. This is a study of the use of risk to inform critical issues of state—but ones that must ultimately be resolved by subjective inference and conjecture about the likely future course of events and the potential effect of alternative courses of action. Such issues are fundamentally different from more discrete institutional choices—the optimal helicopter to replace an aging one in service today, the schedule of insurance benefits most likely to produce a given revenue stream from an actuarial point of view—that can be partially or completely resolved through objective calculations. It is the difference between challenges that have an identifiable “best answer” and those on which there will be unresolvable debates over facts, interpretation, and values.

A second term that will be important to this analysis is *outcome*. It is an essential aspect of such judgments that outcomes remain erratic and ambiguous: No matter how much data we gather, in the end we can only guess at what the results might be. If the United States were to deploy major land forces to the Baltics today to “deter” Russian aggression or adventurism, for example, the outcomes could fall across a wide spectrum, from acquiescence by Moscow to paranoid overreaction and military clash. And there would be no way to be sure, in advance, which would emerge.

The issue of outcomes is in turn related to a third concept that will recur throughout this study—*causality*. A major reason why outcomes are so ambiguous is because the causalities at work in a complex, uncertain environment cannot be known. In fact they evolve over time, so that a cause-and-effect relationship in effect at one moment may disappear in the future. A key aspect of complex strategic judgments is that causalities at work in the environment can only be inferred, and never very reliably.

Most of the classic debates in international relations and security studies are, in one way or another, about causalities. Will a given structure of the system produce certain behavior? Will retrenchment generate aggression? What makes the debates so frustrating—and ultimately unresolvable—is that repeatable causalities simply do not emerge in complex systems governed by human perception and the influence of group dynamics. Causal links are utterly contingent: A threat may deter one adversary and provoke another—and those relationships might be reversed in five years’ time. A major reason for this, of course, is that causalities in interactive strategy are governed by perceptions, and the meaning that decision-makers bring to a situation is idiosyncratic and difficult to predict.¹⁸

A fundamental problem in risk management for complex strategic judgment, then, is that outcomes—the foundation of risk—can only be guessed at,

in large part because the underlying causal links are obscure, unreliable, and constantly changing, like the tumbling shapes of a kaleidoscope. Even seemingly decisive pieces of information or intelligence will not always resolve this problem. A signals intercept in which the Russian president was heard forecasting his own likely reaction to the deployment of US forces would not prove with certainty that he would react that way in practice. The United States all but assured Moscow it did not consider Korea a vital interest in 1950, for example, only to turn around and fight a costly three-year war on precisely that basis once the North attacked.

It is now reasonably well established that the financial crisis proves that decision-makers did not fully appreciate these critical limitations to risk management—*aspects of non-determinism, uncertainty, and the fickleness of causality*. They saw issues as technical and technocratic rather than subjective and complex. Many viewed outcomes as substantially predictable rather than highly contingent, and treated the decisions they were making as optimizable choices rather than subjective and complex judgments. Partly as a result, they built up far more confidence in their plans and strategies than was warranted.

To put it simply: Organizations took approaches and models entirely appropriate for very discriminate use and employed them to justify big bets under uncertainty—without an intervening layer of rigorous analysis and careful, informed, self-aware, and self-critical judgment. Those qualities—*rigor, self-criticism, openness to information and alternative perspectives*—in turn represent the antidote to the frivolous treatment of risk. But the avenues to a flippant, overly deterministic use of risk processes did not arise in a vacuum. The context of the financial sector generated powerful incentives—and the culture of specific firms undermined rigorous decision-making—in ways that magnified the perceptual mistakes.

The danger is much the same in national security. Former Navy secretary Richard Danzig has written eloquently of the impulse toward predictive, linear analysis in defense circles. Bureaucracies, he explains, “seek predictability as a means of maintaining order.” Organizations have an institutional tendency to tame complex events with simplified planning procedures and predictive models. He quotes Henry Kissinger to the effect that bureaucracy generates a “quest for calculability.” The modern defense establishment, Danzig explains, is built on a foundation of predictive planning, enshrined by McNamara-era planning policies.¹⁹ This context creates a forceful temptation to domesticate nonlinear, uncertain environments with objective planning processes that generate seemingly objective assessments.

The troubles with risk management

In the late 1990s, one of the most-admired companies in the United States established a sophisticated risk management unit that soon garnered notice as

a best practice for the industry. It was called the Risk Assessment and Control unit, or RAC. At its peak the RAC boasted over 150 skilled analysts—finance experts, accountants, statisticians—and a \$30 million budget. In an apparent reflection of the priority it accorded risk management, the company required formal RAC approval of any significant deal. “Only two things at [this company] are not subject to negotiation,” the CEO once boasted in an interview: “The firm’s personnel evaluation policy and its company-wide risk management program.”²⁰ Outsiders were duly impressed: The rating service Standard and Poor’s declared their faith in the system. “Even though they’re taking more risk,” an S&P analyst said at the time, “their market presence and risk-management skills allow them to get away with it.”

As it turned out, things weren’t so rosy. The company was Enron, and its risk processes were, put it charitably, a sham.

Later, after Enron’s collapse, everything would seem so obvious. Enron executives admitted that RAC analyses were routinely ignored. “I treated them like dogs, and they couldn’t do anything about me,” one former executive told Bethany McLean and Peter Elkind. The man in charge of the RAC was well-liked but reportedly hesitated to confront senior leaders determined to make deals and take risk—and sometimes overruled subordinates who impeded favored projects. CEO Jeffrey Skilling reportedly said that this was exactly the way he wanted it; he bragged of having the foresight to choose someone so compliant for the risk management post. The bottom line was simple: As the anonymous executive told McLean and Elkind, “The process was there, sure, but the support wasn’t.”

The central argument of this study is that process itself means very little. A large number of human factors, from wishful thinking to groupthink to skewed incentives to imperative-driven thinking to risk-embracing cultures that punish dissent, can—in *any* operationally oriented, can-do culture—conspire to undermine effective thinking about risk. Risk analysis in support of complex strategic judgments is (or should be) all about consequences, what could go wrong from an organization’s choices. But a range of human factors tend to dim the image of the future and impede an unbiased consideration of outcomes. The Enron case represents perhaps the apotheosis of this phenomenon, a situation in which the future hardly mattered.

Risk management, the financial crisis strongly suggests, is about creating *a culture of rigorous analysis and habits of risk-aware judgment* in organizations. (This is one of the conclusions of the study that seems to apply equally well to discrete and big issues, deterministic and uncertain contexts.) But this turns out to be painfully difficult. It is only a slight exaggeration to conclude that risk management processes in a context of true uncertainty are destined to fail, if we judge the activity in terms of its ability to prevent risk disasters—tragedies that unfold because risks were not sufficiently taken into account.

These limitations to risk management call for a more explicit discussion of its purpose. Risk analyses crop up at various points in the development of a strategy or policy, sometimes without any coherent relationship. Some analyses of risk almost equate it to strategy. A critical goal, in this regard, will be to understand what we mean by “effective” or “successful” risk management, as opposed to the basic elements of a strategic planning process, and to find a role and purpose for the activity that is precise, targeted, and shared across an institution. In part the challenge is to distinguish a “failure” of risk management from an entirely reasonable judgment call under uncertainty, given that we know many such judgments will end up being wrong.²¹

This study contends that successful risk management for complex strategic judgments involves taking seriously the potential consequences of proposed strategies, assessing those dangers honestly and with eyes wide open, and then developing powerful and rigorous mitigation strategies once a strategy is put into effect. To be clear at the outset, then, when I refer to risk, I will ultimately be thinking about potential dangers inherent in the outcomes or consequences of proposed courses of action. It is through this approach that risk can make the most important contributions to strategy—a conclusion that emerges partly from the experience of risk management in the 2007–2008 financial crisis.

Risk in the financial crisis

The concept of risk management had become well established in the US financial sector by the mid-2000s. It was, in fact, a deeply entrenched, highly institutionalized management specialty. In pre-crisis polls conducted by Deloitte Consulting and others, the vast majority of firms reported having a Chief Risk Officer. They claimed to have Enterprise Risk Management processes. Most, by 2006, proclaimed themselves either very or extremely confident in the risk management procedures in their firm. Ben Bernanke, newly installed as the heir to Federal Reserve chairman Alan Greenspan, sang the praises of risk management in June of that year. Retail lending had become “routinized,” he proclaimed, because “banks have become increasingly adept at predicting default risk by applying statistical models to data, such as credit scores. ... [Banks have] made substantial strides over the past two decades in their ability to measure and manage risks.”²²

In the years before the crisis, risk management had also become a highly quantified and probabilistic discipline. The goal of such mechanisms as Value at Risk was to offer leadership detailed projections of the exact probability, to a very narrow range of confidence, of some damaging event or other threatening a company’s position.²³ Risk managers were fond of speaking in very detailed percentages and probabilities: *There is a 1.5 percent chance of a loss of more than 25 percent of our investment.*

Taken together, these factors led many financial firms to develop imposing levels of faith in their ability to manage risk. “A belief had arisen during the late 1990s,” Gretchen Morgenson and Joshua Rosner have written, “that bankers had so improved their risk-management and loss-production techniques that regulators could rely on them and their financial models to develop capital standards.”²⁴ Fears declined as the industry decided—amazingly, just a decade after the collapse of Long-Term Capital Management, a failure grounded in similarly undue faith in the precise estimation and management of risk—that they had cracked the code. Reserves could be cut, leverage grown, and potentially dangerous financial instruments developed, all because procedural risk management could be relied upon to sound the necessary warnings. And if these institutions followed these perceptions, they became hugely leveraged in part *because* they became so confident in their ability to manage risk.

Part of the problem, once again, was a dangerous habit of mistaking uncertain contexts for deterministic ones. In a deterministic context or system, inputs equal outputs, the initial conditions set the parameters for the outcomes, and the information currently in the system is a good guide to future developments and trends. As we’ll see, some risk environments have strong elements of determinism—as in the population data used by actuarial analysts. But nearly all really strategic decisions take place in non-deterministic contexts in which non-linear dynamics and ambiguity about initial conditions means that future possible worlds or scenarios become unmanageable. In such contexts, there is not one potential future world from today’s starting point—there are hundreds of them. Human factors provide major elements of uncertainty, so that non-deterministic environments are also what have been called “transmutable,” meaning that they are constantly evolving and emerging under the influence of judgments and choices.²⁵

These problems, however, should not have been a surprise to senior leaders in the financial sector. They were certainly well-known to professional risk managers, who appreciate only too well the limitations to their approaches and are deeply schooled in issues of determinism, probability, and uncertainty. The problem was that, while in theory risk processes could be applied precisely and carefully, in *practice* they were not. And the reasons have everything to do with the human factors discussed in Chapters 4 through 10. The result was that the most sophisticated financial enterprises in the world could not internalize their own warnings about the dangers inherent in their strategic choices²⁶—a potential flaw whose potential implications for national security are only too real.

Such dangers were on display in the financial crisis, in the mismatch between probabilistic approaches and a context of deep uncertainty. A year after the Deloitte survey, on the very cliff-edge of the crisis, Merrill Lynch CEO Stan O’Neal crowed about recent profits—\$2.1 billion in a single quarter—and

promised smooth sailing through what even then could be seen as choppy waters ahead. Part of his confidence stemmed from the quantified promises of risk managers. Recent profits, he wrote in an e-mail to the firm, “reflected the benefits of a simple but crucial fact: we go about managing risk and market activity every day at this company. It’s what our clients pay us to do, and as you all know, we’re pretty good at it.”²⁷

Within 15 months of that boast, Merrill Lynch was effectively out of business, dumped on Bank of America in a fire sale. One animating question for this study is how this state of affairs could possibly have developed among firms who saw risk management as such a priority. My answer, in the most general terms, is that the failure of risk management in 2007–2008 reflects nothing more than the revenge of both uncertainty and human factors in an environmental of deep uncertainty. It’s a well-known feature of human thinking that we look for convenient harbors from uncertainty, ways of navigating complex questions with no readily identifiable right answer. Risk management was viewed as a mechanism for discovering probabilities to guide judgment in circumstances where possible outcomes are known and can be assigned odds. The mismatch between these two approaches—and their assumptions about the nature of the problems they are dealing with—opened the way for a number of specific human factors to have a destabilizing effect.

In order to make this case, this study unfolds in several sections. Part I frames the issue by defining risk and surveying its current use in the national security context to help make complex strategic judgments. Part II examines the specific lessons of the financial crisis for the conduct of risk management—issues ranging from the subjective character of risk to the role of imperative-driven thinking to how incentives shaped risk. Part III returns to the national security realm, applying many of the lessons of the preceding section, developing a proposed framework for risk in complex strategic judgments in national security, and offering—on the basis of that framework—a tentative assessment of risks associated with the current US national security strategy.

Risk and strategy

A leading implication is that risk management cannot be viewed as separate and distinct from basic strategic judgment. Without sound judgment risk will be mangled, but without a clear consideration of what could go wrong—outcome-oriented risk—strategic judgments are incomplete. The goal is not a fully independent risk management process, but rather *risk-aware strategic judgment under uncertainty*. I will argue that risk procedures should be engines for generating thoughtful and creative mechanisms for dealing with challenges to current and prospective strategies. A risk-informed mentality cannot be separated from a broader mindset underpinning good judgment—of open-mindedness, tolerance for dissent, unwillingness to be captured in an

imperative, and other factors. Good risk management is a product of careful judgment, not the other way around.

One step with surprising importance for managing the risks of strategic choices is the rigorous identification of theories of success that explain *why* the selected means will achieve the desired end. Very often strategies involve mere laundry listing of actions or means: The United States intends to deter North Korea, and so it dispatches a carrier, holds an exercise, and sells South Korea additional military equipment. But what underlying theory explains how these actions will work together to achieve deterrence? In the Korean example, the assumed mechanism is relatively straightforward—that “strength deters”—but even such an intuitively appealing strategic concept might not apply in all cases. In more complex examples the link between means and ends is far less guaranteed. In Afghanistan, for example, the United States spent well over a decade taking dozens of specific actions—employing many discrete means—guided by strategic concepts that were ill-defined and grounded in presumption rather than evidence. Effective strategy is about using carefully selected means to achieve well-defined ends, and the only way to understand that in rigorous terms is to know what mechanism or theory links the two.

One possible interpretation of this notion is that risk processes can contribute to strategic choice by focusing on pitfalls in these theories of success—a close assessment of the *potential* dangers of a proposed strategic concept. Risk analyses can inform strategy by focusing attention on what can go wrong with strategic concepts proposed to link ends and means. This study will argue that, by contrast, two typical ways of conceiving of risk—as a survey of environmental dangers, synonymous with “threats,” and as a gap between means and ends—may be critical to a full-informed strategic dialogue, but neither is best conceived as a “risk” exercise.

Finally, this approach points to a critical way of conceiving strategic analysis that goes beyond risk itself. Given the nature of the judgments being made and the character of the context for them, what senior leaders are really doing, I will argue in subsequent chapters, is not managing risk but rather managing uncertainty. This seemingly semantic distinction carries critical implications for the mindset leaders bring to the task, and the specific types of strategies they use to advance their interests. I will propose a specific concept of *managing uncertainty for competitive advantage* to capture the full scope of the needed approach.

Senior leaders of the US national security enterprise have been right to highlight risk as an issue that needs more attention. The chapters that follow first sketch out a number of primary lessons of recent experience, and then offer a number of suggestions for how best to organize a risk-aware strategy process. In the process the analysis will offer some judgments about the status of the US global posture and potential avenues to better mitigating risk.

2

Defining Risk

Toward the end of the 2014 Quadrennial Defense Review (QDR)—the primary long-term strategy document issued by the U.S. Department of Defense—there appears the “Chairman’s Assessment,” a comment on the document by the senior-most officer in the U.S. military (at the time, Army General Martin Dempsey). The federal statute mandating the QDR states that it must be conducted “in consultation with” the Chairman, and requires that the Chairman’s office produce a formal assessment. Over time, the Chairman’s comment has come to focus on the risks embodied in and taken by the defense strategy laid out in the QDR.

“Today the U.S. military can conduct all of these missions,” Dempsey wrote about the defense tasks anticipated by the QDR, “but under certain circumstances we could be limited by capability, capacity and readiness in the conduct of several of them.” He referred to the decision of the QDR to “take risk” in various areas—by which he meant that the proposed defense program offered fewer resources to accomplish largely unchanging objectives.¹ The Chairman’s annex also described a range of things that could go wrong in the overall strategic context, such as threats to US interests and rising technological capabilities of adversaries:

[I]n the next 10 years, I expect the risk of interstate conflict in East Asia to rise, the vulnerability of our platforms and basing to increase, our technology edge to erode, instability to persist in the Middle East, and threats posed by violent extremist organizations to endure. Nearly any future conflict will occur on a much faster pace and on a more technically challenging battlefield. And, in the case of U.S. involvement in conflicts overseas, the homeland will no longer be a sanctuary either for our forces or for our citizens.²

In these and other comments, the Chairman’s assessment portrays risk in a number of ways. It discusses what could be called contextual risk—dangers in

the overall strategic environment, threats to the United States and its interests such as belligerent actors. It stresses the risk inherent in a gap between means and ends, the danger inherent in failing to provide sufficient resources to accomplish stated goals or fulfill requirements. And it outlines developments or trends that could get in the way of intended policies or programs and prevent US strategy from achieving its goals.

Losing hold of a concept

Strategy, one classic framework suggests, is about the interrelationship between ends, ways, means, and risk. Of those elements, the terms “ends” and “means” are broadly understood; their definitions are widely agreed, their meaning and role in the strategy process generally accepted. The notion of “ways” is more abstract, but theories of strategy offer a clear and well-established definition: Ways are the manner in which the means are employed to achieve the end.³

If we ask, on the other hand, what we mean, in the making of strategy, by “risk,” there is no shared answer. Some will say “gaps between ends and means.” Others will say “threats.” Still others will say, “Dangers created by my strategy.” Some might try to approach a definition by simply listing various categories of risk, such as operational, strategic, and institutional.

This basket of meanings is symptomatic of the fact that the term “risk”—even when used in specific contexts, such as business or national security—is now employed in a dizzying variety of ways, often in the same organization or sometimes in the same analysis. Companies take risk when they don’t develop their talent properly. They face the risk of suddenly shifting financial markets or belligerent opponents. They embrace risk in their degree of leverage; if their credit-worthiness slips, they are assuming more risk. They add to risk when they reduce the resources devoted to achieving a given end. Even within the national security realm, or even more specifically within the Department of Defense, risk can refer to many different issues.

In the corporate arena, companies play with multiple risk categories—credit, strategic, financial, market, reputational, regulatory, operational, human resource, and more. A very specific and painfully detailed taxonomy has arisen by which investment banks and other financial institutions conceive of risk. In some firms, in fact, important elements of risk analysis and mitigation are the responsibility of the security department; and while robberies or attacks on buildings are indeed threats, they do not generally reflect the core strategic risks a firm takes with its activities. “Such organizational silos disperse both information and responsibility for effective risk management,” argue Robert Kaplan and Anette Mikes. “They inhibit discussion of how different risks interact. Good risk discussions must be not only confrontational but also integrative. Businesses can be derailed by a combination of small events that reinforce one another in unanticipated ways.”⁴

Each of these forms of risk offers an important category of information to senior leadership. Yet the welter of categories can create as much confusion as understanding, especially when different offices in a firm are responsible for distinct pieces of the risk pie. In some cases the dividing lines between them are not always clear; in others, various forms of risk could easily be seen as sub-sets of other categories, and any office or group responsible for a comprehensive risk assessment in one or the other area could not account for risk without looking at both or all of them.

When risk is employed in multiple, unreconciled ways, it loses any coherent meaning, or purpose in the strategy process. “Managing risk” in national security could refer to anything from dealing with threats to the nation to providing sufficient resources to accomplish particular missions. Risk becomes almost synonymous with strategy itself, encompassing a survey of the environment, the balance between means and ends, the feasibility of proposed courses of action, and a half-dozen other things. There is reason to believe that the proliferation of risk categories contributed to the recent financial crisis by hiding substantial and cross-enterprise risks in a dozen different formulations. When senior leaders refer to “risk” today, it is no longer clear what exactly they mean.

A central challenge is the lack of an agreed-upon definition. Even within the same organization—the Department of Defense, for example—people talking about risk can end up speaking past each other because they are thinking of fundamentally different things. General Dempsey himself recently suggested in an interview that “I’ve discovered that the two hardest words to adequately articulate in my line of work are ‘risk’ and ‘readiness.’” The concept of risk was so challenging, he explained, “because the meaning is so dynamic. It’s a combination of capability and intent on the part of those people who would do us ill, and frankly, you can measure capability but it’s hard to measure intent. So risk is extraordinarily difficult to articulate in a way that people understand.”⁵

The problem is similar in the financial sector. There much more thought has gone into defining risk and the practice of risk management is built, as we will see, on an exhaustive typology, such as market, credit, operational, and reputational risk. But in the process the meaning of the overall term—and the effort to assess and mitigate it—can be stretched to the point that it begins to lose coherence. And an important lesson of the financial crisis is that a vague and ill-defined mission for the risk function will not have enough bite when it runs up against the malign influences of human factors. An unintended consequence of the proliferation of risk management procedures is that the assessment and consideration of risk has sometimes become a sort of background noise rather than a specific, and highly influential, moment in the strategy process.

This study contends that risk analyses can best contribute to making effective complex strategic judgments by focusing attention, rigorous assessment

and mitigation efforts on the problem of *what could go wrong with established or proposed courses of action*. The most useful role for risk management at the strategic level is not in assessing the environment, generating “foresight,” massaging algorithms to measure a firm’s capital adequacy, or pointing to the dangers of gaps between means and ends. Those are all worthwhile components of a strategic process, but they ought to be filled by other steps and functions. Meantime, the most likely source of disasters—a failure to think seriously and honestly about the potential consequences and feasibility of the strategy itself—often manages to escape rigorous analysis.

To be very clear, my argument does not presume that organizations never use risk management processes to examine the consequences of possible strategies. Many do it all the time, assessing what could go wrong with a proposed move. One can find risk-related sections in various national security documents. Case study accounts of national security choices often reflect a persistent attention to what we would call outcome risk—if we make this choice, what could go wrong?

But a review of the role of risk in informing complex strategic judgments in the financial crisis and in current national security practices suggests three things. First, risk is not *primarily* used for consequence management: It is employed for a whole range of objectives, and few organizations focus it entirely on this one purpose. This tends to dilute the effect and, most importantly, means that organizations do not obtain a mindset of rigorous outcome management. Second, much of the outcome focus that does emerge is tactical, operational, and program-level rather than strategic. Organizations assess the possible results of a specific investment and trace ongoing risk in a particular program. It is more uncommon to find true enterprise-level strategic risk management with a focus on consequence management. And third, outcome-oriented risk assessment at the strategic level is often casual and haphazard rather than organized and rigorous. The result is that major strategic judgments often take place without a careful and organized consideration of what could go wrong.

The essential character of risk

In the broadest sense, risk refers to potential negative events in relation to the interests and objectives of a group, dangers that must be kept in mind while developing plans and programs. Most sources thus agree on a rough definition of risk: It is the potential for something to go wrong in relation to a country or organization’s interests. One source defines risk as “the probability that a particular adverse event occurs;”⁶ another calls it “probability and magnitude of a loss, disaster, or other undesirable event.”⁷

Put another way, risk is the threat of harm or loss that may adversely affect the ability of the organization to accomplish its mission.⁸ The Department of

Homeland Security defines risk as “the potential for an unwanted outcome resulting from an incident, event, or occurrence, as determined by its likelihood and the associated consequences.”⁹

A common technical definition of risk places emphasis on the informal sum of the “probability times the consequences” of a potential danger. A devastating possibility could still be of moderate concern if it is exceptionally unlikely, and vice versa. But probability and consequence are broad categories, and together offer only a generic framework for assessing any potential event. It could be an opportunity as well as a risk. The essence of risk relates to something that can go wrong, which can then be assessed by its probability, consequences, or a number of other factors.

Risk has a number of other typical aspects. First, it is only of concern in relation to something of importance to a state or organization.¹⁰ Risk is the possibility that something bad will happen to *something we value*.

Second, risk is integrally related to chance. Risks are always possibilities or probabilities, never certainties. When an organization can fully anticipate and plan for an outcome—“this start-up will require fifteen staff to begin at the following likely salaries”—that is a planning function (such as cost). Or an analysis could point out an inevitable, unavoidable result of a choice: Ending arms sales to Israel will anger the Israeli government and its supporters in the United States, for example. Again, that is a certain and expected outcome of the action, a planning factor. But it’s not a risk. Risks by nature are only *possibilities*.

One U.S. Defense Department risk guide puts the distinction this way. “Risks should not be confused with issues,” it suggests:

If a root cause is described in the past tense, the root cause has already occurred, and hence, it is an issue that needs to be resolved, but it is not a risk. While issue management is one of the main functions of [Program Managers], an important difference between issue management and risk management is that issue management applies resources to address and resolve current issues or problems, while risk management applies resources to mitigate future potential root causes and their consequences.¹¹

Even in its true, probabilistic sense, the term “risk” is too often used as a mere synonym for “chance.” A statement that refers to “the risk that a satellite could be hit by space debris” is really talking about the simple potential that it could happen, not a specific probability. Risk is often used in this commonplace manner, in place of chance: the “risk” that a card will be dealt in a poker game, for example, or that someone will fall down a flight of stairs. This may be perfectly acceptable in general usage, but as a component of strategy, risk means something far more than chance.

James March and Zur Shapira have emphasized that actual decision-makers often conceive of risk in ways different from its theoretical foundations. They

do not consider it as part of a comprehensive, outcome-based evaluation of dangers against opportunities, for example, but see risk as an inherently negative concept. More importantly, they do not, often, think in probabilistic terms, even when relying on quantitative risk assessments. Executive decision-makers appear to be much more influenced by the scale of potential risks, for example, than their probability.¹² Many of the issues laid out in Chapters 4 through 10 add more factors to this list of things that distinguish risk as a theoretical construct from the actual notion that decision-makers hold in their heads: Real decision-makers are more influenced by many subjective factors than the objective results of risk processes. A major lesson of the financial crisis, indeed, is that this distinction can lead to disaster.

Risk and reward

An important way in which the conception of risk in national security differs from that in finance is the relationship between risk and reward. Investment decisions, as with many business choices, reflect a balance between the two: more risk is generally associated with greater potential reward, while less risky investments tend to have more modest returns. In the risk-reward balance, the risk side of the equation is sometimes viewed as the possibility of loss, often permanent loss, in the underlying value of an investment. Other models treat risk as volatility, the potential for dramatic swings in value. Either way, integrating risk into strategic choices involves weighing the balance between potential loss and potential gain; and in the process, to the extent that risks can be anticipated and priced, they can be made a formal part of the calculations that attempt to objectively compare the two sides of the equation.

Many risk management models have therefore used market histories and other long-term data sets to assign a specific price to the risk of a given investment. Such risk pricing, in fact, was a major function of financial industry risk managers before the crisis—to generate reliable estimates of the relationship between reward and the potential for loss. It was when such modeling became too deterministic, when it underestimated volatility, or when it failed to account for “fat tail” swings in markets that it helped produce disaster. But the basic concept, that the price of risk involves some *relationship* between potential gain and potential loss, remains at the heart of financial considerations of risk.

This relationship differs from the most common conceptions of risk in national security. While there may be opportunities to be seized, the dominant ways of conceiving national security risk focus on the danger or threat side of the equation and pay little attention to the balance with possible gain. Investments in national security are typically made in order to forestall threats. In financial calculations, high levels of risk can be accepted, even encouraged, when they are part of an overall strategy of achieving gains; in

national security, states mostly want to avoid risk. The problem is that there are usually insufficient resources to do so adequately. The challenge then becomes to estimate the relationship, not between risk and reward, but between *resources and risk* in order to design a minimally acceptable posture and an efficient allocation among various threats.

In finance, simply put, investors are taking and balancing risk as part of a strategy to earn rewards. Managers of national security are allocating resources to cancel out risk. This is not the whole story, of course: Many national security strategies are undertaken to gain something; states are not always in a defensive posture. But the risk and reward sides of the equation are not always directly compared, and certainly not rolled up into a single revenue or pricing estimate.

This is not surprising, because it is exceptionally difficult to apply objective values to elements of risk in national security. When considering the degree of threat, for example, senior national security leaders always look at the same information and come to very different conclusions about the level or degree of threat posed by a given nation or issue. A risk management process could try to resolve these problems through surveys of experts, but it is not clear that the median of such estimates would be more accurate than any specific estimate. Even a majority agreeing on a general range might only validate conventional wisdom.

Another variable central to any risk-reward comparison—the question of what interests the United States has at stake in an issue—is also subjective, and often impossible to know in advance. Washington spent decades thinking it had huge stakes in Vietnam, only to shift and decide that it did not, and could safely withdraw. Interests are a function of visions of the US role in the world, which are grounded in theories and worldviews rather than data. There is likely to be no objective way to resolve these disputes.

Nor could such a process generate a reliable estimate of a third variable—the likely effectiveness of specific actions designed to reduce risk. Indeed this may be the very crux of the problem of estimating risk in national security: The general calculus of riskier or less risky actions simply cannot be known with any reliability. Those who worry about Russian aggression in Europe today see large deployments of NATO forces in the Baltics as a risk-reducing action, because of the deterrent effect they believe it will have. Those less worried about such aggression, or more concerned about Russian paranoia, see such reinforcements as highly provocative and therefore far more dangerous. There will never be a consensus about which strategic choice embodies more or less risk, and the risk/reward calculus will always be a function of belief.

This is yet another reason why simplistic means-ends sufficiency measures do not adequately capture risk. Some observers believe that reducing the size of the U.S. Army or Navy will create dangerous risk; others do not. Either way the beliefs are typically a function not of detailed analysis but of the

observers' larger theories about the strategic context and causal relationships involved. Indeed, national security judgments are so heavily based on theories and worldviews precisely because such complex social patterns where (as we will see) each case is fundamentally unique offer less basis for reliable probabilistic assessment.

There is yet another problem, related to another fundamental aspect of complex strategic judgments: The existence of multiple values or objectives. When pricing risk in finance, leaders are dealing (usually) with a single goal—profit. A relatively straightforward approach can calculate risk against a single variable. But national security leaders will have multiple values to consider, and multiple objectives, with any decision: This is another fundamental aspect of complex judgments in this field. In making a choice to reduce national security risk, for example, a president would have to think (at a minimum) about domestic political reactions, allied perspectives, the effect on military readiness and operational tempo of new deployments, the costs involved, the perspectives of senior military leaders, and more. What emerges, then, is the need for a whole series of overlapping risk assessments: Risk to security, risk to political viability, risk to the health of the force, and more.

Ultimately these distinctions between financial risk-reward approaches and the way risk has been treated in national security suggest three things. One is that risk assessments cannot provide a singular basis for judgment. There are too many different goals and variables, and there is far too much uncertainty involved. At best these efforts offer insight and perspectives to help guide what will ultimately be a subjective, largely intuitive decision. Coming chapters will have more to say about that.

Second, approaches to risk in national security have more in common with project-based risk than financial risk. Projects have risks such as unforeseen changes in the market, new regulations, rising cost of raw materials, and so on. Conceiving national security risk more in terms of strategies, projects, or initiatives would generate a more direct parallel. Project-oriented risk assessments often try to inform judgments under uncertainty rather than conveying objective risk prices. This mindset is closer to what we find in national security, where senior leaders are constantly trying to evaluate the risk involved in various options.

Third, and finally, these considerations reinforce the need for outcome-oriented risk approaches in national security. The essential approach must be to place a focus on the potential consequences of actions—the risks of making a specific choice, rather than the risks in the environment (which can be evaluated as threats) or the risks of applying insufficient resources (which can be addressed as sufficiency). The most important way that dialogues on risk can contribute to national security strategy is by focusing attention on an oft-neglected question in strategic decision making: If we choose this option, what could go wrong?

This approach would have something in common with the risk component of the risk-reward calculus in finance and other industries. There, too, senior leaders are considering a choice—whether to invest or change markets or develop a new product. They want to understand the possible consequences of their choice, in terms of both danger and opportunity. That essential mindset, a strategist’s conception of looking forward to anticipate how the world will unfold under the influence of their choices, ought to inform the employment of risk.

Risk versus uncertainty

A major theme of this study is that the context of uncertainty (which will be defined more fully in Chapter 4) must decisively shape the understanding and analysis of risk in complex strategic judgments. Under uncertainty, it becomes more difficult to assign meaningful probabilities in the manner demanded of a deterministic view of risk assessment. In some sense, one-half of the traditional understanding of risk—the concept of likelihood—disappears altogether. We simply cannot know the “probability” that arming Ukraine will prompt a Russian escalation, or that offering North Korea a concession will generate one behavior as opposed to another. Too many variables and unknowns intrude in the causal chains.

Suppose that the United States is set to invade Syria to depose the Assad regime. What is the *risk* that Iran will react with herculean efforts to undermine the US and allied presence? Various forms of information, overt and covert, will offer some clue as to whether it might be high or low, but that information will rarely if ever reach a degree of precision to allow a probabilistic estimate.¹³ This is due in part to the simple fact that *the action itself will change the context* in ways that make it difficult if not impossible to forecast the world in which consequences will unfold. These sorts of risks share many of the characteristics of deep uncertainty, and they are the most common risk variables that senior decision-makers in business and national security have to grapple with.

This is not to suggest that risk models, quantitative approaches, and data-based analysis have no place in complex strategic judgments. They do, but their limits must be carefully understood. A detailed assessment of the possible effects on a given conflict of a strategic choice in force structure can play an important role in informing that choice—but the outcome of that assessment is not *equivalent* to the risk judgment being made. In any complex strategic judgment, that decision will incorporate dozens of highly subjective factors. The problem, as reemphasized by the financial crisis, is not in using highly procedural risk management based on quantitative models: It is in using them as a substitute for strategic judgment, rather than an input to it.

It is important, then, to take seriously the potential implications of confusing the relationship between risk and uncertainty.¹⁴ As the risk analyst John Adams puts it, uncertainty “is the realm not of calculation but of judgment.”¹⁵ My argument is that *all* important complex strategic judgments have such a character. This does not mean that there is no risk involved, or that thinking about risk cannot be valuable. It merely suggests that risk remains highly uncertain and subjective. The distinguishing feature of risk is not that it is measurable, but that it focuses specifically on the chance that something could go wrong.¹⁶

Yaacov Vertzberger agrees that all complex national security decisions contain substantial uncertainty, and so to distinguish classic probabilistic risk from uncertainty in that field makes little sense. Uncertainty, Vertzberger suggests, has to do with a decision-maker’s level of information and foresight, while risk has to do with our expectations of uncertain *outcomes*. He defines risk as “the likelihood of the materialization of validly predictable direct and indirect consequences with potentially adverse values, arising from events, self-behavior, environmental constraints, or the reaction of an opponent or third party.”¹⁷ Uncertainty is the context in which such judgments play out, whereas risk has to do with one’s expectations about outcomes or consequences.

The conundrum of the general and the particular

One of the most important reasons why strategic judgments are immune to procedural risk assessments is the role of what can be called the general and the particular. Risk management models—and indeed data-based modeling of all sorts, including that used by hedge funds to make short-term bets and by “big data” analysts who forecast election outcomes based on exit polls and other data—thrive when they can employ large amounts of data reflecting thousands of relatively comparable events. A study of ten thousand similar men taking a given drug can obtain reliable results of its likely effects, especially when the number of participants and their behavior allows for isolating specific variables.

The problem is that complex strategic judgments are usually trying to make sense, not of the latest example of a sea of similar cases, but of one unique event. Every large-scale sociopolitical event is built around a momentary set of variables; no two events will be comparable in the way that measurements of height or weight or intelligence might be. In such a context, decision-makers have no reliable set of examples from which to draw patterns.

The role of uncertainty points to the critical importance of this gap between the universal and the particular. A perfectly balanced roulette wheel allows for probabilistic statements about the ball coming up red or black, or landing on a specific number, because of the physical design of the system. A sample of the height of American men for fifty years can produce a reliable portrait of

the distribution of heights, the probability of how many six-foot-eight men will appear, and so on. In other words, in some systems, the individual and the general stand in some meaningful relation to one another, one that can be sketched out with probabilities and even in some cases forecasted.

In systems characterized by high degrees of uncertainty, however, such as the economy, the financial market, and the international security context—a critical defining feature is the *incomparability of the specific and the general*. If we ask, What is the next year likely to look like on the stock market, the fact that the last five years might have seen limited volatility and substantial overall gains tells us little of value. Asking the same question before 2007 would have produced misleading answers. The context is *nonergodic*, a system in which future patterns cannot be inferred from present facts. There will be broad trends and “truths” that are widely accepted, but even these will be violated by out-of-the-norm outcomes whose potential is obscured by those same conventional wisdoms.

There are limits to this problem, even under uncertainty. As changeable and nonlinear and subject to random fluctuations as the market may be, it still provides billions of data points in thousands of categories.¹⁸ Regularities sometimes can emerge, especially in regard to short-term fluctuations. Some corners of international politics allow for the more effective accumulation of reliable data sets of comparable cases—in terms of society-wide issues such as disease or public opinion, for example.

Subjective probability assessment

One way that rationalistic decision models try to escape these various implications of uncertainty is to offer a less deterministic, fuzzier version of probabilistic judgment. A “subjective” probability evaluation is always available, in which decision-makers take the information they do have and make a best-guess about the potential for various outcomes.¹⁹ As opposed to objective, data-based probability, subjective probability theory (sometimes called “the personal theory” of probability) is based on individual experience and beliefs. It measures how probable someone *believes* something to be. Such beliefs are then taken as a meaningful proxy for a true deterministic analysis. We can’t “know” whether Iran will respond violently to a US invasion of Syria—but based on our experience, the intelligence and other sources of information at hand, and other sources of data, we can make a subjective judgment that’s probably not too far off.²⁰

As useful (and frankly inevitable) as they are, subjective assessments of risk are also the window through which human factors squirm into the equation.²¹ Decision-makers do assign subjective values to many potential outcomes; the problem is that those assigned values are often wildly inaccurate or are demonstrably based on some specific bias. As even defenders of subjective probability admit, moreover, the values are entirely dependent on the

worldview of the person making the estimate. Before 2007, the subjective, gut-level probabilities assigned to various risks by senior leaders of the investment banking sector did not begin to capture the looming dangers of derivatives and subprime mortgages. The errors resulted from a number of identifiable biases and group dynamics, from overaggressive organizational cultures to simple wishful thinking, that skewed the subjective estimates.

The distinction is to some degree one between a laboratory or theoretical setting, and the actual context of real decision-makers. As even some leading researchers in the heuristics and biases tradition have demonstrated, when deeply educated about issues like base rates and potential biases, individuals can become reasonably good at subjectively estimating probabilities even under a degree of uncertainty. Philip Tetlock's work with a series of "good judgment" projects has demonstrated something similar: Groups of "super-forecasters," working in a highly disciplined, peer-review setting and often using a range of complementary decision aids, can become reasonably good at anticipating near-term outcomes of uncertain events.²² This effort actually provides an outstanding example of the ways in which reasonably objective analyses can usefully inform strategic judgment.

But in the actual world of real senior decision-makers—under the pressure of many of the incentives that will be examined later in the book—human factors tend to skew and bias foresight and probability assessments, and make subjective probability assessments an extremely chancy and even dangerous proposition. Institutions are not led by super-forecasters. Their senior leaders have neither the time nor the inclination to employ such disciplined processes. Often, as we will see, they refuse to listen to those that do—many skilled analysts forecast the collapse of 2007 well in advance, but their warnings fell on deaf ears because of the same human factors impelling the risk-taking in the first place.

But when dealing with the far edge of uncertain environments—highly nonlinear interactions between swirling, complex environments and the possible outcomes of various actions—even highly disciplined subjective probability assessments still amount to guesswork. (A careful and nuanced scholar, Tetlock recognizes the limitations of his forecasters in terms of nonlinear issues and long-term time horizons.) Predicting large swings in the economy is precisely such a "cloudlike" endeavor, which may help to explain why studies of formal government or quasi-governmental predictions have found them to be routinely terrible at forecasting. One IMF study, for example, discovered that of the 60 recessions that struck the global economy or specific nations during the 1990s, only two were forecast in advance—and even those were not seen until about six months beforehand.²³

Objective models and processes can actually be very useful even to inform complex strategic judgments, when they are used to check, question, and correct instinctive leaps rather than to justify and exacerbate them. Models such

as Value at Risk—and the results of any forecasting process—can be used to bolster preconceived notions as well as to spur a rigorous analytical approach. The deciding factor isn't the quality of the model. It is the leadership, analytical rigor, and organizational culture of the enterprise making the strategic choice.

Risk management: strategies for dealing with risk

The comprehensive discipline of risk management is about much more than assessing and warning of risks. It is also about mitigating them. The field of risk management has outlined several kinds of strategies for mitigating risks.²⁴ The literature generally mentions a number of leading approaches. These include:

- *Avoidance*. Taking a risk as unavoidable and making changes—in goals, program design, operational concepts—to sidestep it.
- *Transference*. Shifting risk to other parties through a variety of means: partnerships, subcontracting, alliances.
- *Mitigation*. Taking active steps to reduce the (1) likelihood or (2) impact of a given risk.
- *Acceptance*. Simply consenting to the risk and taking account of its possible effect on operations.

National security processes seldom examine risk mitigation with such specific categories. Typically, mitigation strategies are assumed or discussed only in the most general terms—which often reflect proposals to throw more resources at a problem, either to respond to a threat in the strategic context or to close a supposed ends-means gap. Some of these strategies are commonly used, though not necessarily referred to in these explicit terms.

Each of these approaches could be appropriate in a given circumstance, given the nature of the risk, the importance of the plan being tested, the degree of involvement by others, and other factors. Deciding which risk management approach makes most sense is a key judgment in the process. Psychological dynamics involved in risk processes make it very likely that leaders will fund subconscious ways to ignore risk. But the subconscious justification often precedes the formal analysis, meaning that the decision-makers are really just finding excuses not to take the risk seriously. We find this pattern in cases ranging from the financial crisis to the Iraq war, and it is a leading source of human-factor risk calamities.

The dangers of relying on risk analysis

The economist Jack Dowie has made an iconoclastic argument that we should abandon the use of risk in strategic dialogues. “There’s no need to use the

word 'risk' in the process of identifying the best course of action,"²⁵ Dowie has argued. Elsewhere he concludes that, "If the intended aim of improved understanding of how risk or risks are perceived, assessed, communicated, amplified, attenuated, or so on, is to produce better decisions ... then much time, energy and money is being wasted."²⁶

Dowie offers a number of reasons why this is the case. The notion of risk, he contends, "is not an operational concept. It is a word with a cluster of meanings that is used in a multiplicity of ways to refer to, or embrace, three operational concepts" such as probability and utility. To the extent that risk managers are contributing to strategy, he suggests, it is because they are working on issues that bear on probability or utility.²⁷ (The analogous argument of this study is that concepts like sufficiency and threat are operational, from the standpoint of strategy—but thinking of them as risk does not add anything to those issues.) Misunderstandings about the nature of risk, he suggests, are "especially likely to occur when people are in the emotionally-charged situations of serious decision-making."²⁸

Dowie contends that there are "perfectly good, precise and accurate terms" for all the concepts that risk is used to assess.²⁹ But by using risk instead of more precise and accurate terms—such as threat or sufficiency—risk management processes cloud the true meaning of the issues and obstruct effective responses. These shortcomings of risk analysis contribute to a problem that will be discussed in some detail in Chapter 8: Warnings of the most common class of risks—what can be called "gray swans"—frequently go unrecognized, in part because of the fact that risk assessment processes do not come to grips with the complications of uncertainty. "Simple risk warnings," Dowie recognizes, are "likely to be ineffective" because they are necessarily ambiguous.³⁰

The financial crisis validates many of these dangers. It was in fact the very proliferation of risk management terms and procedures—along with a creeping ambiguity in understandings of the term—that created a sort of analytical fog around risk. We see the same phenomenon in national security, where, as we will see in the next chapter, risk has become a term that means everything and nothing at the same time.

Interestingly, Dowie believes that, far from desiring risk management to generate disciplined, precise quantitative estimates, senior decision-makers actually prefer a process that floats on a vague confusion of terminology. He argues that senior decision-makers are far more comfortable making "gut calls" by applying their experience to largely intuitive judgments. They are therefore happy with a risk process that remains somewhat imprecise, because it allows them to make the emergent intuitive judgments they want to make from the beginning.

In several chapters of this study I will make more or less the same argument—that complex strategic judgment under uncertainty is dominantly

(and necessarily) an emergent process of “feeling-of-knowing” impressions, and that risk processes fail in large measure because they cannot penetrate that process. This is another way of stating one of the central themes of the study—that human factors defeat procedural risk management. The answer is not a highly deterministic or quantitative process to generate objective answers, except on the most discrete engineering-style problems that are amenable to a “correct” answer. But Dowie might agree that a leading purpose of strategy processes is to help improve intuitive judgment with greater discipline; what we think of as “risk” can most helpfully be seen as various elements of strategic logic. This study contends that, to the extent we preserve the term “risk,” it should be used to energize a focused discussion of the possible consequences of strategic actions.

A concept of risk under uncertainty

The main reason to be clear about what risk means is to create a process that meaningfully adds to the development and execution of strategy. Some possible concepts of risk merely duplicate other elements of strategy-making. If we define risk management as identifying threats and generating mitigation strategies for them, for example, it would appear to amount to the same thing as a general environmental assessment, and hardly seems worth the trouble. At the same time, examining the relationship between means and ends is essential but already ought to be built into the basic analysis of a strategy process.

An implication of this volume’s main theme—the dangers of human factors in strategy and risk analysis—is that formal risk processes have one dominant purpose: to safeguard against the ways in which human perception and cognitive bias can go wrong in the strategy process, specifically in downplaying or ignoring outcomes. Arguably the most common route to strategic disaster—and to failures of risk management—is the tendency of senior leaders to ignore potential consequences as they rush excitedly into some favored plan. Once senior decision-makers have decided on a course of action, reflective thinking about the risks of that course is far less likely: This is a fundamental lesson of the financial crisis, and it is a primary feature of most national security debacles as well.

Taking all this seriously points us in a single direction for conceiving of risk and risk management: the identification, assessment, and mitigation of outcome-oriented risk, taking a current or proposed strategy as a starting point and asking what could go wrong in its consequences. This approach would leave the assessment of threats to the environmental assessment, and leave the means-ends gap (mostly) to the basic strategy analysis itself. (Some aspects of that gap might count as true success-obstructing factors, and could be treated as risk.) It would use risk management mainly to force a discussion

of *what could go wrong* with one or more proposed strategies. We ought to be discussing risk *in relation to* a policy objective or proposed policy: *identifying and developing mitigation plans for possible negative outcomes, unanticipated consequences and second and third order effects of national security strategies adopted to mitigate vulnerabilities and threats*. And I will argue that it comes closest to the most useful way of employing risk to enhance strategic judgment.

3

Approaches to Risk in National Security

The use of risk as a concept in defense planning, and the use of procedural risk frameworks, have proliferated in the national security field over the last decade. These approaches tend to cluster naturally into a few basic categories—program, force management and operational risk being three common ones. The question, in the context of the experience with risk in the financial crisis, is whether these approaches are likely to improve strategy—or pose a danger to it.

Many of these efforts are targeted, discrete, and provide useful input into Defense decision-making at various levels. Flying squadrons, for example, use operational risk management to mitigate the perils of operating aircraft in very demanding environments.¹ Some services have developed excruciatingly specific risk management frameworks; one of the most well-established was the Air Force's Capability Review and Risk Assessment (CRRRA) process² that derived various desired effects to be achieved under Air Force concepts of operations, and then traced required capabilities and identified gaps and shortfalls.³ Acquisition programs use risk to keep track of potential threats to long-term development programs.⁴

This chapter surveys the dominant current ways of conceiving risk in national security. Its focus is at the strategic level; dozens of program- and service-specific risk processes are in place, but this analysis is about the use of risk to enhance strategic decision-making. It argues that two default ways of thinking about risk—as synonymous with threat, and as the gap between means and ends—while they reflect critically important steps in strategic logic, have significant shortfalls. The chapter concludes that a new approach is required. The study will then turn to the lessons of the financial crisis to suggest what that approach can most usefully be.

An important distinction is between operational or program-level risk and strategic risk. One of the central themes of this study is that risk management is often applied with great rigor and detail to the lower end of this spectrum—exhaustive, engineering-based risk assessments of a new shipbuilding

program, for example. But it is used less frequently and effectively to support truly strategic decisions involving enterprise-level decisions about broad direction: Should a company dive into a whole new market in fundamental ways; should the United States intervene in Syria; should the Army be reconfigured entirely for counterinsurgency warfare. These big bets usually involve far more important stakes than program-level or operational decisions, and yet risk assessment is much less reliably employed to inform these choices. And while the US national security enterprise has elaborate and well-practiced risk mechanisms for its more prosaic choices, more attention is needed to the role of risk in shaping strategic decisions.

The DOD framework for risk

Within the overall Defense enterprise there are literally hundreds of program-level risk frameworks in place, to measure everything from safety in military flying squadrons to operational safety in nuclear plants to program risks in the acquisition process. Many of those issues—and their associated conception of risk and risk management—differ in important ways from complex strategic judgments under uncertainty. The analysis here is not meant to apply to such discrete, program- or operations-level risk, and this section does not examine such frameworks.

For the purposes of this study, the most appropriate statement of institution-wide and strategic-level risk appears in the 2001 Quadrennial Defense Review. It introduced a new framework for assessing risk in the national security enterprise that has since become a standard approach throughout various parts of DoD, and subsequent defense planning documents.⁵ “Managing risk is a central element of the defense strategy,” the report contended, arguing that, in part, this process “involves balancing the demands of the present against preparations for the future consistent with the strategy’s priorities.” The QDR outlined four specific categories of risk management, which have become standard elements of the DoD risk framework:

1. *Force management* risk, which the report defined as threats to “the ability to recruit, retain, train, and equip sufficient numbers of quality personnel and sustain the readiness of the force while accomplishing its many operational tasks”;
2. *Operational* risk, dangers to the capacity to achieve military objectives in conflict;
3. *Future challenges* risk, threats to an ability to “invest in new capabilities and develop new operational concepts needed to dissuade or defeat mid- to long-term military challenges”; and
4. *Institutional* risk, threats to the ability to manage defense resources efficiently.

These same categories were employed in the 2010 QDR, a signal that the defense establishment is seeking consistency in long-term risk management approaches.⁶ The 2010 QDR defined risk as “those key shortfalls or complex problems that threaten the Department’s ability to successfully execute its priority objectives.” Similar frameworks crop up in various Joint Staff and Service analyses and approaches.

A related framework is used by a number of the services and other commands and offices within DoD. Its two basic components emerged in a comment by the commander of US forces in Afghanistan in 2014, General John Campbell, in just one example, who told NPR that “I constantly, every day do assessments on ‘risk to force’ and ‘risk to mission.’”⁷ Risk to mission is a form of program or operational risk; risk to force is institutional. If the Navy cannot fulfill its intended concept of operations in the Pacific, that is a risk to mission; if it cannot recruit enough sailors of high enough quality, that is a risk to force.⁸

The 2010 QDR also added a relatively brief mention of a new category—“strategic, military, and political risk.” It defined *strategic* risk as “the Department’s ability to execute the defense priority objectives in the near term, midterm, and long term in support of national security.” *Military* risk was seen as “the ability of US forces to adequately resource, execute, and sustain military operations in the near- to midterm, and the mid- to longer term.” And *political* risk had both an international and a domestic component: Abroad, it was “the perceived legitimacy of our actions and the resulting impact on the ability and will of allies and partners to support shared goals”; at home, “political risk relates to public support of national strategic priorities and the associated resource requirements in the near term, midterm, and long term.” These were not spelled out in detail and have not played a significant role in subsequent public risk management statements. Again, they reflect a hodgepodge of considerations, from international dangers to the sufficiency of the US national security posture to global reactions to US strategic moves.

The 2010 QDR concluded that “This QDR identified areas of weakness in our defense program, presented options to mitigate them, and made recommendations on where and how to rebalance the Department toward our most pressing challenges.” The basic theme, then, was risk as a function of shortages of needed capabilities to achieve goals. This theme—of risk as a gap between means and ends, or requirements and capabilities—may be the dominant way in which risk is presented in most official Defense Department documents. This remains the theme, for example, of the risk analysis in the 2014 Quadrennial Defense Review.

Interestingly, the analytical frameworks commonly used by non-Defense federal agencies and departments tend to highlight two other ways of thinking about risk. Perhaps the most extensive approach to contextual risk can be found at the Department of Homeland Security, and its approach is focused largely on what a classic strategy process would call “threats”—identifying and

assessing dangers in the environment.⁹ This risk function at DHS prepares one of the most elaborate environmental risk assessments in the government—a “Strategic National Risk Assessment” that basically amounts to a list of possible risks in the context. These include flood, disease, hurricane, volcano, wild-fire, and so forth.¹⁰ The approach is built around identifying dangers in the context.¹¹

The space agency NASA has an even more exhaustive and detailed framework for risk management. Its focus, however, is what might be called program or mission risk—dangers to the conduct of NASA’s planned space launches and related activities.¹² Like Air Force or Naval aviators, NASA employs concepts of risk to mitigate dangers to its flight operations. It conceives of risk management as an avenue to systems safety, and approaches it in part through quantitative and engineering-based methods.¹³ This is an important and useful approach, one that NASA has used to imbue a deeper mindset of appreciation for risk. But it represents a different approach, for a different sort of problem, than risk management in support of strategic decision-making.

Twin challenges: fragmentation and big bets

The current practice of risk management in national security appears to have two immediate potential pitfalls. The first is that, as a result of the multiple concepts, the Department of Defense no longer has a coherent concept of what it means by risk—if it ever did. The DoD uses risk to describe dangers in the overall strategic environment, the aforementioned gap between ends and means, endogenous problems with Defense institutions like the military services, and more. Enemy intentions pose risk; shortfalls in capabilities create risk; added resources would help “buy down risk.” All of those things are interesting to know, but if the specific issue of risk is to have analytical value, it must refer to something more particular. A RAND study argues that:

DoD planners are faced with two stark alternatives as they grapple with risk in force planning. At one extreme, *risk* is often used in an undefined way to justify resource reallocations, as in “We are cutting this program and will just have to accept more risk in this area.” The term *risk* is used so often in this casual manner that it is beginning to lose any meaning. At the other extreme are the highly quantitative methods associated with formal risk assessment and management. The latter are framed to address highly bounded problems (e.g., the threat, vulnerability, and consequences associated with a specific threat to a specific facility). The data required to apply their precise algorithms are simply not available to defense planners working at the strategic and operational levels of warfare. Yet DoD and the services have embraced some methods that would suggest that objective measures of probability and harm exist.¹⁴

A good example is the 2010 QDR. Many types of risk float around in the document, whose various meanings are never resolved. The danger of such a kaleidoscope of risk approaches is obvious enough: If risk comes to mean a multitude of different things, it will be less likely to offer a coherent, disciplined advantage for strategy.

A second challenge is that these risk frameworks are mostly oriented toward operational and program-level risk. The application of risk to strategic-level choices is much more episodic and idiosyncratic. Senior leaders definitely think in terms of outcome-oriented risk when dealing with big strategic bets, but as we will see there are a whole range of human factors than tend to get in the way of rigorous assessment of such dangers. But because there is such a proliferation of lower-level risk processes in place, there is a tendency to think that the national security establishment “does risk.” In many ways, it does—but much less so for the most important choices the nation makes.

Risk as threat

Sometimes, as we have seen, national security documents or studies refer to risk primarily as threats in the surrounding environment that could cause trouble for US interests. This approach emerges in comments from senior leaders who refer to the “risk” posed by peer competitors or terrorist groups, or in formal institutional analyses that examine a wide range of things that could go wrong. But this concept of risk ends up duplicating other steps in a typical strategy process and should not be conceived as risk analyses at all, except to the degree that the dangers they outline become specific potential intervening events threatening the implementation of policy or strategy. Moreover, this method potentially distracts attention from the most important contribution risk management can make in avoiding strategic disasters, which is the deep consideration of consequences.

Defining risk as threat

Thinking of risk-as-threat, as is common in many DoD approaches and in the risk framework of the Department of Homeland Security, focuses on general dangers facing an organization, whether resulting from broad trends in the situation (such as climate change) or the malign intentions of other actors (whether business competitors or national rivals or enemies). The approach embraces the closely related concept of vulnerability—it is vulnerabilities, in the dictionary sense of an exposure or flaw that makes an institution subject to danger, that create the opportunity for threats to do harm. A vulnerability could be the absence of port security, for example, while the threat is the potential for terrorist attacks.

In the process, it is important to distinguish between dangers that are simply “out there,” trends or realities in the general context (including other

actors), and perils that apply specifically to a strategic action. One is existential and passive, the other reactive and tied to behavior. The danger of contagious disease or Russian ambitions in the Baltics are *threats*; sending US peacekeepers into a violent country poses *risks* to the deployment and to US strategy. The line becomes blurred when we consider the risks to an existing strategy, which can look like threats: US troop deployments in Korea carry the risk of embroiling the United States in a war, which is perhaps another way of understanding the threat of North Korean aggression.

Some approaches decline to make this distinction and treat risk and threat as essentially synonymous. The Oxford Dictionary defines a *threat* as a “person or thing likely to cause damage or danger,” or a “stated intention” to cause harm. And the dictionary defines *risk* as the “possibility that something unpleasant will happen,” which could be viewed as merely a subset of threat. This indeed is a common way in which these terms are used. All too often, especially in the national security realm, they are seen as one and the same: Many analyses refer to the “risk” of a North Korean invasion, or a terrorist attack, or Russian occupation of the Ukraine—meaning the threat of such malign outcomes.

The limitations of risk as threat

There is probably no way to entirely depart from this commonplace use of the term “risk.” And there is no question that assessments of threats and dangers in the strategic environment is a crucial element of strategy. In order to fully inform leaders about the environment, a strategic process must survey leading threats and dangers.

Yet equating risk with threat raises a number of problems. The approach brings nothing new to the overall strategic process. The basic threat-assessment stage is a part of any good strategy. Most organizations have elaborate processes in place to survey the strategic or competitive environment and inform senior leaders about the basic trends, including threats and dangers. Rebranding this approach “risk” does not add anything new to the analysis, and as we have seen it does not target the most common sources of risk disasters.

Second, the very nature of threat analysis in national security processes almost guarantees that using risk as threat will skew, and often exaggerate, the perils involved. The biggest problem in strategy is not that strategists ignore the surrounding context. In fact threat analysis is, if anything, massively over-determined in national security processes: Security professionals typically see a wide range of potential threats to the nation’s well-being. Recent national security documents are a powerful reminder of the fact; senior national security leaders have spent much of the last several years arguing that the world is more dangerous.

A third problem with threat-based risk analysis is that it can be very unclear what exactly something threatens. Lack of precision encourages misunderstanding or vagueness in the analysis of threats. As an example, many might

agree that Chinese fishing-vessel provocations in the South China Sea carry some danger. They might be termed a “risk.” But what exactly is the target of this risk? Regional stability? The security of peace? The interests of US allies? US credibility in the region? Equating risk with threat can encourage a proliferation of threats under the guise of being “comprehensive” in the assessment of risk. Risk management ought by nature to have a degree of overprotectiveness; threat assessments should be more balanced. The advantage of an outcome-oriented conception, as will be discussed below, is that it forces a specific application of potential dangers to specific actions.

Fourth and finally, this approach distracts attention from the much more important gap in strategy-making. A major insight from the lessons of risk management in the financial crisis—and from parallel cases in national security—is that the bigger problem is in taking seriously potential consequences and second-order effects of actions. National security processes tend to obsess over threats and pass quickly over the potential consequences of actions. This is where risk can play a much more important role, and focusing on threats places the emphasis of risk management at the wrong place in the strategy-making spectrum.

A default approach: ends-means gaps

A second and even more common way of thinking about risk in national security is as an unfilled gap between what defense planners think they need to do and the resources they have to do them. This is the notion of risk-as-insufficiency or risk-as-infeasibility of plans. It is generally employed in reference to specific requirements that are developed, for particular scenarios or general Joint capabilities, through war plans and modeling.

As a practical example, consider the war plans that are maintained by regional commands—Central Command (CENTCOM), European Command (EUCOM), and so on. The Defense Department tasks them to develop a plan to win a war against some potential regional adversary. Their “planners” will develop the concept and lines of operations that will win the war, and from those generate requirements—lists of forces, resources, supplies, and so on that they’ll need. And they kick these requirements up the chain to the Joint Staff and Defense Department civilians, who are managing similar requests from a host of commands and other entities. Often the national inventory won’t cover all the demands. In a hypothetical example, the sum total of requests for Air Force squadrons might be 50, and the United States only has 40 of them—the gap of 10 will be called “risk” and placed onto a Powerpoint chart.

Specific military services or Joint Staff offices also employ operational risk to measure the gap between requirements and means. A service will take the requirements generated by operational planning demands—say, a given number of Brigade Combat Teams (BCTs) or squadrons. It will compare them

(either for a specific scenario or for the whole force) to its capacity, and any gaps will become “risk.” In more specific terms, risk for particular capability areas can be measured by specific capacities: For example, having a reconnaissance capability at certain levels will be assessed to require a given capacity; shortfalls from that capacity will be outlined as “risk.”

This sufficiency gap analysis is probably the dominant approach to risk today. It is the primary spirit of the Chairman’s Risk Assessment. The main body of the 2014 Quadrennial Defense Review (QDR) examines risk in much the same way, as a function of declining budgets and inadequate resources.¹⁵ The Army’s Training and Doctrine Command developed a new Army Operating Concept in 2014. Its section on “Risk and Mitigation” is all about shortfalls that could obstruct the strategy—potential gaps between objectives and resources, such as “insufficient funding and inadequate capacity” or “insufficient strategic lift.”¹⁶ When the former Commandant of the Marine Corps testified about the risks of a smaller Corps in relation to its missions, he referred to the outcome as a “moderate risk” force.¹⁷ Former Army Chief of Staff General Raymond Odierno testified that the reduced readiness rates of Army units left them unable to fulfill requirements and “translates into strategic risk.”¹⁸

This mechanism for understanding risk relies on the idea that predictable relationships can be discovered between particular capabilities and capacities and certain outcomes. But there are arguably precious few linear relationships between inputs and outputs in war. Too much uncertainty creeps into various stages of the process: how the enemy will behave, how well units will fight, random effects like weather, the leadership qualities of key officers, and far more. Making a reliable risk assessment based on such relationships is to invite error and misunderstanding of the foundation for the final numbers.

Risk assessments will spell out reasons why the United States needs a certain number of brigades, or aircraft carriers, or airplanes to avoid “failure” in a given contingency, for example, but these claims will be based on elaborate assumptions that don’t always stand up to scrutiny. When a regional command such as CENTCOM declares that it “needs” 12 or 14 or 18 brigade combat teams (BCTs) to successfully prosecute a war plan, the number has perceived objectivity—but in fact any such estimates are the highly contingent products of very rough calculations that sometimes amount to little more than very well-structured guesswork. And when the war actually occurs, it turns out that the commands can often win with far less than they thought they needed. A given threshold of capability, for example, might not imply failure; it might mean only that the force would fail to meet a self-imposed timetable for winning. The United States would still likely win, just more slowly and perhaps at greater cost. It can be misleading to define and assess risk strictly as gaps in the ability to perform established and narrowly chosen war plans or operational concepts.

The definition of those concepts is also a function of strategy; there is no objective measure of requirements, but only in relation to the strategic concept being used. What you need depends in part on *how* you propose to fight. Fighting a given campaign through attrition warfare may generate one set of requirements, whereas a guerrilla war campaign of indirection would generate a very different set. Different ways of conducting a war or campaign can have radically distinct implications for needed capabilities, and yet ends-means based risk assessments typically assume a single, standard concept and can end up misleading about the level of risk involved. An additional problem is that the choice of those stressing scenarios to guide risk is usually a function of expectations about the future more than any objective set of criteria for how or why they should be used as the benchmarks of risk.

In sum, then, thinking about risk as the gap between means and ends suffers from many potential shortcomings. Dozens of variables help determine a level of sufficiency relative to interests, and most risk processes can only capture a handful of them. “Risk is about the consequences of policy and resource decisions,” one source argues:

Many leaders and their staffs speak of “taking risk” and “accepting resource shortage” as being synonymous. Although it may be true in some instances that a high-level decision to remove resources from an Air Force activity increases risk to the Air Force and nation, this is not necessarily the case. It also could be that little or no increased risk is associated with the resource decision. Even if risk is increased, the decision may be necessary to free resources to address other priorities within the USAF or elsewhere in DoD. The broader issue is “What is the net effect on risk to the nation?” It is not helpful—that is, it is not informative in the public policy debate—simply to speak of “risk” when Air Force resources fall. Risk of what? What are the policy consequences when such resources fall? Is national security damaged or compromised when this occurs? In what way? Under what circumstances? How likely is the damage to occur if the resource reallocation persists? If the damage in fact occurs, how severe will it be? And compared with what?¹⁹

There may also be bias in the foundation of requirements on which such assessments are based. Such requirements are typically generated by the combatant commands who are asking for forces. These headquarters have a natural incentive to inflate the threat, and requirements: All organizations tend to ask for more than they “really” think they need in such resource games, in part because they know they’ll get less than they ask for.

Finally and perhaps most importantly, thinking of risk as the relationship between ends and means encourages a misunderstanding of what it takes to mitigate risk. This approach implies that, if a state matches its stated needs

with its capabilities, it has driven risk to zero. That is not, of course, the case, and we need to look beyond the relationship between ends and means to appreciate the full scope of risk. Planners for the Iraq war in 2003 may have judged the operational risk to be low based on the capabilities they had on hand, but that did not reflect a true, and comprehensive, assessment of the “risk” involved in that calamitous leap of faith.

This is not to suggest that gaps between ends and means pose no danger. At some point and in some ways, they will: Any time a strategic actor undertakes a mission with far fewer resources than its military estimates suggests that it needs, it courts failure. What this is really about is what could be called sufficiency analysis—does an institution have what it needs to accomplish a given mission. Such analyses are important and can and should be part of any strategic analysis. They should be tied closely to the statement of the strategy or strategic concept being used, because as we have seen only by linking ends and concept can a process derive meaningful requirements, and then outline capabilities needed to achieve the goal. But this line of thinking is not, fundamentally, about risk—it is about the feasibility of a proposed strategy with the resources available.

The mechanisms of risk: quantification and stoplight charts

Another typical characteristic of current DoD risk management processes is that they seek data-driven, quantified outputs, and they are often ultimately expressed in some version of a stoplight chart (or the equivalent, a red-yellow-green spectrum of risk). Both of these approaches carry the danger of creating misleading risk analyses.

This study, again, is examining the specific challenge of risk management under uncertainty, not in some of the more programmatic areas mentioned above. Those cases—for example, assessing risk in Naval Aviation—can make more productive use of data-driven risk frameworks. Under uncertainty, as we will see in more detail later, trying to reduce key variables to numbers can mislead.

A temptation to false precision

There can be few better examples than the terribly misleading system chart created to express key variables in the Afghan war.²⁰ It was covered with factors, ranging from “territory not under government control” to “ANSF execution and capacity” to “potential attractiveness of government versus insurgent path.” These are, of course, highly qualitative considerations: The “potential attractiveness of the government” is a nuanced issue with many sub-variables of its own, something that could occupy long hours of philosophical debate. It is, moreover, an inherently subjective variable, entirely dependent on the perceptions of the viewer: Some Afghans could (and of course did) view

exactly the same situation (for example in a specific sub-district) as reflecting very different degrees of government attraction. Many of the factors on the chart had the same quality—they were dependent on one's perspective (which, as we will see in Chapter 5, is one of the leading factors tending to skew risk perception).

And yet ultimately the authors of this chart—which was really a model—assigned a value from zero to one for each of the factors, assigned relationship values among the pairings, and ran the model to produce an overall answer. One can easily see how bankrupt such a process can become: The values assigned have huge arbitrariness built in (even if they are grounded in some objective information), which then becomes compounded by the interactions. And hanging over the whole process is the crushing weight of uncertainty: The architects of this model cannot know which variables will end up being decisive, or why. The result is an answer that is essentially meaningless.

False precision—sometimes called overprecision or fake precision—is simply the depiction of quantitative data in a manner that implies more accuracy is involved than is actually the case. A common joke in this regard involves a museum employee who insists that a fossil is “one million and three years old,” because it was said to be a million years old when he started his job three years before. But the same principle can emerge in apparently complex and sophisticated risk models. If the underlying data are based on subjective or estimated values, then the final products of a model or calculation will usually involve far more arbitrariness than the risk presentation will imply.

Much the same result can be expected from risk frameworks that attempt to build toward quantified answers for complex strategic judgments. An example might be the sort of risk assessment for deterring Russian intervention in the Baltics outlined in the thought experiment earlier in the chapter. A framework might generate an answer of “high” risk for the deterrent mission based on a range of quantified variables: ratio of forces on the ground, for example, or numbers of Russian threats in the previous month. The result will incorporate many interesting and useful pieces of data, some of which, examined explicitly and on their own, could provide important insight. But the sum total resulting picture has a major chance of being substantially mistaken. How Russia views certain numerical realities, for example, will be dependent on many uncertain factors. Ultimately, trying to assign a numerical value to a qualitative reality poses serious dangers.

Such effects can be especially pronounced in aggregated risk models, such as the Afghan chart. It is almost inevitable in a situation in which many variables are summed into intermediate ones, and then several intermediate factors are joined into a top-level risk estimate for some category. If each variable involves a wide margin of error and is based in part on subjective estimates (or reflects an objective value for something that is really a proxy for the intended factor), the resulting top-level number can be almost meaningless. To take

just one example, the Afghan model made highly subjective assessments of a number of qualitative factors—the degree of governance, public perceptions of the government—or else it used indirect proxies, like miles of new road construction as a stand-in for governance. In every such example, a little more arbitrariness and inaccuracy was introduced into the model, but it (and others like it) portray their results in seemingly objective terms. There is likely a tipping point in such situations at which the accumulating inaccuracy totally undermines the reliability of the assessment, but this is seldom estimated or kept in mind as the values are used to inform choices.

The perils of stoplight charts

The same can be said for summing up an organizational risk picture in a so-called “stoplight chart,” one of the most common mechanisms for portraying risk in DoD. Typical capability-gap risk assessments become even more ambiguous and potentially misleading when they roll up large numbers of such subjective judgments into a universal number (or color coding) which conceals as much as it clarifies. For one thing, the process usually chooses an often-arbitrary level to count as “failure,” often a generic percentage not grounded in any clear criteria. The risk briefing will ultimately present a single color-coded judgment—“moderate” risk, or “severe”—which is the product of many subordinate judgments whose relation to one another is not always clear. The models need to make assumptions about the overall effect on the judgment of individual categories, and these are as often as not based again on subjective judgment.

The question in such presentations is what is going on underneath or behind the colors—what judgments and assumptions have been made to get to that final, simplified, often significantly arbitrary presentation. Such underlying assumptions and judgments are seldom made explicit in risk presentations, and yet they often reflect the most important issues at stake, and the ones that demand the most urgent deliberation by senior leaders. Whether it is something as operationally basic as the exchange ratios between aircraft in a conflict, or something as complex as whether “good governance” should be 20 or 50 or 80 percent of the effect of a given variable, these decisions determine risk outcomes. They will tend to reveal the most about the underlying dynamics and thus go the furthest to inform complex strategic judgments. They should therefore be the heart and soul of any risk dialogue.

For example, a hypothetical Air Force analysis might suggest a “severe” level of risk in a given scenario. This might actually mean, in a given case, that the vast majority of risk levels remained low, and one or two key categories drove the overall assessment in a negative direction. But whether the Air Force could mitigate those individual capabilities with some fairly easy steps is not clear from the overall assessment. Rarely are things so black-and-white. Even with many shortfalls identified by a risk analysis, a nation or institution could

still win, or succeed—just at greater cost, or over a longer time, or in different ways. In this sense the message of ends-means analyses dressed up as risk can be profoundly misleading: Instead of focusing on the real risks, it drives attention onto arbitrarily chosen metrics of sufficiency whose absence creates “failure” according to the model.

These portrayals of risk suffer from the same flaws as most quantitative models under uncertainty. For example, they tend to conceal dozens of assumptions: In order to determine whether a given capability or quality is at “yellow” or “green,” for example, will demand multiple judgments that are ultimately subjective. Yet the model implies a degree of objective analysis that does not exist.

Such spotlight charts can also have the effect of shutting down dialogue and closing off awareness rather than the other way around, especially when combined with some of the cognitive flaws that will be outlined in the coming chapters. Combined with elaborate briefings, they can allow senior leaders to passively embrace risk judgments made on their behalf; this certainly happened a great deal during the financial crisis. This does not have to be the case—some leaders clearly use such presentations as an invitation to a highly detailed discussion that questions the underlying assumptions. But the default message of these approaches is the opposite—they offer seemingly objective truths rather than asking questions.

This chapter has argued that there are a number of problems with many currently popular approaches to risk. The question we are left with, then, is what approach to risk would offer a more decisive support to strategy. In order to inform that choice, this study now turns to the specific lessons of the recent financial crisis—in broad terms, the role of human factors in undermining procedural risk management.

Part II
Lessons of the Crisis—
The Character of Risk

4

Risk and Uncertainty

A decade before the financial crisis of 2007–2008, the world financial system was rocked by a more limited but still revealing event: the collapse of the hedge fund Long-Term Capital Management. That trauma sent shock waves through global finance and hinted at many of the dangers residing in the over-leveraged, highly complex investment instruments that would later threaten the global economy. One of the most important lessons of the Long Term case had to do with the danger of placing inordinate faith in quantitative models that promised to tame the unruly complexity of markets.

Long-Term Capital Management (LTCM) had a seemingly faultless pedigree. Its founding chief was John Meriwether, who had amassed a brilliant investment record as head of bond trading at Salomon Brothers. He built the new fund on the research of two prominent economists, Myron Scholes and Robert Merton. Their work had derived a complex formula for estimating the value of derivatives, which won them the Nobel Prize in economics in 1997. After racking up returns of over 40 percent in its second and third years, the fund came crashing down amid the 1997 Asian financial crisis and Russia's economic calamity the following year. On September 23, 1998, a coalition of major banks was forced to come together to inject cash into the listing ship of LTCM in order to safeguard the global financial system.

Once the fund had collapsed and nearly brought down a number of other financial institutions with it, it became clear that the LTCM approach suffered from notable methodological flaws. A number of them will feature in coming chapters. But the leading source of failure was a single factor: the belief that the firm had transmuted an uncertain environment into a deterministic one. "If you can make risk disappear—poof!—in a quantitative sleight of hand," Scott Patterson has written of the LTCM approach, then "you can layer on even more leverage without looking like a reckless gambler."¹

These essential elements of the LTCM disaster would reappear in the financial crisis of 2007–2008. Once again, brilliant financial minds would develop

seemingly fool-proof models for forecasting markets and anticipating risk, and using those models to substitute for judgment. Once again they would be ruined by arguably the central reality suggested by the crisis: the dominant role of uncertainty in shaping the context for strategic risk management.

This chapter lays the groundwork for the remainder of the study by discussing the foundational principle of complex strategic judgments—a context of deep uncertainty. The failure to acknowledge the significance of this context has been a major source of error in risk management, and it is essential to understanding the character of risk in national security. After introducing the basic problem, the chapter will survey the character of deterministic contexts and then define several essential elements of an uncertain environment.

It is important to stress, again, that this analysis accepts a perfectly appropriate place for quantitative approaches and complex models, even under uncertainty. The models at LTCM were brilliant and, for a time, entirely effective. The firm racked up massive profits for an extended period. The problem wasn't that their models were "wrong," or that they did not provide any window onto markets. A proximate issue was that the models did not account for large swings in the market which, by the models' own probabilistic estimates, ought to be essentially impossible—but which had occurred in the markets just a few years before. More broadly, that blind spot was a symptom of the larger problem: using quantitative models as a replacement for subjective judgment.

Understanding uncertainty

The problem for risk management, and ultimately strategy, is not that models built on probabilities and equations always fail. They are perfectly appropriate to some questions—if Starbucks is analyzing the relationships among purchases of different items, for example, or if an actuary wants to estimate the risk of a car accident in Midtown Manhattan for 45-year-old single men driving Tesla roadsters. Models can even provide somewhat accurate forecasts for highly discrete issues on a very specific time horizon within a larger context of uncertainty. Those are cases, among other things, where large data sets of roughly comparable events will be available. But probabilistic models do not work nearly as well, and are often dangerously misleading, as the dominant guide to choices in complex situations involving high degrees of ambiguity of information, the nonlinear interaction of a tremendous number of variables, and the constant churn of human choice and chance.

Uncertainty as I will employ the term here has three basic aspects.² First, uncertain contexts have a high level of *ambiguity*: Decision-makers lack key information they would need to make the sort of calculations expected by rational analytic methods. In uncertain situations they will never have the perfect information they would need to make conclusive judgments. This is

sometimes called “epistemic” uncertainty because it relates to our degree of knowledge about the world.³

One of the most fascinating examples of the problems of ambiguity for judgment can be found in the so-called Ellsberg Paradox. Before he became famous for challenging the official narrative on the Vietnam War by releasing the Pentagon Papers, Daniel Ellsberg was a brilliant social scientist, and he produced some very interesting and original work on judgment and decision-making. One case Ellsberg used was the following: You confront two containers, each filled with red and black balls. You are told to reach into a bin and grab a ball, and asked to make a bet on the likelihood of drawing a red one. Both bins contain 100 balls; in one, you know the ratio is 50–50, but in the other you have no idea. You’re told you must draw a red ball to win \$1,000— which bin do you choose?

With this invented situation, Ellsberg created a situation in which probabilistic analysis won’t help. Because you have zero knowledge of one of the containers, “it is *impossible*,” Ellsberg concludes, “to infer probabilities from your choices.” Players in his game are left having to guess in regard to a fundamentally ambiguous situation. When subjected to this test, most people do not rely on any strict principle of probability—they tend to decide based on intuitive preference rather than objective standards.⁴

The greatest challenges for judgment emerge in cases such as the Ellsberg Paradox, when dealing with severe uncertainty and ambiguity. This fundamental situation, the need to infer meaning in a fundamentally ambiguous situation, is the author of many cognitive biases and the reason for many heuristics we use to make sense of reality. (Ellsberg himself drew one particular implication from his own case: Very often our gut judgment under ambiguity will be biased toward the option with more known risks.) This sort of a context could be described, as Nate Silver has done, “initial condition uncertainty”⁵: Because decision-makers do not fully grasp the nature of the starting point, projecting forward becomes speculation.

If the first defining characteristic of uncertainty involves the lack of sufficient information, the second draws on the insights of complexity theory. It depicts the world as a “complex adaptive system,” an emergent product of vast numbers of intersecting variables in a state of constant change and flux.⁶ There are too many random, unpredictable factors, too many billiard balls on this dynamic pool table, to have any sense of how their collisions will play out.

The actors or forces in such systems swirl and interact, and patterns or behavior emerge in a sort of creative chaos that can only be partly understood, let alone forecast.⁷ In nonlinear systems, inputs do not necessarily equal outputs, and tiny changes in the starting variables can, through these swarms of interactions, make a massive difference for the outcome. The causalities involved are so dispersed, complicated and sometimes masked that causal relationships can only be inferred or guessed at.⁸ A good example is the

weather, whose vast number of interacting variables defies prediction past a few days. As is sometimes said, such complex systems are more “cloudlike” than “clocklike,”⁹ billowing into unpredictable patterns rather than following linear, deterministic rules.

Such a context carries some fairly obvious implications for risk management. For one thing, environments characterized by nonlinearity will not produce reliable patterns. “The ‘system,’” as the RAND Corporation scholar Paul Davis has pointed out, “is not a constant for which one can prepare straightforwardly. Rather, it includes human beings and organizations that think, behave, and adapt to events in myriad ways.”¹⁰ Emergent, unpredictable patterns and behavior will be the norm, and anticipating risks, or understanding the likely outcomes of our steps to mitigate them, becomes far more difficult.

An interesting example of this recognition from the period of the financial crisis is the Federal Open Market Committee. One study found that, in its economic forecasts between 1992 and 1999, the FOMC was wrong 96 percent of the time. Alan Greenspan spoke to the problem at a 2003 meeting: When presented with some very elegant models for the behavior of elements of the economy, he downplayed the ability to model effectively over the long-term for such a dynamic system. “We start with a degree of uncertainty that is very high,” Greenspan admitted. “It is much higher than it is for those who take the data and put them into a model and do projections. Most modelers are dealing with a controlled environment in which the number of variables is well short of a thousand. In the real world there are a million, and we don’t know which ones are important.”¹¹ He made this statement four years before the financial crisis struck.

A third major aspect of uncertainty is related to the issue of dynamic effects, but more focused on human choice and distinct shifts in the environment. It holds that a context of uncertainty is “transmutable,” meaning that it “undergoes incessant change as the consequence of human agency.”¹² People interact based on mutual expectations of choice and agency, which introduces unpredictable dynamics beyond simply lack of information or the structural principles of complex systems. The economist G. L. S. Shackle calls these “crucial experiments,” by which he means cases in which “the very act of performing the experiment may destroy for ever the circumstances in which it was performed.” He gives the example of individual chess moves, which “inevitably change the whole course of relevant future events for the individual.”¹³

In sum, uncertainty relates to a condition in which neither the present situation (initial conditions) nor potential future states (scenarios) can be known with any precision. In such a context, there will never be any reliably objective basis for a decision. Judgments under uncertainty can and should be *informed* by various objective measures. But the ultimate choice will necessarily be grounded in an intuitive, common-sense, experience-grounded sense of the “rightness” of one or another course.

Risk, probability and uncertainty

Fundamental to the whole strategy of LTCM was the fact that its analysts believed that they had solved the problem of probability. They had found the beauty of predictability within the larger chaos of markets, they believed, the pool of deterministic order in the larger ocean of complex uncertainty of the financial markets. Their letters to investors reflected this ambition, claiming that they could estimate the odds of failure down to the fraction of a percentage point. “Long-Term,” Roger Lowenstein has written in his classic account of the collapse, “was an experiment in managing risk by the numbers.”¹⁴

Various assumptions must be built into such models if they are to provide reliable forecasts. One is some version of regression to the mean—that markets constitute a long-term comparable series of events that betray standard deviation. Specifically, the LTCM models assumed that a market collapse big enough to threaten their firm would be what is known as a ten-sigma event, a one-in-ten-to-the-twenty-fourth-power random accident. Yet an obvious problem for such thinking should have been that collapses of precisely such magnitude had occurred in the not-distant past, in 1987 and 1992. LTCM, however, was using statistical models whose reach back into history stopped just short of 1992.¹⁵

Another baseline assumption of the LTCM approach was the “efficient market hypothesis,” the idea that investors have sufficient information and that stock values at any time reflect an objectively “correct” or valid level. This, too, was widely undermined by experience even by the time of LTCM’s founding—or perhaps better put, the extent and manner in which it was relied upon proved disastrous.¹⁶ Models making the assumption of an efficient market may be able to amass a record of reasonable accuracy on some very discrete and narrow areas of focus; likely price swings in a specific commodity given a welter of specific background conditions, for example. But when the strategy becomes vulnerable to swings in the market as a whole—when macro uncertainty can destroy accurate measurement of micro-determinism—then even precisely targeted models can serve up calamities.

In deterministic environments, issues behave according to certain rules: There are broad, predictable patterns in the data. Swings tend to self-correct because the system regresses to the mean—it is stable in the technical sense, governed by dynamics that pull it away from extreme fluctuations back to a typical pattern. In such circumstances, history will generally offer a reliable guide to future events.

These criteria do apply to certain economic phenomena at given times—rates of car crashes, for example, or demographic patterns. But markets as a whole, and other large-scale, complex, nonlinear fields driven by human agency and group effects like herding, diverge from expected patterns, sometimes refuse to return to the mean (at least very quickly), and embody wild

swings that seem unimaginably unlikely from the standpoint of pure statistical analysis. The evidence for this was already ample in 1998, but LTCM chose to ignore it, or at least to place little emphasis on the crucial distinction between discrete and general assumptions of determinism. Their counterparts in the financial industry would do the same thing a decade later.¹⁷

The economist Benoit Mandelbrot, who has used complexity theory and an appreciation for uncertainty to hammer away at traditional economic models that assume deterministic contexts, was making this argument long before the crisis. “The seemingly improbable happens all the time in financial markets,” he points out. Just in a handful of cases from 1987 through 2002, markets rose and fell in spikes that were, according to probabilities derived from standard distribution, essentially impossible: one in 50 billion, one in four trillion, one in 10 to the *fiftieth* power. These are “odds so small they have no meaning.”¹⁸

After the LTCM debacle, it might have seemed obvious enough that data-driven models trying to force probabilistic patterns onto nonlinear, mania-infused markets could be dangerous and misleading. Instead, such models proliferated in the wake of LTCM. Hundreds of new hedge funds appeared using variants of quantitative modeling. And the Fed chief of these years, Alan Greenspan, while apparently somewhat skeptical of radical quantitative models, nonetheless remained fixed—despite the experiences of LTCM and others—on the basic ideology of efficient markets. There was a fundamental overconfidence involved on the part of leaders who believed they had tamed complexity.¹⁹ But events would bear out the critique of Nassim Nicholas Taleb, who described the Black-Scholes equation as trying to be “mathematical and airtight rather than focus on fitness to reality.”²⁰

A leading lesson of the crisis, then, and the essential foundation for the analysis to come, is that risk management and nonlinear, volatile, fundamentally uncertain contexts make an uncomfortable pair. In her brilliant study of the Challenger launch, Diane Vaughan argued that “Risk assessments at NASA went on under circumstances that made risk fundamentally incalculable: large, complex organizations regularly and systematically fail because their parts interact in unpredictable ways that defy planning; large-scale technical systems cannot be tested to environmental conditions, because all environmental conditions cannot be foreseen.”²¹ This was true, in that case, in large measure for the same basic reason that cropped up in the financial crisis—the way disruptive human factors arose to cloud judgment.

The situation is much the same (in fact arguably even worse) when we consider strategic choices in national security. The variables are even more numerous, the interdependencies more intense, the lack of an ability to test options under real-world conditions even more obvious. Risk will be “fundamentally incalculable” in such circumstances, despite the recurring efforts of defense institutions—such as the systems analysts of the Vietnam War years—to impose rationalistic models on complex realities. Joshua Cooper Ramo

suggests that such characteristics are integral to the global strategic environment. In *The Age of the Unthinkable*, he maintains that the world isn't a "math problem," but rather suffused with "contingency and irrationality." Ramo's argument is that the "complex physics" of the emerging interdependent world create more nonlinearity than ever before. He describes what he calls the "sandpile effect" in which a few additional grains added to a pile can create transformative change—a system "constantly poised on the edge of unpredictable change."²²

The baseline: probabilistic challenges

Risk is a probabilistic concept, at least as it has evolved over the last several centuries. The basic concept of risk, at least one understanding of it, has to do with something measurable that can go wrong.²³

The economist Frank Knight offered the paradigmatic version of this distinction. Risk, for Knight, involves inherently measurable possibilities—situations in which a large number of cases has things in common, characteristics which are "both few in number and important enough largely to dominate the situations." The shared aspects of those few characteristics "enable us to reach an approximation to the law of the situation as a whole," and to identify "tendencies" of what ought to happen.²⁴ As the political scientist Stephen Nelson has put it, "In environments characterized by risk a decision maker can attach a probability distribution to the range of plausible outcomes resulting from a decision."²⁵

This perspective is inherently positivistic. Positivism, the idea that every claim can be rationally proved or disproved with objective data, is a child of the Enlightenment, which helps to explain its essential claims—it is an intellectually imperialistic philosophy, out to conquer and subdue reality. It embodies an implicit faith in the ability to find truth grounded in "objective, controllable, and universally verifiable experience."²⁶ The story of risk management, in this sense, is a sub-set of the modern era's larger effort to transcend fate and take control of events.²⁷

A critical background condition for positivism is the idea that all the information necessary to mastering reality is available—the question is just how hard it may be to get. The future, in this view, is resident in current facts, if a decision-maker could only know them. In 1814, the mathematician Pierre Simon de Laplace put the basic concept very neatly: "If an intelligence, at a given instant, knew all the forces that animate nature and the position of each constituent being," he wrote, and if it could analyze the data adequately, then "to this intelligence nothing would be uncertain, and the future, as the past, would be present to its eyes."²⁸ This assumes a deterministic world, in which the variables that produce outcomes are theoretically discoverable, and causalities can, with enough study, become clear. It is a world in which "our

personal qualities and the properties of any given situation or environment lead directly and unequivocally to precise consequences.”²⁹ The financial crisis was driven in part by derivatives that were designed to have this effect—to create a perfectly efficient market, based on probabilistic analysis, that could anticipate futures with astonishing clarity.³⁰

Probabilistic models work very well in the context of some problems, including those governed by physical laws, or the pursuit of simple efficiencies of allocations when most variables can be controlled. When determining how to patrol an airport, or how to search for drug smugglers at sea, an algorithm can help suggest search patterns that optimize resources. Another example is actuarial science, the forerunner in some ways of risk management: Insurance actuaries can use millions of examples, and the expectation that future trends will generally follow established patterns, to build fairly reliable quantitative models.³¹ Even under uncertainty, probabilistic models can be deployed to offer a perspective on very specific slices of the overall issue. But such models cannot fundamentally resolve complex strategic judgments, which deal with high levels of uncertainty and emergent patterns that make deterministic models unreliable and dangerous.

If *objective* probabilities cannot be discovered for various outcomes under uncertainty, decision-makers, as we saw in Chapter 2, can still assign *subjective* probabilities—the probability we anticipate through our individual perspective. Probability in this case becomes a measure of belief, not objective fact; subjective probability becomes what any “rational” person could assign to a given situation.³² This allows some degree of objectivity to the degree that rational people will properly calculate certain aspects of a situation.

The problem is that rational people do *not* assign common or reliable probabilities to events. A theory of probability cannot be subjective and collective at the same time. Once we open the door to the idiosyncratic assignment of probabilities, all manner of cognitive biases flood in—overconfidence, anchoring, wishful thinking. Decision-makers’ worldviews and ideologies will skew their perception of risk. Since judgment, including risk assessment, under uncertainty is (as we will see) a fundamentally imaginative and creative enterprise, so will be the subjective assignment of probability.

The distinction between deterministic and uncertain contexts is Nassim Taleb’s fundamental argument in *The Black Swan*. He makes a fascinating distinction between what he calls “Mediocristan” and “Extremistan” to illustrate the basic difference. Mediocristan is his image of a place where normal distributions prevail, and any outliers fall within certain parameters. More importantly, the number of events is so large that any one outlier won’t dramatically change the outcome. Examples could be height, or weight: The distribution falls within set parameters—no human weighs a thousand tons. Mediocristan includes such issues as height, weight, car accidents, mortality rates, and IQ. It is the realm of risk, in Knight’s distinction—the realm of probabilistic analysis.

Extremistan, on the other hand, is rife with massive lurches away from normal distribution. An example is net worth—one person might make a million times as much as another. The whole idea of an average becomes questionable in Extremistan, because calculations “can depend so much on one single observation.” In this realm we find such issues as wealth, book sales by author, the impact of natural disasters, company size—and, Taleb claims, financial markets. This is the realm of uncertainty, and Taleb later offers a parallel list of fields that allow true expertise and those that don’t, which is a function in part of their place in one of these two worlds. Experts are more common among test pilots, physicists, accountants, insurance analysts, and livestock judges; experts are absent among psychiatrists, college admissions officers, judges, intelligence analysts—and financial experts. As he explains, “Almost all social matters are from Extremistan.”³³

At some point it’s guesswork

One distinction that helps to clarify what I have in mind is the difference between epistemological and ontological uncertainty. Under epistemological uncertainty, the problem is not that the future is theoretically unknowable—it is that decision-makers cannot get enough information to figure it out. Ontological (or “aleatory”) uncertainty, on the other hand, is much more fundamental: It is the notion that the future does not yet exist, but is created by an impossibly dense set of interacting phenomena. Because it is constantly being created, the future is not resident in currently available information—no matter how much we have.³⁴

Scholars of risk sometimes explain such contexts by saying that the “ergodic” assumption doesn’t hold. This concept refers to a situation in which the future can be expected to be essentially like the present, and so data and relationships examined today can be taken as an analogue to future ones. But an uncertain system is fundamentally “nonergodic,” meaning that decision-makers cannot extrapolate outcomes, or even consistent cause and effect relationships, from the current situation. “Variables such as the price of cotton, employment in the healthcare industry, or the co-movement of mortgage default rates, cannot be reliably forecasted from existing information because the scheme of cause and effect that will define their future values *is yet to be created.*”³⁵ Nonergodic, ontologically uncertain environments are constantly in the process of being created. The most powerful source of this nonlinear unfolding in social contexts is human behavior, preferences and interactions.³⁶ As Frank Knight has put it, when “The decisive factors in the case are so largely on the inside of the person making the decisions,” then the case is “not amenable to objective description and external control.”³⁷

One of the defining characteristics of complex strategic judgment in such circumstances is that it involves the application of subjective, qualitative

values. Science, statistics, and risk assessment cannot offer final answers to such questions. The choice involves trade-offs among competing considerations and values that cannot be objectively resolved: Is a higher risk of conflict, for example, “worth” bolstered US credibility? Is severe damage to the US world reputation “worth” degrading terrorist groups by a certain amount? Numbers could be assigned to such goals and interests, but they would be arbitrary and ultimately subjective. Making such trade-offs calls for judgment, not probabilistic analysis. The way a decision-maker balances among competing variables will depend on the context,³⁸ and the way different decision-makers make the trade-off will differ. Preference is constructed by specific leaders in particular circumstances.

Risk, in this sense, is a collective construct, not an objective property or value. Risk is what we make of it, from the standpoint of social norms and values. Two scholars have gone so far as to argue that the test of an effective risk framework is not its objective validity but its ability to appeal to enough people to become generally accepted.³⁹

Characteristics of uncertain challenges

Complex strategic judgments therefore take place in a context of deep uncertainty. Such contexts have a number of important characteristics.⁴⁰

One is that *each case is fundamentally unique*. “There is no typical member”⁴¹ of any set in such a context. We can only understand an event in its own circumstances on its own terms.⁴² As Shackle suggests, observers can undertake probabilistic accounts when they can repeat experiments in “suitably uniform circumstances” that are “sufficiently numerous.” What he calls a “non-divisible experiment,” on the other hand, involves comparing a series that is neither uniform nor frequent enough for probabilistic conclusions. These are “serial” or “isolated” trials and experiments.⁴³ There is no “standard” example of a set; games of chance are a series of independent operations in which one action has no influence on future ones, and empirical verification as generally understood—as repeated trials of similar series—becomes problematic, if not impossible.⁴⁴

As Nate Silver has written, probabilistic analysis relies on a reliable, consistent data set.⁴⁵ Making general inferences from single events in an uncertain context becomes exceptionally dangerous. It becomes, in another guise, the problem of mistaken analogies, when one case of a given causality or chain of events is presumed to be a guidepost for others. Here we run across, once again, the disjunction between the general and the particular mentioned in Chapter 2.

A second and related characteristic of such contexts is that *the past is not a reliable guide to the future*. It is a basic principle of ontological uncertainty that futures cannot be read from present, or past, realities. In the financial industry

before the crisis, as quantitative models were employed to estimate risk, the distinction between past and future could be very sharp, and yet ignored by modelers who needed some sort of data. Roger Lowenstein gives an excellent example: Because mortgage policies had been so much more stringent in the years before 2001, broad-based mortgage bonds almost never defaulted. Using that data set as a measuring stick for the vastly more risky subprime loans being pushed out in later years made no sense, because the mortgage policies had changed—but the modelers had no alternative. So they guessed at the relative risk functions between two very different contexts.⁴⁶

Third, *causalities become difficult if not impossible to know for sure* under uncertainty.⁴⁷ As we have seen, understanding causal chains is critical to probabilistic risk analysis. And yet causal analysis is extremely chancy for political events, even in retrospect.⁴⁸ It is a fundamental characteristic of nonlinear, transmutable environments that causal relationships will be fickle and changeable.

Fourth, under uncertainty *judgment becomes an emergent process of imagining meaning in situations*. If an uncertain, nonlinear situation does not surrender obvious, objective meaning, decision-makers must read their own truth into the situation. Risk assessment becomes an individualized and personalized process. It becomes subjective, idiosyncratic, and intuitive. This will be investigated more fully in Chapter 5.

Fifth, as a result of the subjective character of interpretation under uncertainty, *perception, not reality, guides risk assessment*. When decision-makers are reading meaning into situation, it is their perspective, more than any objective facts, that determines what they see. The result is a system of interdependent perspectives highly sensitive to bias, rumor, and information cascades.⁴⁹ In an interconnected system, perception drives outcomes. Opinion can become fact faster than an organization can respond, and this constitutes a form of risk all its own.⁵⁰ Research has shown various ways that people and organizations attempt to cope with uncertainty in the context of strategy-making, and many of them demonstrate significant potential for cognitive bias, from wishful thinking to overconfidence to an illusion of control.⁵¹

A sixth and final characteristic of uncertain environments is that *warning of risks becomes more difficult*. Uncertainty obliterates the hard and fast evidence that a risk manager or dissenter could offer to bolster their case for an emerging danger. The timing and specificity of warnings are critical determinants of how well it is received. Warn too early, and a decision-maker can brush off the issue as a long-term problem not yet ripe for solution. Warn too late, and they will not believe they can do anything about the problem. Offer vague and ambiguous signals supporting the warning and the decision-maker will discount its significance.

Uncertainty complicates all these factors, making it very difficult to know when an issue is ripe and denying warning analysts sufficient detail and

specificity in the information to compose a persuasive message. A majority of warnings under uncertainty—as we’ll see in Chapter 7’s discussion of “gray swans”—appear to fall prey to this problem, with leaders asking for more detail before committing to mitigating actions. A common response from decision-makers focuses on the uncertain character of the information, and the central problem of not being able to anticipate outcomes. “I know this *could* happen,” they will say about a particular warning. “But so could a lot of things. Why should I alter my course of action on a mere possibility?”

Uncertainty and risk: lessons for national security

The first general lesson of the crisis for risk management in national security, then, involves the fundamental importance of uncertainty. Barry Schacter has argued that even in the wake of the crisis, “It seems likely that the new risk management framework will be defined by the old risk management paradigm”⁵²—that is to say, a linear, probabilistic belief that events could be tamed.⁵³

Again my argument is not that models are useless or always misleading, but that their specific role in risk management under uncertainty must be carefully understood and rigorously controlled. It is when they substitute for judgment—or, as we will see, help to rationalize excessive risk-taking—that they become dangerous. Scott Patterson cites a manifesto penned by two quantitative modelers in January 2009, men who believed strongly in the power of data but believed that its use had run out of control. “Physics, because of its astonishing success at predicting the future behavior of material objects from their present state, has inspired most financial modeling,” the experts argued. Physicists run repeated experiments to discover repeatable relationships and solid laws. But “It’s a different story with finance and economics, which are concerned with the mental world of monetary value.” As hard as finance tries to “emulate the style and elegance of physics,” the unavoidable truth is that “there are no fundamental laws in finance.”⁵⁴ Roger Lowenstein writes that when Merrill Lynch brought in a physicist to measure risk, he “tried to explain to his peers that the laws of Brownian motion didn’t truly describe finance.” The problem, Lowenstein concludes, was that “home-owners weren’t molecules.”⁵⁵

Even one of the most ardent proponents of Big Data, Nate Silver, makes very clear that the data-based approach only works for a very specific set of problems. In chess, for example, looking ahead at a game through 10 or so moves, the number of possible combinations of moves and countermoves becomes impossible to forecast. In fact Silver specifically calls out national security issues: In his own work on baseball, he writes, he had a sufficiently robust data set to work from. National security issues are much more case-specific, and the critical data are often hidden.⁵⁶

A number of the specific risk management tools in the financial world, such as Value at Risk (VaR), suffer from these flaws. VaR presumed probabilistic distribution and gave false confidence, making ultimately unreliable predictions about likely moves in the market. As one source concludes, “The VaR methodology makes sense as an efficient mechanism for handling risk only if the world of finance is one of risk rather than uncertainty.”⁵⁷ Value at Risk “can be gamed,” Joe Nocera found in a 2009 investigation—once managers knew it would be the key metric, they could grab investments specifically designed to artificially lower the firm’s VaR, while the bottomless risks inherent in some extremely risky correlated bets remained.⁵⁸

Simon Johnson and James Kwak have pointed out the explosive connection between models and human factors in their discussion of the role of Value at Risk approaches. Quantitative models are “shaped by the human beings who create and use them,” they point out, “and those human beings have their own incentives.” In particular, VaR models “brought concrete benefits to specific actors in the banking world by helping them to rationalize bad bets.” If “common sense” would cause a firm to call of a highly risky trade, then the trader—who hopes to profit and will not be liable for losses—“is better off if the risk manager uses VaR instead. ... In other words, models succeed because they meet the needs of real human beings.”⁵⁹ This is exactly the problem: The theoretical value of any objective risk process exists only at the discretion of the human factors that ultimately shape judgment.

Models like VaR, or any data-based risk assessment tool, can potentially be used with sufficient constraints to be helpful without being dangerous. The challenge becomes one of leadership and judgment—using tools and decision processes appropriate to the situation. The really explosive danger is the *intersection* of quantitative models and human factors. Ideally one should be able to counteract the other, but when both become tools working in the same direction you have problems. VaR and related tools had the effect of creating a false sense of security and allowing firms to exaggerate the value of the risk process.⁶⁰ Such critiques point again to the dangers of stoplight charts and similar simplified representations of supposedly objective, data-based risk: They will often mislead, suggesting far more certainty than is warranted and often concealing dozens of key assumptions.

A critical lesson of uncertainty and the experience of the crisis is therefore that risk processes must be designed to generate an informed conversation among senior leaders in which they examine risks in a direct and qualitative sense. This has significant implications for how we view risk and risk management. It is a highly imperfect tool and cannot be counted upon to generate effective anticipation or mitigation strategies. It can, on the other hand, serve as a meaningful discipline—a requirement for senior leaders to focus on the implications of uncertainty for their projects, and to discuss risk in extended and direct terms.

A second leading lesson is that the effects of uncertainty reinforce one of the basic arguments of the volume, which may be a key lesson of the crisis: Risk management is destined to fail; our mechanisms for taming it are simply not up to the task. Conventional wisdom says risk management should have worked; it's just that firms did a terrible job of understanding the true character of the underlying assets (such as subprime mortgages), and so risk was "mispriced." But there's another interpretation—that probabilistic risk instruments based on the idea that markets are efficient and actors are utility-maximizers will *always* make such mistakes, because they are confusing risk and uncertainty and trying to quantify what cannot be quantified.⁶¹ As the business scholar Barry Schachter has argued, "These aren't bugs," in the risk management system—"they're features."⁶²

Finally, the role of uncertainty points to the chief value of risk management as being outcome-oriented. The psychological effects of uncertainty are aimed precisely at rigorous outcome analyses. In uncertain situations, it is very easy for senior leaders to abandon a rigorous outcome analysis (and mitigation effort) on the excuse that it's just not possible. Because the environment is so uncertain, they can't adequately judge what might happen.

An appreciation for the role of uncertainty should lead us to rethink the way in which we prepare for possible dangers to our projects. It may be that major institutions operating in such contexts should conceive of their activity not as managing risk, but managing uncertainty. That mindset would focus not on quantifying specific potential risks in an expected future, but at identifying many potential outcomes and building resilience against the broadest range. If dealing with risk in an uncertain environment is a radically different process than controlling it in a deterministic one, then the approach and mindset ought to reflect the difference. The primary task on the specific kinds of issues under consideration here—complex strategic judgments—is indeed managing uncertainty. Chapter 14 will spell out what this might mean.

5

Risk Is What We Make of It

As the financial crisis grew and matured into 2008, markets became highly volatile, ready to jump at the slightest indication of trouble. Traders began reacting, not to the reality of a firm's economic situation, but to rumors and skewed perceptions of it. In March of that year, Bear Sterns felt the wrath of this accumulating wave of perceptions.

That month a broad-based run began on Bear stock, triggered by the perception that it was low on capital. This was not the case, but the impression spread that it was. As the financial journalist Brian Burrough has written, "There has never been anything on Wall Street to compare to it: a 'run' on a major investment bank, caused not by a criminal indictment or some mammoth quarterly loss but by rumor and innuendo that, as best one can tell, had little basis in fact."¹ Once it began, the mass hysteria was accelerated by media outlets with an incentive to maximize controversy and perceived harm: CNBC anchors cited anonymous "sources" telling them apparent untruths.² If some leaders of Bear and other firms are to be believed, the trend was also shepherded by devious short sellers, anxious to undermine confidence in firms against which they had bet.³

To be sure, there were plenty of toxic assets floating around the financial sector, and some firms were legitimately in trouble. The perceptions that brought Bear low had a basis in reality. But very quickly the difference between fact and rumor, perception and reality was lost by markets desperate to avoid huge losses. The same phenomenon was very much in evidence in the Asian financial crisis of 1997–1998, in which the "Electronic Herd" (so named by Thomas Friedman) helped drive the collapse by overreacting to financial concerns in South Korea and elsewhere and stripping capital out of these countries.

Once perception gathers momentum, it becomes very difficult to head off. Even prominent steps announced by financial regulators and public banks—promises of injected cash in 2008 by the Fed, IMF rescue plans during the Asian crisis—did little to restore confidence. "Time and again," the financial

journalist Paul Blustein has written about the Asian case, “panics in financial markets proved impervious to the ministrations of the people responsible for global economic policy-making. IMF bailouts fell flat in one crisis-stricken country after another.”⁴

They fell flat, in large measure, because of the central theme of this chapter: the subjective character of risk. A major lesson of the crisis, one with ready parallels in the national security field, is that different people can see the same situation—the same event, the same investment, the same business or strategic environment, the same war plan—as involving large risk or small, likely or unlikely, catastrophic or not that big of a deal. Decision-makers do not simply read or measure the meaning of the situations. They *impose* meaning in an inferential and creative act of judgment, one influenced as much by emotion as reason. This substantially complicates the process of risk assessment and management.

As we have seen, in largely deterministic contexts, leaders can make risk judgments grounded in much more comfortably objective criteria. If determining how to structure an aircraft safety regimen, or how to place ships in a maritime search for missing sailors, or whether to charge \$500 or \$750 for a particular form of insurance, data-based analysis can play a decisive role in determining the decision. In this chapter as in the larger study, I am speaking specifically about complex strategic judgments, as defined by the criteria laid out in Chapter 1. In such cases, when confronted with situations of high uncertainty that offer no clear objective truths, human beings manufacture a subjective understanding of the situation, and of risk.⁵ Risk is what we make of it, in an interpretive encounter with reality—a fact that has substantial implications for national security planning.

This chapter makes this argument in three stages. First, it sketches out the role of perception in governing risk assessment under uncertainty. Second, it draws in some of the evidence from the financial industry case studies, arguing that the experience of the crisis strongly suggests the perceptual nature of risk. Finally, the chapter summarizes recent research on the nature of emergent, intuitive judgment—a “feeling-of-knowing” process of reaching conclusions which, although not based on objective data, come to have powerful influence over decision-makers’ thinking. The chapter concludes with some implications for thinking about risk in the national security realm.

There is no objective risk to manage

The basis for this conclusion flows out of the previous chapter, and the character of judgments under uncertainty. When a situation does not offer up a potential “right” or “true” answer, decision-makers are forced to fashion the basis for a judgment, in effect to create one. They will use reasoning and refer to evidence, but the same set of facts and even the same broad interpretations

and arguments can produce radically different interpretations of the same events. Kenneth Boulding has argued that the very definition of a decision is a “choice among alternative perceived images of the future,”⁶ and this is very true of judgment under uncertainty—it’s the *perceived images* that govern the conclusions.

An example is the threat of a North Korean invasion of the South—a risk to the current US forward presence strategy. In assessing this risk, strategists could take into account North Korean history, the statements of the regime, the levels of its seeming desperation, the status of its relations with the South, and so forth. Yet looking at these same factors, different observers come to very different conclusions. Some consider the danger to be substantial, believing North Korea could be tempted into an attack; others dismiss the risk and suggest that North Korean leaders know they would be committing national suicide. The debate continues if we shift the subject to the risks attendant to a policy action, such as withdrawal of US forces: Some believe such a step carries severe risks of conflict; others do not. It all depends on the various perspectives, or lenses, they bring to the issue.⁷

In the realm of complex strategic judgment, risk is a function of perception more than fact or data. Diane Vaughan, in her magisterial history of the Challenger disaster, explains that “Risk is not a fixed attribute of some object, but constructed by individuals from past experience and present circumstance and conferred upon the object or situation.” People assess risk “as they assess everything else—through the filtering lens of individual worldview. ... Risk is in the eyes of the beholder; it can be present or absent in the same situation or object, or, if present, present to a greater or lesser extent.”⁸ One implication is that decision-makers in real situations do not always focus on the classic components of risk assessment—likelihood or consequences—but instead will manipulate information to support preconceptions.⁹

A famous example from the psychology literature dates from the 1950s. After an especially brutal football game between Princeton and Dartmouth, surveys asked students who had started the rough play. Naturally, large majorities from each college blamed it on the other side. But even when students were shown an actual film of the game and asked to note penalties as they saw them, the answers remained just as biased. “It is inaccurate and misleading,” the researchers concluded, “to say that different people have different ‘attitudes’ concerning the same ‘thing.’ For the ‘thing’ simply is not the same for different people whether the ‘thing’ is a football game, a presidential candidate, Communism, or spinach.”¹⁰ Or, as we are coming to understand, risk.

Even preference setting is governed by subjective and invented factors. People are not very good at knowing what they want, and so their preferences emerge, or are constructed, as they collide with choices and judgments.¹¹ And when decision-makers are inventing meaning, they are not necessarily disposed to be careful about truths, and actually tend to be more concerned about

reinforcing their existing beliefs.¹² The lack of objective anchors allows observers to make whatever meaning they like—an invitation to cognitive bias.

Divergent interpretations then become an engine of the sort of “transmutability” mentioned in the previous chapter. Decision-makers interpret the same situations in very different ways, and then act on those perceptions in ways that can produce different futures. Because of such subjective responses, the same basic situation can give rise to very different results. Mark Blyth has explained that in the Great Depression, “outcomes as diverse as Japanese fascism and Swedish social democracy appeared in various places under broadly similar economic conditions.”¹³ The same inputs can cause radically different outputs, making it impossible to generalize. In complex and uncertain situations, each case is a one-off event that cannot be included in large data sets.

This subjective, interpretive approach to judgment tends as well to look for narratives and stories to give meaning to events. Nassim Nicholas Taleb refers to a “narrative fallacy” that is based on our “limited ability to look at sequences of facts without weaving an explanation into them, or, equivalently, forcing a logical link, an *arrow of relationship*, upon them. Explanations bind facts together,” but the risk arises when this habit “increases our *impression of understanding*.”¹⁴

Interpretive risk in the financial crisis

We see this subjectivity of risk crop up over and over again during the financial crisis. During the buildup of the bubble, for example, leaders of most financial institutions perceived as manageable the risk involved in highly complex derivatives built on top of shaky mortgages. They did so in part because of assurances from those who designed these instruments, who assured them that they could not go wrong—for example, that housing prices would not all fall at once. They did so because they had accumulated faith in such instruments, based on years of experience. And they did so because they were motivated to see safety in profit-generating strategies.

But these were subjective judgments dressed up as a financial truth, and a large number of skeptics pointed this out beginning years before the crisis hit and continuing right up to 2007. Once the crisis began, risk perceptions flipped, and investors cast a critical eye at any financial institution, regardless of the true numbers on its books. The primary issue was not the objective reality as much as the perception of it.

The argument here is not that objective economic facts—or the realities of a national security situation—have no value. Nor am I suggesting that it’s pointless to get the best objective portrait of a situation. In fact a rigorous assembly of information is critical to avoid, as best possible, the malign influence of skewed perceptions. (Had a more objective and less profit-chasing mentality been operating in the early 2000s, the rush to risky derivatives would have

been far less severe.) The argument here is that perception *overwhelms* reality, that created meaning is more important than objective facts. This is true in part because, in a situation of true uncertainty, *the available facts will never be enough to justify an objective risk judgment over a subjective one*. Only rigorous analytical techniques, self-aware organizational culture, and other routes to effective risk management can cure the problem.

Part of the challenge is an essential problem with induction or inference.¹⁵ We tend to fill the space between facts and interpretations with confirmation bias, invented patterns, and other products of our imagination. We can think of the problem as the bridge or gap between facts and conclusions, and it is especially dangerous when we are inferring truths under conditions of great uncertainty.¹⁶

Because the situation offers no objective answer, the decision-maker must fill the yawning gap between facts and interpretations. Under uncertainty, inference is an act of imagination rather than logic. And in doing so, people engage in all manner of cognitive shortcuts and potential errors: skewed risk perception, unwarranted faith in their own skill. The result is a process of envisioning futures, a creative and imaginative development of meaning. The political scientist John Steinbruner has argued that, in complex strategic decisions, "the burden of establishing a stable pattern of beliefs falls to a greater than usual extent upon the imagination of the decision maker."¹⁷ As the scholar Leonard Mlodinow has written, "Human perception ... is not a direct consequence of reality but rather an act of imagination. Perception requires imagination because the data people encounter in their lives are never complete and always equivocal."¹⁸

This simple insight is in many ways the centerpiece of this whole analysis. In complex strategic judgments, decision-makers are not reading objective facts to derive linear conclusions. They are manufacturing inferential judgments through imaginative leaps, stoked by unconscious impressions over which they have limited awareness.

Feeling-of-knowing and emergent judgment

Judgment, then, is in many ways a process of allowing the subconscious to gradually produce an emergent belief in the absence of direct access to verifiable truth. Judgments represent a gradually arising feeling-of-knowing, a gut sense, that eventually allows a decision-maker to feel comfortable about a conclusion.

Suppose you have a complex decision to make, involving many different elements, significant stakes, important personal values, and a high degree of uncertainty. Suppose, for example, that you are deciding where to attend college. You will have a number of clear facts to play with: the relative cost of various schools; their location; some aspects of their character (are they

studious or not, are fraternities and sororities the center of social life?). You'll discover a few basic statistics about graduation and job placement rates and probably be able to harvest some perceptions from surveys of current and former students.

Ultimately, though, in a choice like this, the most important factors in the decision will remain largely guesswork. Merely defining "success" in the decision will be challenging enough. How would you be able to know, 10 or 20 years later, if the choice was the "right" one? By whether you're "happy" during your time at school? (How would you measure that, anyway?) Whether you think the education is the highest quality you could have received? (Again, measured how, exactly?) The best value (according to what criteria)? The highest starting salary after graduation? There will be some obvious categories of failure—if you're miserable, feel out of place, become convinced that the place is second-rate. But knowing if you made the "best" or "right" choice among your many options will be extremely difficult. Even if you end up transferring away, you might still admit that the decision to attend the school in the first place was the best you could have made at the time.

This is true in part because you're not just trying to get a single thing out of college. Your objectives with the choice are multiple and somewhat conflicting. You want to enjoy the experience, challenge yourself, set yourself up for success in career and life, come into contact with talented professors and students, discover a new part of the country, emerge with a minimum of debt, and more. Selecting a college is very often about weighing, balancing, and managing the trade-offs among multiple goals, where the variables are not all known at the outset and causal relationships are often mysterious. There is simply no way to do that in anything like an objective, quantified, probabilistic, or indeed reliable fashion. There can be no algorithm that can capture all the relevant factors, in part because so many are qualitative and because the relationship between and among the variables is unpredictable and variable. The process of balancing competing values and preferences is by definition both *subjective* and *intuitive*, reflecting more of a gut feel than an analytical calculation—which story or narrative about the balance feels "right" to you.

Early on in the process you figure out that if you try to learn and master every key fact relating to the decision, you'll drive yourself nuts. You may start making lists, keeping spreadsheets, assigning values—then find yourself going back and changing numbers you'd assigned to make the results match what you "feel" they ought to be. Every week, some new category of information is likely to rise to your attention (the quality of the food at the school, or the placement rates of one program or department as opposed to the whole college) that you hadn't even thought to consider before. You confront epistemological uncertainty: There's just no way to gather all the facts relevant to the decision.

Thinking further about the matter, you realize that, more fundamentally, you are also dealing with ontological or aleatory uncertainty, the "transmutability"

of the situation. You can't make a truly objective judgment about the outcome in part because the causal links between your choice and the outcome are swimming in a vast sea of variables that are constantly evolving, partly under the influence of the choices of others. It's as if you're plotting a course from Point A to Point B—except that the map is constantly changing, with landmarks appearing and disappearing. In the context of strategic choices, hundreds of actions, behaviors, changing attitudes, trends, and events refashion the context for your decision over time; a professor whose fame and skill attracted you to a school might leave after your freshman year.

Decision-makers have no objective way to calculate whether the decision will be the right one because so much will have altered by the time they are in a position to make that call. The world when they arrive at their "outcome" will be different—both objective facts and their perception of them—from the destination that existed when they began.

About the only judgment method that makes sense for this sort of a decision, wallowing in uncertainty, is to gather a broad range of facts, "sleep on" the issue—letting the unconscious process all the considerations—and gradually gain a gut sense about the best course. We have all experienced this sensation: mulling something for hours, days, or weeks; taking in hundreds of pieces of largely suggestive and uncategorized information; and then gradually having a "sense" that one answer is "right." The result will be highly intuitive and subject to various analytical errors, but it's the answer that will feel right, and the decision-maker will be able to defend it with powerful arguments (which may or may not actually reflect the reasons why their unconscious settled on that preference).

It is typical to associate this sort of thinking with informal decisions, or personal ones—what sort of dishwasher to buy, what college to attend. But part of the argument here is that senior decision-makers in business and government employ the same basic approach. They gather information about an issue and allow their subconscious to suggest a feeling-of-knowing solution. Sometimes the reaction will be very rapid, if the possible answer connects with long-standing belief systems or worldviews. With more unfamiliar or confounding cases, it may take time for an answer to emerge. But the basic process is the same, because it could not be otherwise: There is no objective utility function that can resolve an issue, and so the matter comes down to subjective judgment.

Feeling-of-knowing judgment

This mental process has a number of characteristics, all of which were evident in the events of the financial crisis. It is, for one thing, a simplifying mechanism that aims to boil down the complex variables of an uncertain situation to one or a few basic decision rules. As in Steinbruner's cybernetic theory, the mind "imposes an image and works to preserve that image" to make sense

of data, settling on a “single course of events” that is the default pattern and then molds information to support it.¹⁹

This is a very different mechanism from the one assumed by neoclassical economics. In economic theory, people are generic utility maximizers.²⁰ “Rationality” is defined as congruence between ends and means, or the consistent and purposeful pursuit of well-defined ends.²¹ As we have seen, though, under uncertainty these sorts of assumptions no longer hold. “In the world of risk the assumption that agents follow consistent, rational, instrumental decision rules is plausible. But that assumption becomes untenable when parameters are too unstable to quantify the prospects for events that may or may not happen in the future. The past is not a prologue. ... In this new environment there is no basis for agents to settle on what the ‘objective’ probability distribution looks like.”²² Under uncertainty, the feeling-of-knowing dominates.

The economist Frank Knight offered a wonderful description of this process. Because of our fundamental lack of knowledge about the future, rendering a judgment, Knight contends, becomes a matter of opinion rather than analysis. When we approach a judgment under uncertainty, Knight concludes:

we are likely to do a lot of irrelevant mental rambling, and the first thing we know we find that we have made up our minds, that our course of action is settled. There seems to be very little meaning in what has gone on in our minds, and certainly little kinship with the formal processes of logic which the scientist uses in an investigation. ... There is doubtless some analysis of a crude type involved, but in the main it seems that we “infer” largely from our experience of the past as a whole, somewhat in the same way that we deal with intrinsically simple (unanalyzable) problems like estimating distances, weights, or other physical magnitudes, when measuring instruments are not at hand.²³

In other words, we make an intuitive guess—a judgment. “We act upon estimates rather than inferences, upon ‘judgment’ or ‘intuition,’ not reasoning, for the most part.”²⁴ The economist Peter Earl has called this sort of thinking “very much in the nature of a leap in the dark.” A decision-maker’s mind “will cobble together” a picture of events “until something clicks into place.”²⁵

Bruce Mangan has similarly written about nonconscious, intuitive feelings of familiarity and a “feeling of knowing,” which he also calls a sense of “rightness.” These perceptions are based on nonconscious inputs, but “feel” very strong—which come to us as “gut feelings,” “just knowing,” hunches.²⁶ The psychiatrist Robert Burton hypothesizes that the source of the process resides in the brain’s limbic system, in stimuli that generate unconscious patterns that eventually give rise to convictions and senses of things—certainties, a feeling that something “just is” right—without a decision-maker’s fully knowing why.²⁷ The resulting process is not based upon a careful weighing of facts

and conflicting evidence. It is based on a very rough and loose absorption of data which coalesce into a sense of a right answer (or decision), a sort of subconscious magic elixir whose precise mechanisms remain elusive and whose workings differ remarkably from one person to another.

A related and important lesson of the financial crisis—and a concept that ties the issue of meaning in with the role of incentives, which will be examined in a later chapter—is that the subjective risk perceptions of individuals are always grounded in some social context. The meaning that a decision-maker brings to a situation will be a mixture of his or her own experiences, beliefs, worldviews, and so forth, and internalized values, norms, and beliefs from larger groups—companies, military services, nations. Individual perspectives are shaped by institutional norms, Diane Vaughan explains from her research into NASA decisions. These “create unreflective, routine, taken-for-granted scripts that become part of individual worldview.”²⁸ A person’s perception of risk will likely reflect some degree of larger truths that are widely accepted within a country, an industry or an organization, although every person’s perspective will be a unique mix of influences, both social and personal.

Powerful examples of social influences on risk perception were on display in the financial crisis. These included, obviously enough, the priority on short-term profits, as well as the related notion that finance was being revolutionized with complex new instruments, and that firms which would succeed would embrace rather than reject them. Particular institutions inculcated higher risk tolerance than others, and specific leaders responded to the influences.

What we see when we consider risk, then, is a subjective truth we have imagined for ourselves. A great example is Long-Term Capital Management, in which the leading scholar-investors had a faith-like belief in their models. Even afterward, they wanted to perfect their algorithms rather than abandon them. One partner continued to insist that their calculations of the probability of extreme market events remained accurate—even after being reminded that “it had happened, not once in a hundred years but many times—in Mexico, on Wall Street, in stocks, in bonds, in silver, in Thailand, in Russia, in Brazil.”²⁹

This reflects one of the primary dangers with feeling-of-knowing forms of judgment: They can be immune to reflective consideration. Intuitive judgments arise smoothly and easily, and leave decision-makers with a perception of confidence that may not be justified by the evidence available or the relative merit of their favored option. Such intuitive confidence is not characteristic of every judgment, they explained, and can be counteracted. But when it is in evidence it will dim alternatives and make the intuitively appointed choice far more attractive than an objective consideration would warrant.³⁰

All of these patterns were in evidence in the financial crisis. The judgments involved seem to have had the character of emergent, feeling-of-knowing impressions—that the whole industry was profiting from derivatives, for

example, and few firms felt able to stand aside and watch the profits flow by them. Once in place, these beliefs proved impervious to contrary evidence or warnings. As we will see, there was more than enough reason to believe that the derivatives would end up causing mayhem. But senior leaders in most firms simply refused to listen, in part because the meaning they brought to events had been so firmly established.

Lessons for risk management in national security

The fact that risk is what we make of it carries a number of implications for understanding the nature of risk management in national security. Most essentially, while the seeming goal of risk management is to standardize risk, risk is inherently subjective in a way that threatens to overwhelm our efforts to assess and manage it. Risk appetite will vary according to leaders and circumstances.³¹ So the challenge is to manage the perceptions of senior leaders as much as it is to assess risk. “Failures” in risk management are often not failures in risk management at all, but failures in overall management and strategy by senior leaders.³²

Perhaps the most important aspect of imaginative, feeling-of-knowing forms of judgment, and a key lesson for thinking of risk in national security, is that such mechanisms become highly subject to all manner of cognitive errors. An imaginative process of judgment opens the door for all manner of self-justifying efforts, and overconfidence. The other side of imagination is delusion.³³ Understanding the process by which decision-makers import meaning to subjective events therefore lays the groundwork for the subject of the next chapter, cognitive bias in risk assessment.

Despite their grounding in subjective meaning-making, though—or perhaps because of it—judgments reached through an emergent process of feeling-of-knowing become highly resistant to contrary evidence. Part of the issue is that the whole process is not really based on an objective encounter with facts from the beginning. It’s the emergence of an unconscious sense, a gut feel, a feeling of “rightness” about a certain interpretation or proposed course of action. Decision-makers under the influence of a feeling-of-knowing process of judgment are not terribly influenced by contrary evidence. A fleeting reaction is subject to change; an expectation that has gradually emerged over hours, days, or weeks of information accumulation and intuitive feeling-of-knowing will be very stubborn indeed. Once we have made the initial, feeling-of-knowing choice, our broad, instinctive reactions are highly resistant to outside pressure, indeed to evidence of just about any sort.

We end up thinking our own views are more objective and generally held than is actually the case. Because of this tendency, our preferences “simply seem ‘right,’” as two researchers conclude.³⁴ Emergent judgment shapes our perception of facts, and filters incoming information to match the accumulating

belief. Senior leaders can become immune to warning, in part because they trust their sources of judgment over those of the people offering the risk analyses. The resulting process of judgment has much more in common with the development of belief and even faith than it does with formal logic and reasoning, or calculations of probabilities. In its emphasis on an emergent sense of what is “right”—guided by a combination of subjective perspective and social influences—it tends to generate a nearly moralistic, rather than purely logical, judgment. And the resulting judgments become even more immune to challenge.³⁵ Imagination combined with subjectivity generates belief, self-created ideas about the world to which we then become loyal.³⁶

It turns out that such terms may well capture an important character of the way in which human perception works. As it brings in incoming data and churns it over in the hidden layer, the mind is not merely “thinking”; the process is creating an intermediate stage of analysis and judgment that then adds a level of determination that eventually more closely approximates belief and faith. “It is impossible to overlook the shared qualities of the *feeling of knowing*, a *sense of faith*, and feelings of *purpose* and *meaning*,” Robert Burton writes. “All serve both motivation and reward at the most basic level of thought.”³⁷ The result of emergent feeling-of-knowing judgments is something well beyond a contingent possibility—it is a form of faith.

A major lesson of the crisis, as well as parallel cases of risk disasters in national security affairs, is that these skewed perceptions, confidence in formed belief, and elements of motivated reasoning can be especially powerful when directed at the deeply held ambitions or plans of individuals or groups. As we shall see, it is the combination of uncertainty, meaning-making, and incentives that produces wishful thinking and other forms of motivated reasoning. And a common result is to dampen concerns about outcomes. Senior leaders do not give adequate attention to the risks of their hoped-for plans because they have imposed an interpretation on events that suggests they simply don’t have to. When warnings do come, as we will see in Chapter 7 on “gray swans,” they are routinely ignored.

This was the story, in part, of the financial crisis. It is the story of dozens of national security debacles. And it points to one critical role for risk management: focusing disciplined attention on possible consequences of proposed courses of action. While it will never entirely control the tendency to impose personalized and sometimes biased meaning, such a risk process can at least create a context in which leaders will approach their decisions having been forced to confront potential risks more seriously.

6

Indifferent to Consequences

In March 1965, Lyndon Johnson agonized over a request for Marine battalions to defend US missile and aircraft sites in Vietnam. In a March 6 telephone call with Secretary of Defense Robert McNamara, captured on the White House recording system, Johnson worried that the Marines would end up “fighting with the Vietcong and really starting a land war.” He summed up the discussion with what must be one of the most tragic remarks ever uttered by an American president: “My answer is yes,” he told McNamara. “But my judgment is no.”¹

Just two and a half hours before that conversation, Johnson had been on the phone with his political mentor, Senator Richard Russell. “I guess we’ve got no choice,” Johnson said of the troop request, “but it scares the death out of me. ... A man can fight if he can see daylight down the road somewhere. But there ain’t no daylight in Vietnam. There’s not a bit.” Three months later, in June, he would tell McNamara that “I see no program from either Defense or State that gives me much hope of doing anything.”² And yet he continued to dispatch endless streams of reinforcements to a war that he strongly sensed could not be won.

Leading models of rationalistic decision-making and risk management presume that leaders are profoundly *outcome-oriented*, or consequentialist.³ They develop clear objectives or values (“utilities”), assess options to determine which is most likely to maximize those goals, and then evaluate risks relative to their goals and plans. A fundamental assumption of rationalistic models is what has been called the “prospective orientation”: As the scholar Raanan Lipshitz puts it, “options are chosen by considering the likelihood and attractiveness of future consequences.”⁴ Risk analysis is, or ought to be, particularly attuned to future consequences: Its essence involves the identification of what can go wrong, which is a question all about outcomes.

But in the spring of 1965, Lyndon Johnson was not doing this, at least not in the manner presumed by classic utility maximization or risk management. To be sure, his decision process had some consequentialist ornamentation—discussions of possible dangers, passing consideration of obviously unacceptable

options such as massive escalation or outright withdrawal. But none of these merited serious consideration, and at the end of the day, Johnson took key Vietnam decisions for one dominant reason—because he felt he had to, almost regardless of the consequences. “[W]e know, from Munich on,” Johnson blustered at Senator Everett Dirksen in February 1965, “that when you give, the dictators feed on raw meat. If they take Vietnam, they take Thailand, they take Indonesia, they take Burma, they come right on back to the Philippines.”⁵ Johnson worried as well that his political future was at stake in his response to Communism, and his well-documented urge for dominance made him unwilling to back down.

All of this—the sense of Cold War strategic urgency, an obsession with global credibility, Johnson’s sense of *machismo* and belief that he’d be comatose politically if he abandoned a war—created a sense of inevitability around the choice.⁶ The administration never rigorously considered any option other than continuous escalation. Johnson’s National Security Adviser, McGeorge Bundy, readily admitted that he “did not even consider it necessary to posit a precise military objective for the escalation he was advocating.”⁷ The result was that Johnson was substantially blind to consequences—and thus to risks.

This same phenomenon crops up again and again in cases of strategic judgment and risk management. In a particular category of high-stakes decisions sharing a number of characteristics, senior decision-makers subordinate outcome orientation to many other considerations, including urgently felt imperatives. They do not engage in a true comparison of the costs and benefits of a range of alternatives, and they are not guided by the expected future utility of their actions. They become obsessed with the act itself, and far less concerned about its likely outcomes.

For risk management, this is a serious problem. Risk processes are designed to protect organizations from potential consequences. When senior decision-makers become immune to outcome-oriented thinking for one reason or another, they will not give serious consideration to risk. They may continue to give it rhetorical emphasis, talking about what could go wrong, but the trajectory of their *judgment* will never substantially vary. Considerations of risk will become a feeble ornament on a choice made for pressing reasons. No matter how extensive the structure of risk analysis may be, risk will become immaterial.

Varieties of nonconsequentialist decision-making

Researchers have known for some time that peoples’ assessment of future values is subject to a wide range of curious and sometimes distorted influences. Decision-makers can discount consequences for a whole range of reasons.⁸ They can be interested in justifying a past choice more than maximizing future utility. They can place more emphasis on procedures than outcomes.⁹

They can be driven by an imperative to act in a certain way, regardless of the consequences. They can feel themselves under the shadow of a moral obligation.

A number of theories and concepts speak to the limits on consequentialist thinking, or the role of a single, dominant variable in determining behavior. One example is the “prominence effect,” a theory in which one dominant factor determines the decision-maker’s preference.¹⁰ This theory suggests that, when confronted with a complex and uncertain choice, people in various roles, such as consumers or risk assessors, will elevate a single decision rule above all other considerations, and adhere to it even when it does not appear to maximize stated goals. The prominent variable overwhelms the choice.¹¹

A second related concept is the psychological notion of affect, a term whose broadest sense refers to any emotional reaction but is sometimes defined more precisely as “the specific quality of ‘goodness’ or ‘badness’” connected to a situation. Outcome-denying judgments can be grounded in an emotional reaction to a situation, and the literature on affect suggests that such impulses offer a “quicker, easier and more efficient way to navigate in a complex, uncertain, and sometimes dangerous world” than more deliberate analytical methods. A group of researchers including Paul Slovic have described the particular effect on risk management of what they call the “affect heuristic,” a judgment tool built on the quick, emotional, “gut” response to issues. “Intuitive feelings are still the predominant method by which human beings evaluate risk,” these scholars have concluded.¹²

Affective decision rules can take the form of “visceral” influences or factors, things like cravings for drugs and sexual desire.¹³ Visceral influences “can give rise to behaviors that look extremely impatient and even impulsive,”¹⁴ which is precisely what we see in strategic actions that disregard consequences and risk. Research suggests that when people feel a strong immediate emotion about an issue, perceptions of risk simply glance off their well-armored justifications. Smokers, for example, craving the immediate reward of the immediate nicotine rush, ignore evidence of long-term health hazards.¹⁵ George Loewenstein, a leading researcher in the field, concludes that “visceral factors can be so powerful as to virtually preclude decision making,” and one could easily add appreciation of outcome-oriented risk. “No one *decides* to fall asleep at the wheel, but many people do.”¹⁶ The parallel is clear enough: No CEO *decides* to bankrupt their company, and no president *decides* to commit their nation to a destructive quagmire—but many do.

A third set of factors that can impede appreciation for consequences and risks derives from economic research on future-oriented activities like saving and other “intertemporal” choices.¹⁷ Studies have asked people to compare immediate profits to potential long-term payoffs and guaranteed money versus probabilities of future payoffs. Two factors seem to be at work: immediacy and certainty. One scholar explains that “ample evidence shows that decision

makers care ‘too much’ for things that happen immediately, and seem incapable to attribute proper value to delayed events.”¹⁸ Immediate-payoff choices also seem more certain, and thus engage a distinct but closely related decision rule. As Daniel Kahneman and Amos Tversky put it, “people overweight outcomes that are considered certain, relative to outcomes which are merely probable—a phenomenon which we label the *certainty effect*.”¹⁹

A critical implication is that, when outcomes are dimmed, choices are never analytically evaluated to determine if the assumptions behind the utilitarian justification are really true. Will a failure to respond lead to aggression? Will responding deter it? When other imperatives wash away a concern for outcomes, decision-makers do not need, and often fail to develop, any theory of how their actions will produce the intended effects. The viability of a proposed option, as well as the risks it involves, are secondary. The action itself serves a felt need.

Such patterns are often called “nonconsequentialist” decision-making,²⁰ because people are guided by a sense of what they must do rather than an effort to maximize utilities. Outcomes become irrelevant in the face of deontological, absolute values that veto other considerations. In contrast to the goal-directed consequentialism of utility theory, deontological choice is value-driven, concerned more with satisfying certain normative requirements. Evidence is abundant for the situations in which people violate utility-maximizing outcomes in favor of ethical or cultural norms. My argument, in effect, is that either short-term business imperatives or historically grounded national security maxims can serve as ethical imperatives for senior leaders. It would simply be “wrong” to violate such principles as the need to drive quarterly profit margins or the international credibility doctrine, in the same way that it would be wrong to trade children’s lives for the efficiency of a school bus route. Some utilities trump others, and the hierarchy of utilities can be so steep and extreme that the one or few at the top count as imperatives. In such situations, the act of following the norm becomes its own utility.²¹

Jonathan Baron and other researchers have examined a particular variant of such decision processes called “protected values.” These are absolute decision rules that resist any form of trade-off with other considerations. They are deontological in the sense that they dictate a certain outcome “whatever the consequences.”²² Such values have a quality of “absoluteness” that makes them function as decisive rules for choice. People work to rationalize protected values, suggesting again that there is a halo of utilitarian thinking around even such absolute rules. But the fundamental mechanism at work is an immediate requirement rather than considerations of utilities or risks.

Philip Tetlock has discussed a similar concept, of “sacred values,” which refers to “cultural-identity or moral-religious imperatives.” Tetlock offers the example of a decision-maker who thinks in terms of the “opportunity

costs” of professional integrity. Someone asking such a question is likely to be told that they “just don’t get it”—the similar response, ironically, to those who challenge imperatives in a national security decision-making context.²³ “Research on sacred values,” Tetlock concludes, “suggests a supplementary perspective that posits people to be intuitive theologians struggling to defend sacred values from secular encroachments.” The risk is that “intuitive theologians are suspicious, and unapologetically so, of the classic Enlightenment values of open-minded inquiry and free markets. Opportunity costs be damned, some trade-offs should never be proposed, some statistical truths never used, and some lines of causal/counterfactual inquiry never pursued.”²⁴

In his encyclopedic account of the origins of World War I, Christopher Clark describes senior decision-makers obsessed with almost everything *but* the assessment of risks. French leaders were convinced of the urgent need for “firmness” with Germany. In Austria, the “perplexing narrowness” of the debate over war was a product of the fact that “Austrians were so convinced of the rectitude of their case and of the proposed remedy against Serbia that they could conceive of no alternative to it.” Many actors in the crisis, Clark suggests, believed that they were “operating under irresistible external constraints,” and as a result refused to take responsibility for the events they themselves had set in motion.²⁵

The scholars Eldar Shafir and Amos Tversky have written that “people often fail to consider the possible outcomes and consequences of uncertain events.”²⁶ A range of empirical evidence suggests that this is a common mental process in strategic judgments. To the extent that it is in evidence, risk management will become largely irrelevant. This cognitive trap was very much in evidence before the financial crisis, and helps to explain how otherwise rational and goal-directed leaders, equipped with the world’s finest risk management procedures, could plunge ahead without significant attention to consequences. They were in the grip of an imperative.

Characteristics of imperative-driven thinking

A specific form of nonconsequentialist judgments involves following, without much analysis or rigorous thought, an intensely felt requirement or obligation that creates a sense that the nation or firm “must” take some action. The imperative can be strategic, political, or personal.²⁷ The result of imperative-driven thinking is to seriously degrade consequentialist perspectives—decision-makers under the influence of an imperative are scratching an immediate itch, not thinking about the possible outcomes of their action.²⁸ As in the case of nonconsequentialist thinking in general—of which imperatives are a specific variety—such decision-makers will largely disregard outcome-oriented risks.

I would suggest a number of indicators to look for in order to recognize when an imperative is in effect, and threatens to undermine risk management. Imperative-driven thinking might be in evidence when:

1. The language surrounding the decision is rife with a sense of urgency and descriptions of what “must” be done.
2. The justifying language has a high quotient of affective terms pointing to the emotional resonance or quasi-moral significance of the choice.
3. Policy memos and dialogues focus on analysis of the requirement to act rather than the cost-benefit calculations of acting.
4. There is only one serious option on the table; the process does not engage a deep consideration of multiple alternatives.
5. Any analysis of consequences or outcomes is perfunctory.
6. Dissent within the circle of senior officials actually making the choice is disregarded or actively mocked or punished.

Imperative-driven thinking arguably lies at the intersection of deliberative and intuitive, or System 1 and System 2, thinking. It does not reflect pure intuition or unconscious, automatic reactions—a baseball outfielder reacting instantly and without thought to the flight of a batted ball, for example,²⁹ or a person leaping out of the way of an onrushing car. There is more reflection and, especially, justification going on than that. But imperatives are not fully deliberative, either, because they bypass key parts of the rationalist model. Imperatives do stem from intuitive, feeling-of-knowing reactions not terribly different from the example of a baseball outfielder: Any search for alternatives or risks will be minimalist and largely unreflective, and will take place under the shadow of the automatic preference already in place. Imperatives will take on the self-justifying appearance of a deliberative process, but they often remain firmly grounded on a largely intuitive immediate reaction.³⁰

The emergence of an imperative is often sudden, but it need not be. Imperatives can arise over time, weakly felt at first but gradually impressing themselves upon decision-makers. In the case of Vietnam, for example, the sense of an imperative to escalate was not fully matured in 1955 or 1960. The more gradual a decision process is, the more it has the opportunity to embody aspects of reflective thinking—being explicit about goals, comparing alternatives according to predicted outcomes. But unreflective imperatives, with all the characteristics outlined above, can also emerge bit by bit.

Imperatives can reflect the immediate application of some long-established doctrine or belief: Strength is necessary for credibility; companies must seek the maximum profits on a quarter-by-quarter basis; allies must be supported; aggressors must be checked. In the case of Vietnam, for example, the imperative flowed from several firmly held national security doctrines. In the financial crisis, it was a product of basic capitalist theory. And the argument here is

not necessarily that such doctrines are incorrect or irrelevant to the situations at hand: Companies *do* exist to make a profit. But once such principles have transformed from a general worldview into an imperative, decision-makers can become insensitive to consequence, and risk.

Yet imperatives will often emerge quickly, and part of the reason relates to the environment for high-level strategic choice. Decision-makers at the top of the US national security process, for example, generally have too little time available to anticipate issues, examine an issue in depth, and debate consequentialist issues at great length. They tend to face a long series of crises or challenges that crop up suddenly, demanding a quick answer. The same is true, of course, of leading financial institutions: Accounts of decision-making at the top of the corporations sucked into the recent crisis are rife with stories of executives too harried, overloaded, and exhausted to step back and become more deliberate. Such a context is tailor-made for imperative-driven thinking, because it confronts decision-makers with rapid-fire senses of urgency and a need to respond. They simply do not have time for classic outcome-oriented utility calculations.

Once it is in evidence, imperative-driven decision-making differs from rationalistic or deliberative models of choice in a number of ways. At the outset of the decision process, for example, *imperative-driven thinking is likely to obstruct careful analysis of utilities or objectives*. Decision-makers guided by an imperative will be focused overwhelmingly on satisfying the felt obligation. They are likely to make little effort to critically analyze goals or objectives, and when they do it is likely to have the character of window dressing rather than serious deliberation. One hallmark of an imperative-driven process will therefore be that decision-makers and those eventually implementing a policy will end up repeatedly asking themselves what they were trying to achieve in the first place.

Second, experience suggests that *imperatives are likely to generate subjective and shifting utilities rather than constant and objective ones*. Utility functions are meant to be stable and consistent over time. Yet imperatives are entirely contextual. The felt need to act will depend on many factors unique to a specific situation, and they can reflect changing and inconsistent objectives.

Third, *imperatives will be a function of personality, style, and strategic culture more than an objective assessment of utilities*. This argument suggests that they emerge from a largely intuitive, emergent, feeling-of-knowing reaction to a situation of radical uncertainty. Such largely unconscious, reactive decision processes are likely to be highly dependent on the personalities, perspectives, and approach of the leaders, organizations, and nations. Different leaders will conceive of imperatives very differently. Their perception of risk will vary wildly based on their individual perspective, a danger we will encounter in a later chapter.

To make things worse, a fourth hypothesis suggests that *decision-makers responding to imperatives will not engage in a legitimate comparison of alternatives*.

In rationalistic models, decision-makers lay out multiple options and compare their effectiveness in promoting the utility functions. The rule of the imperative could in theory still allow this—decision-makers could compare ways of meeting the obligation—but far more often it will nominate a single default option that seems like the only acceptable way forward. Without alternatives, risk analysis becomes much less meaningful.

Partly as a result, and in one of the cardinal implications of imperative-driven judgment, *decision-makers under the influence of an imperative will be blinded to many potential consequences and risks outside the scope of the imperative*. In rationalistic processes, outcome orientation and trade-offs among competing objectives are assumed: Only by considering possible consequences and their likelihoods can decision-makers determine which course will best advance their utilities. But imperative-driven thinking will camouflage and discount many potential dangers. Even worse, under the influence of imperatives, *discussion of potential risks and second-order effects is likely to be downplayed and even actively discouraged*. Those who raise risks will be either ignored or dismissed with quick rationalizations. They are also likely to be resented, because they are questioning the uncertainty-simplifying tool on which the decision-maker is relying.

Needless to say, these hypothesized effects are not in evidence on every decision or issue, or to the same degree every time they do appear. Further analysis could develop specific criteria to suggest when the rule of the imperative will be more in effect, and when it will tend to be displaced by more legitimate deliberation and outcome orientation. As an initial hypothesis, based on the case studies reviewed for this analysis, decision-makers may be most likely to fall under the influence of an imperative when:

1. The perceived threat appears extremely urgent;
2. The threat engages vital or serious national interests;
3. The danger has already demonstrated its seriousness through some event or action that creates a panic or mania among the public and decision-makers;
4. The issue engages simplified beliefs or decision-rules resident in the world-views of most decision-makers (such as “aggression must be punished”);
5. Decision-makers have constrained windows of time in which to consider an issue, and limited mental energy—amid dozens of competing challenges—to devote to it;
6. The political stakes of the issue are high or even existential; and
7. There do not appear, at least on initial discussion, to be any serious alternatives to address the threat beyond one default choice.

On longer-term, less urgent, less politically salient questions, on the other hand—such as the gradual development of a new product line or weapons system that does not reflect an answer to an urgent danger—imperatives can

give way to deliberative decision-making. Imperative-driven thinking is a very particular form of judgment, and will only cloud the view of the future in a modest number of cases. But they can be critically important issues on which imperatives can generate disaster.

As suggested in the final characteristic above, a consistent implication of imperative-driven thinking is to magnify aspects of groupthink and the punishing of dissent. A shared imperative becomes a collective belief, which must be defended. This is especially true because imperatives often have powerful roots in deeply held unconscious values, whether individual, social, or institutional. Because they emerge from a mostly unconscious, emergent judgment, they tend to be grounded in both an intuitive sense of “rightness” connected to personal worldviews and experience, and to accepted social and strategic norms—anything from the predominance of short-term profit to the critical importance of national credibility. Partly as a result, feeling-of-knowing imperatives are highly resistant to contrary evidence. Individuals become unwilling to consider evidence suggesting that the imperative could be dangerous or mistaken, for obvious reasons: They have committed themselves to a truth, a decision rule, and admitting the potential of risk would undermine their conception of reality—and, in many cases, their social standing or identity within their group. No matter how serious an organization thinks it is about identifying and mitigating risk, no matter how elaborate its risk processes may be, such imperative-driven thinking can override it all on the path to disaster.

Outcomes, risk, imperatives, and the financial crisis

Many decisions in the financial industry leading to the crisis of 2007–2008 reflect such nonconsequentialist, imperative-driven thinking. Many of the reasons suggested above produced a widespread lack of appreciation for outcomes, and specifically risky ones. The result, as we now know, was to grease the wheels of disaster.

In the years before 2007, dozens of leading financial institutions made terrible bets on subprime mortgages and the various derivative instruments built upon them. In most cases they did so because of an overriding imperative: the need to maximize short-term profits in the context of quarterly reporting demands and a business environment becoming more harshly competitive every year. The chief obligation of senior leaders in the financial industry in the mid-2000s was keeping pace with industry revenue and profit levels. They could not be seen as falling behind the profit-generating innovations of others. The potential risks of this course became secondary to the perceived need to respond.

A leading culprit was the financial products division of American International Group—AIG FP. It had certain policies in place designed to ensure a rigorous focus on outcomes and serious risk management. Its traders, for

example, had to defer half of their pay packages for several years to reduce the incentive for get-rich-quick schemes. But somehow even these policies did not prevent the overall imperative of maximizing profits from skewing decision-making in very risky ways. The hard-driving head of Financial Products, Joe Cassano, reportedly pushed people to get astronomical results and employ risky derivatives—and people did not raise objections because “the money was simply too good.”³¹

In a parallel case, the mortgage giant Fannie Mae found itself in a much more competitive landscape dotted with hard-charging mortgage factories like Washington Mutual. In order to keep up, even to survive, Fannie Mae had to innovate—which meant taking higher risk. But the true character of that risk, and the perils it embodied, were never rigorously understood or even deliberated. And when worrisome signs emerged about the new mortgages, Bethany McLean has written, “senior executives disregarded internal warnings, because the lure of big profits was too great.”³² Fannie Mae chairman Jim Johnson pressed for lax regulations that would allow his firm to become more leveraged—a route to profits but also greater risk, an odd bargain for a publicly held organization. His successor Franklin Raines took the ambitions a step further by setting a rash goal of doubling share prices. An “unhealthy pressure throughout the company” emerged, McLean has documented, to meet the targets. One senior executive gave an internal speech in which he fulminated, “You must have a raging fire in your belly that burns away all doubts” about the goal of boosting share price.³³

Andrew Ross Sorkin catalogues numerous examples of the basic problem: Companies felt an urgent demand to match industry-wide profit seeking. The chance for leaders to become revered for the resulting returns, piled on top of the competitive pressure that everyone else would take advantage of risky new financial products if they didn’t, made the imperative irresistible. And in the process, serious risk analysis was avoided, brushed off, or actively suppressed. “Intellectually,” Sorkin writes, Lehman Brothers chief Richard Fuld “understood the risks associated with cheap credit and borrowing money to increase the wallop of your bet. ... But, like everyone else on Wall Street, he couldn’t pass up the opportunities. The rewards of placing aggressively optimistic bets on the future were just too great.”³⁴ In the run-up to the crisis, the competitive environment had become so intense, and the potential for revenue offered by complex derivatives so enticing, that in a number of cases the leaders appear to have believed that they effectively had no choice. They could either plunge into the emerging areas or suffer a fatal competitive disadvantage—and perhaps personal disgrace.

Imperatives can thus create an environment in which risk management and assessment of consequences become an almost entirely rhetorical enterprise. In the case of the massive, doomed hedge fund Long-Term Capital Management, whose collapse predated the recent financial crisis by a decade,

the fund's corporate investors—mouths watering at the prospect of LTCM's astonishing, and ultimately unsustainable, growth rates—gave little sustained thought to risk. In one case, as chronicled by Roger Lowenstein, a Merrill Lynch executive tried to cut off LTCM from further investments, citing risk. He was quickly overruled by a senior executive: "You *gotta* do these trades, or we lose the business."³⁵ When policies to achieve some end—say profit—generate new risk, as Michael Lewis has written, they beg to be "either honestly accounted for or disguised."³⁶ In an atmosphere of imperative-driven thinking, they will often be disguised or brushed off. "The whole market was pressuring us," one senior Merrill executive told Roger Lowenstein. "To suffer the organization telling you that you are losing business—it takes a tremendous amount [of courage] to stand up and say, 'I'm not going to do it.'"³⁷

This case appears to reflect the criteria outlined above for imperative-driven thinking. For many of these financial institutions, a single dominant decision rule—the requirement to match the profits of competitors, a means to the end of driving quarterly returns—governed their choices. It certainly created an urgent perception of a requirement to act in a certain way, specifically to embrace various forms of derivatives and mortgage-backed securities. There was little discussion, in many of these cases, of alternative courses: Once the imperative to join the parade was established, it pushed aside other options. We have evidence, moreover, that the choices in many cases—while sometimes ornamented with formal risk analyses based on models—were gut calls, based on the aggressive instincts of key financial leaders. And the decisions were certainly guided by a sense of immediacy that overrode serious consideration of long-range effects. The result was to render procedural risk management merely a sort of justificatory decoration on a process obsessed with profit.

The decisions leading up to the financial crisis also display a number of the signs or indicators of nonconsequentialist thinking outlined above. Internal discussions and narratives in the financial industry were filled with expressions of urgency and requirement—"we gotta do this trade." Such dialogues were frequently emotional, reflecting passionately felt urgencies. There is some evidence that internal memos and communications reflected such ideas, though we have limited information on this so far. Several case analyses, however, do appear to make clear that when confronted with the seeming attraction of derivatives and mortgage-backed securities, the leadership of various firms did not give deep consideration to alternatives, and seldom made rigorous analyses of the potential outcomes. The strategy was its own utility. And we know, from a range of evidence, that dissent was ignored and sometimes directly suppressed.

The result of these multiple avenues to the suppression of outcome-oriented thinking is a serious, and sometimes fatal, threat to effective risk management. Leaders unconcerned about consequences will be indifferent to risk. Especially in the manner in which I am arguing for here, as a rigorous analysis

of outcomes, true risk management cannot be seriously pursued in an environment of nonconsequentialist urgencies, imperatives, and decision rules. The experience of the recent financial crisis, as well as dozens of national security cases, suggests that indifference to consequences is one of the most common avenues to the failure of risk management. The solution, as we will see in Part III, is to cultivate what Jonathan Baron has called “actively open-minded thinking,” specifically toward outcomes.

7

The Swans to Worry About Are Gray

For almost a decade now, since the publication of Nassim Nicholas Taleb's brilliant, discursive rumination *The Black Swan*, conventional wisdom has been that strategy—in national security as well as areas like finance—has most to fear from the sudden and the unexpected. A black swan is, at its core, a shock, a surprise. It is an “outlier,” Taleb writes, “as it lies outside the realm of regular expectations, because nothing in the past can convincingly point to its possibility.”¹ He goes on to claim that such events are the engines of history. “A small number of Black Swans explain almost everything in our world,” he argues. Social life “is the cumulative effect of a handful of significant shocks.”²

There is no question that genuine surprises do crop up, and pose a significant challenge to strategists. Taleb was right to focus attention on black swans, and building resilience against unanticipated shocks is a key priority for organizations and nations alike. But in the excitement about one category of strategic challenge—out-of-the-blue threats beyond imagination—there has been too little attention paid to what ends up being a much more common problem for strategy. This is the black swan's little cousin: the “gray swan.”

A number of recent crises and calamities, from 9/11 to the financial meltdown, suggest that most factors that upend strategic intent aren't surprises that no one had anticipated. The much more common problem for complex strategic judgment is risks that *can* be anticipated and that are discussed, debated, and sometimes measured, but which are seen—often incorrectly—as improbable, and are therefore disregarded. This distinction plays a major role in explaining the risk failure of the financial crisis, and leading national security calamities. As national security institutions struggle to deal with the implications of an increasingly complex and unpredictable world, it will be critical to keep this distinction in mind.

Categorizing the swans

As so often occurs when a phrase becomes a cliché, the true meaning of a black swan has been somewhat lost.³ But Taleb was very precise, and he meant to refer to a specific category of strategic surprise: wholly new and unexpected threats that had not been, and in many ways could not have been, anticipated. “Nothing in the past,” he writes, would urge us to expect them. Yet once they occur, leaders feel the need to concoct post-facto rationalizations and claim that they’d seen them coming all along.⁴

By definition black swans reflect non-normal events falling outside the parameters of standard distribution. In a critical phrase, Taleb writes, “If you know the stock market *can* crash,” “then such an event is not a Black Swan.”⁵ But even a brief look at a number of recent disasters suggests that, understood this way, black swans are extraordinarily rare. Much more common is a situation in which a risk was well understood; decision-makers *knew* it could happen, and discussed the possibility. The problem wasn’t that they couldn’t conceive it—it was that they didn’t do anything about it. This category of event can be understood as a gray swan, and it was the dominant form of risk failure in the financial crisis.

Most sources that define gray swans point to the idea of challenges are unlikely but well-appreciated. One defines them as “unlikely occurrences that are just likely enough that they should be anticipated;”⁶ another refers to a gray swan as an “event that can be anticipated to a certain degree, but is considered unlikely to occur and may have a sizable impact ... if it does occur.”⁷ A gray swan can therefore be understood as *an unlikely but fully conceivable risk that lies well within the bounds of experience and has been openly discussed, but becomes discounted and fails to generate mitigating actions*. My definition marries two distinct ideas into a single concept: risks that are conceivable but unlikely, and which fail to produce effective responses.

These two factors need not go together. We could imagine an unlikely event that, once discovered, generates immediate and decisive action. Joining the factors, though, gives us particular insight into the origins of many tragedies, in national security as well as economics. A very common and specific category of strategic challenge is characterized precisely by the combination of these two characteristics: They are known, but unlikely; and *for reasons closely related to this character*, they are dismissed. A gray swan fades into the background of the daily rush of strategy, policy, and process.

This is not quite the same way in which Taleb thinks about what he calls “Mandelbrotian Gray Swans.” He defines them as events “that we can somewhat take into account—earthquakes, blockbuster books, stock market crashes—but for which it is not possible to completely figure out the properties and produce precise calculations.” He seems to have in mind things that we know can happen, but not when or where—things characterized by

somewhat predictable but still ambiguous nonlinear patterns. He calls gray swans “rare and consequential, but somewhat predictable” events.⁸

The concept I am proposing includes that element of unpredictability, but adds the idea that these issues remain ambiguous and tentative enough that nothing is done. They sit at the uncomfortable boundary of the predictable and the uncertain, and don’t carry enough immediacy to generate action. Taleb suggests that “If you know that biotech companies can deliver a mega-blockbuster drug ... then it won’t be a black swan, and you will not be surprised, should that drug appear.”⁹ But that’s exactly the problem: *Awareness does not eliminate surprise*. Knowing that something *could* happen is not enough. We are routinely surprised by things we have long considered as possibilities—financial crises, cyber attacks, Pearl Harbor. Somehow, we wrote them off, so that their arrival caused shock and disaster.

Just about every major “surprise” risk of the last half-century in finance and national security affairs was well within the range of the conceivable. Indeed they *were* conceived—warned about, deliberated. They had been outlined in passionate memos and argued over by top leadership. City planners in New Orleans and elsewhere in the Gulf Coast had worried about a hurricane of the strength of Katrina for years, as Adrienne Lafrance argued in a recent essay in *The Atlantic*. “The city’s vulnerabilities had been well-documented and understood.” The barrier to being prepared wasn’t understanding. It was translating understanding into action, taking the risk—the warning—seriously. Lafrance argues that this has been true for a whole range of calamities, including the 1989 San Francisco earthquake and the 2011 Fukushima nuclear disaster.¹⁰ The problem was not that these events were bolts from the blue. It was that, for whatever reason, the key leaders who considered them in depth did not see a reason to respond in time. And this is a major problem for risk management, because it turns out that, almost no matter how well an organization is assessing and trying to mitigate risk, the human factors that prevent gray swans from being taken seriously enough will thwart effective risk management.

Risks understood—and ignored

The terrorist attacks of 9/11 are a good example. Taleb categorizes 9/11 as a Black Swan: “[H]ad the risk been reasonably *conceivable* on September 10,” he contends, “it would not have happened.”¹¹ But of course the risk *was* conceivable—in fact, many senior officials considered a large-scale al Qaeda attack on the United States a near certainty in precisely that time frame. Counterterrorism czar Richard Clarke’s memoir of the prelude to 9/11 makes clear that he and his team had been warning about the rise and intentions of al Qaeda for years. When the Bush administration came into office, Clarke writes, he briefed all the senior officials that “al Qaeda is at war with us, it

is a highly capable organization ... and it is clearly planning a major series of attacks against us.”¹² Clarke followed up that stark claim with a January 25, 2001, memo in which he sketched out the risk from al Qaeda and called “urgently” for a principal-level review of the threat.¹³ Some warnings even discussed the possibilities of terrorists using airplanes as weapons; the World Trade Center itself had been the subject of a botched previous attack.

This is why the 9/11 Commission report didn’t speak of a black swan. It identified the reason for US vulnerability as a “failure of imagination,” an inability to take *known* risks seriously. “The 9/11 attacks were a shock,” the Commission concluded, “but they should not have come as a surprise. Islamist extremists had given plenty of warning that they meant to kill Americans indiscriminately and in large numbers.”¹⁴

Another example of the perils of gray swans is the emergence of post-war chaos in Iraq. This catastrophe lay entirely within the realm of experience, given the post-war instability that had occurred in the wake of so many other interventions. Analysts, officials, and military officers in the Joint Staff, the State Department, the RAND Corporation, and elsewhere had written of the potential for stability to collapse in the aftermath of the invasion. The issue was discussed at the highest levels of government. Specific proposals to mitigate the risk were made and considered. But nothing was done, and when chaos did erupt, it was hardly a surprise to the dozens of people who had warned about it and urged action.

The statistician Nate Silver has argued that when we study the biggest violations of risk expectations, the issue wasn’t that “nobody saw it coming.” People most always discuss and at some level comprehend the risks. It’s just that they don’t adequately appreciate them, or act to mitigate them, for reasons ranging from overconfidence to motivated reasoning to avoidance to herding. The gray swan lies in wait to ruin us as we stroll blithely ahead.¹⁵

Financial gray swans in 2007–2008

The same pattern held true in the 2007–2008 financial crisis. This event again appeared entirely within the range of experience: Speculative bubbles followed by financial crises have been a recurring part of the economic landscape for centuries.¹⁶ In the mid-2000s, plenty of warnings cropped up, many of them high profile and some coming from the very government officials charged specifically with avoiding the accumulating risks of debt, subprime mortgages, and complex derivatives.

One of the most famous of these cautions was Warren Buffet’s statement, as early as February 2003, that complex derivatives were “financial instruments of mass destruction, carrying dangers that, while now latent, are potentially lethal.”¹⁷ Prominent economists like Robert Shiller and Nouriel Roubini made well-publicized statements of accumulating dangers, and even Treasury

Secretary Henry Paulson argued in 2006 that a financial bubble posed serious risks. Senior risk executives in Lehman and Merrill Lynch pushed tough accounts of growing risks at their bosses. Andrew Ross Sorkin describes one example, a Merrill official named Jeffrey Kronthal, whose warnings “put him directly in the path of [Merrill leader Stan] O’Neal’s ambition to be the mortgage leader on Wall Street.”¹⁸ At the insurance giant AIG, people in and around the financial products division became deeply concerned about the Ponzi scheme being assembled by the financial unit. Top AIG leaders brushed off the concerns.¹⁹

In the earlier case of the giant, doomed hedge fund Long-Term Capital Management, there was more than sufficient evidence available to allow LTCM investors to comprehend the dangers and get out in time. After all, LTCM’s elaborate models were predicated on the near impossibility of exactly the sort of divergence from normal patterns—a generalized financial crisis—that the world had seen as recently as 1987. The expectations of market stability built into LTCM’s models were dangerous, and a number of skeptics pointed this out.²⁰ No one listened. As the business writer John Cassidy has argued, the biggest contributing factor to the recent crisis wasn’t that the risks were unpredictable. The problem, he argues in an echo of the 9/11 Commission, “wasn’t so much a lack of timely warnings as a dearth of imagination.”²¹

This was true of just about every major financial institution that felt the impact of the crisis. “Warning signs were certainly evident” in the mortgage giant Countrywide by the mid-2000s, Gretchen Morgenson and Joshua Rosner have written. Indeed, by late 2004, the company’s risk team was sounding the alarm as more of its loans were going bad. The risk management unit told company leadership, as Morgenson and Rosner summarize the message, that “Risky lending practices were imperiling the company.”²² But the profits were too great, and Countrywide actually accelerated its issuance of high-risk subprime loans after receiving the warning.

An important corollary of the role of gray swans is that organizations typically do not learn well enough from near misses—situations in which disaster is narrowly avoided. People compartmentalize near misses and see them as reasons why future risk went down, when in fact they are harbingers of serious or growing risk. Robin Dillon-Merrill of Georgetown University has demonstrated this through a series of studies on near misses in the space program and other large organizations.²³ They are a form of warning—like those of risk analysts and dissenters—that is ignored for many of the same reasons.

Why don’t people respond?

The most perilous challenge to complex strategic judgment, and most serious threat to effective risk management, are not that risks arrive out of the blue. It is that the warnings, when they come, are ambiguous, and ignored. The key

question is why, and the answer might be found in an interconnected range of psychological and human dynamics that skew risk perception.

For one thing, it can be costly to respond to risks, which can call for expensive, time-consuming mitigation measures. Sometimes, as in the warnings of post-invasion chaos in Iraq, the cost of taking dangers seriously might have been relinquishing a treasured concept of how a strategy will unfold: To take the risk seriously would mean rethinking the whole design of the operation, and perhaps its feasibility. Sometimes, as in the pre-2007 financial warnings, the cost of mitigating risk can be in lost opportunities and profits foregone.

Second, senior decision-makers appear to suffer from a sort of “warning overload.” Leaders juggle dozens of major issues every day, and on at least a few, chances are they’re hearing someone’s breathless claims that a disaster is in the offing. Part of the task of senior leaders is to apply considered judgment to such doubts and avoid overreacting to sky-is-falling Cassandras who would find a way never to take any action at all. But one possible implication of this steady diet of warning could be to dampen the effect of specific ones. Such a phenomenon could also be related to what’s been called “decision fatigue”:²⁴ The mental energy required to make decisions is a finite resource. Confronted with a day full of choices, senior leaders naturally seek energy-conserving strategies—and avoid time- and energy-consuming processes, such as an in-depth analysis of a particular warning.

A third factor obstructing action is closely related to warning overload, and has to do with the institutional culture of operational organizations, whether business or government. They tend to be characterized by a can-do culture of action in which senior leaders are judged by their ability to make things happen, not avoid danger. Bold achievements are revered; avoiding risks long before they emerge doesn’t typically advance the reputation of a senior leader. It’s difficult to take credit for sidestepping problems that “could have” cropped up. The result is to create a culture inherently resistant to warnings—and one that all too often treats those who offer them as annoyances or, at best, well-intentioned fussbudgets.

Such cultural dismissal of people in warning-related positions is knitted together with one of the more fundamental psychological biases affecting decision processes, and a fourth major barrier to taking risk seriously—wishful thinking. One of the best-established empirical facts in cognitive research is that people will believe what they want to believe, in order for their intentions and projects to be realized. One scholar who has reviewed a wide range of evidence summarizes the results simply: “Our cognition is driven and directed by desire.”²⁵ Some researchers have termed the phenomenon “motivated reasoning,” by which people unconsciously *will* themselves to believe certain things.²⁶ Once they are in such a state, trying to convince them otherwise can have the opposite effect to that intended: Asking people to defend their opinions causes them to hold on to them more strongly than if they

were never challenged—making the process of tearing down subconsciously defended belief systems tremendously tricky.²⁷

This sort of motivated or wishful thinking empowers a whole range of related reasoning flaws and cognitive errors. “Once armed with a hypothesis,” people look for confirming evidence, constraining their information search.²⁸ This is sometimes termed “bolstering,” the practice of giving unconscious approval to facts or concepts that serve our goals and disapproving information or ideas that go against our favored outcome.²⁹ Repeated consideration of a possibility makes it appear more feasible—and because we tend to think about what we want, this establishes the connection between desired outcomes and ones considered likely.³⁰ In sum, desire “functions as an information filter, permitting retrieval of those bits of information and the use of those inferential rules for the seeming justification of the target belief. ... People are more likely to arrive at those conclusions that they want to arrive at.”³¹

Motivated reasoning and wishful thinking are not universal. It’s not as if human beings never conduct an objective or rigorous search for the truth. There is a somewhat counteracting desire to be right and accurate, which can lead to detailed and unbiased analytical assessment of issues rather than wishful thinking or closed-mindedness. Various studies have found that people prompted to be motivated by accuracy could overcome some bias effects—for example, when seeing information that challenged existing perceptions, they were more able to treat it fairly.³² Similar results have been found when people were held accountable in some way for the outcomes of their analyses.³³ People have, in fact, many interests, and among top officials being accurate is often one of them; to speak too simplistically about “feeling good” as the singular goal of thought or decision-making is not accurate.³⁴

This was surely in evidence throughout the financial crisis. In firm after firm, the desire to gain profits from complex derivative instruments caused a filtering of risk-related information, and a rejection of warnings about gray swans. Within AIG, for example, senior leaders trusted the aggressive, profit-spewing Joe Cassano and accepted his explanations of the investment vehicles he was building without much question. They *wanted* to believe that it all made sense, and so they did. Wishful thinking also magnifies the effect of decision fatigue, because mentally exhausted leaders seem to actively resist the complex analysis demanded of gray-swan warnings, and assure themselves that everything will turn out fine. This can then become a form of avoidance: When confronted with unhappy possibilities, our minds close off, extending a metaphorical stiff-arm to consideration of inconvenient facts.

Fifth, warnings can collide with the deeply held incentives that inspired the action in the first place. Those who worried about the risks embedded in complex financial derivatives slammed up against the general belief that no firm could abandon these profit-spewing monsters as long as others kept

using them. Those who pointed to the dangers of post-conflict instability in Iraq confronted the deeply held assumptions and desires of advocates of the invasion—Iraq was a middle-class society ready to be freed from tyranny and embrace democracy.

Sixth, gray-swan dangers can be dismissed because decision-makers are in the thrall of an urgent imperative. As suggested in the previous chapter, in some circumstances decision-makers can feel that they *must* act, because of some organizational or strategic or personal demand. In such cases the risks associated with taking that action become largely irrelevant. Many senior Bush administration officials quickly concluded after 9/11 that Saddam Hussein could not be left in power, and brooked no debate on the issue. Many financial leaders in the run-up to the 2007–2008 crisis felt pressed by a bitterly competitive environment to engage in the same risky derivative speculation as everyone else—because they looked around and saw the profits their rivals were raking in. Once an imperative is in play, judgment is substantially foreclosed, and warnings will have little effect.

The impossible need for certainty

Seventh and finally, warnings of gray swans simply cannot be unqualified enough to generate action because they take place in the context of radical uncertainty as assumed by this study. The difference between deterministic and uncertain environments plays a critical role in complicating the environment for effective warning of gray swans, just as it does for risk management more generally.

Whether consciously or not, senior leaders are instinctively aware that they are operating in a context of true or radical uncertainty, at many levels—the complexity of hundreds of variables intersecting to produce emergent patterns, or human agency constantly reshaping the future. Both because of a lack of knowledge of all the variables and because, even with perfect knowledge, no one can anticipate the intersection of choice, forecasting beyond very narrow parameters is bankrupt under uncertainty. The British economist G. L. S. Shackle asks how a decision-maker faced with uncertainty should respond—and he concludes that, “If we ask what in such a case it is rational to do there is no answer, if rationality means choosing the most preferred among a set of attainable ends.”³⁵

This is, in fact, Taleb’s main argument. His work is a damning indictment of linear, probabilistic models of analysis that aim for a best or “right” answer, and a plea to take seriously the implications of uncertainty. For our present purposes, though, the main implication is that a context of uncertainty hampers the ability to warn. “Sure, that *could* happen,” a CEO or cabinet secretary may think. “But so could a hundred other things. We just don’t know.”

In such an uncertain context, those issuing warnings will never be able to be unequivocal enough to force action—and it is precisely in such an environment that psychological dynamics like avoidance and wishful thinking will be most in evidence to brush aside concerns.

Historical cases suggest that this sort of “ambiguity under uncertainty” has much in common with the reasons why intelligence warnings are often ignored. Frequently, calamities chalked up as “intelligence failures” are really failures of responding to warnings that were delivered—just not specifically enough. Decision-makers want an unambiguous alarm; *the attack will come at three in the morning on this date, with this many divisions*. But the equivocal evidence available to analysts seldom allows such precision. This shortfall will then be magnified by the bureaucracies of warning. As risk or intelligence products run through the required approval chain, they are often edited into nondescript statements of the obvious: *Instability is likely in North African countries under certain conditions and at some time*. This level of warning will seldom be sufficient to overcome the powerful barriers to responding to warnings or gray swans outlined above.

A notable example of the role of uncertainty in dampening warnings can be found in the US intelligence community’s attempts to raise the alarm about the potential for chaos in Iraq in 2003.³⁶ Various high-level intelligence products joined policy warnings in predicting sectarian divides, the lack of a democratic political culture, and typical post-liberation violence posed serious threats to post-invasion stability. The products were distributed to senior officials up to and including the president. And according to discussions with a number of officials who received these products, the warnings were brushed off largely because they were so highly qualified. A highly contingent warning is far easier to ignore, especially when in the grip of imperatives and wishful thinking. Of course something might happen—and then again it might not. All the warnings did, the officials said over and over again, was suggest the *possibility* that something bad could happen. And a mere possibility will not remotely be enough to overturn dearly held ambitions.

Taking gray swans seriously

The true character of the risk-related warning problem facing senior leaders and strategists, whether in business or government, is not that fate will ruin them with bolt-from-the-blue shocks they’d never considered (although such things do happen). It’s that they must constantly balance 20 or 30 or 50 major issues for judgment, each one containing substantial risks. And some of those dangers—even though they are understood, fully in line with many existing trends, discussed and sometimes debated at length—will become the gray swans that lead to disaster. Risk comes when we know in theory about something but cognitive flaws prevent us from fully appreciating them. And as we

have seen it will place in particular danger their ability to take seriously the consequences of their actions.

The primary challenge for large institutions, then, is not making themselves resilient to the arrival of totally unheralded black swans (although this does remain an important goal). What's needed most of all is a set of strategies for better analysis of and responses to gray swans. That is essentially a warning and risk management problem, and there are a number of strategies that might improve institutional batting averages versus gray swans.

1. *Cultivate a culture of valuing warning.* Arguably the single most important characteristic of organizations that avoid gray swans is their dissent-accepting culture. They make clear that they value warning, reward those who offer it, and create mechanisms to ensure that the organization cannot avoid dealing with them. From military units to companies like Berkshire Hathaway to investment firms like Goldman Sachs, there are organizational personalities characteristic of enterprises that treat gray swans seriously.
2. *Include formal risk assessment and mitigation step section in any strategy document or risk management process.* This somewhat mechanistic but still useful action could help to create a habit of dealing with gray swan issues. If the challenge of "possible but unlikely risks" is put front-and-center in any strategy process, participants may be less likely to fall victim to the human dynamics noted above.
3. *Train analytical staffs—risk managers, intelligence analysts, strategists—in the skills of conveying warning.* Offering warnings—the practical challenge of conveying risk—is far from a straightforward or easy task. Much of the literature on risk and warning has to do with large-scale public notices. Intelligence professionals are trained in warning techniques—but even they have great challenges breaking through pre-established thinking, and those concerned with risk in business or other government agencies seldom take the same conscious approach to warning. Conducting research on effective warning techniques, and training risk and strategy professionals in such approaches, could help deal with gray swans by making it less likely that the warnings will be ignored.
4. *In order to address the personal aspects of the issue, post an individual highly respected by senior leaders in key risk and warning positions.* Risk assessment and warning is a highly personalized affair. Senior leaders tend to take seriously those whom they respect; a major lesson of the financial crisis is that a perception of risk—or the lack thereof—is highly dependent on the perceived character and talent of the person or group overseeing the risk. Any position required to offer warnings and make senior leaders take notice should be staffed with leaders personally or professionally close to the organization's chief.

In a strategic landscape characterized by radical uncertainty, disaster is much more often a product of known but underemphasized risks than it is the result of totally unforeseen shocks. Large organizations would do well to cultivate the sort of resiliency required to deal with the unpredictable. But they would benefit on a more regular basis from taking seriously the peril of gray swans, and organizing their risk and decision processes to account for them. And more broadly, they would benefit from a risk process focused on outcomes.

8

Risk Becomes Personalized

Relatively early in the life of the Enron Corporation, the firm established a trading operation for oil under the direction of a man named Louis Borget. “Within Enron, he was a shadowy figure who divulged as little as possible about the details of his operation,” Bethany McLean and Peter Elkind report in their masterful history of the firm’s rise and fall. Borget’s operation began to generate substantial profits, and his confidence in his methods grew along with the profits. He sent a 1986 memo to the board in which he argued that highly trained “professionals” were using “sophisticated tools” to “generate substantial earnings with virtually no fixed investment and relatively low risk.”¹

This basic idea—that there was money to be made by moving other money around, without worrying too much about anachronistic issues like physical investments or the effective implementation of plans—would become central to Enron’s culture, and a major cause of its collapse. At the time, though, in a relatively young firm trying to make its name, the effect was magic, and Borget became known as an untouchable wonder-boy of profit. When hints began to emerge that he and his subordinates might be playing fast and loose with regulations, the worries were brushed aside. Outside auditors who tried to peer into the trading office’s books were stymied. “The Enron executives were terrified of offending Borget,” McLean and Elkind explain. A senior Enron executive sent him a love-note during the auditor’s investigation: “[Y]ou understand your business better than anyone alive. Your answers to Arthur Andersen were clear, straightforward, and risk solid—superb. I have complete confidence in your business judgment.”

McLean and Elkind’s history indicates that Borget’s “business judgment” hadn’t been so fine-tuned after all. The trading operation had made a bet on oil futures that began to go wrong, then doubled down with massive infusions of cash that raced past supposed institutional limits on the size of trades. Only the eleventh-hour intervention of a sensible Enron executive named Mike Muckleroy allowed the firm to liquidate much of the position before it

destroyed the firm. Despite this massive, last-minute effort, Enron still lost \$140 million on the transaction.

Before these risks became apparent, though, no one in the company had been willing to challenge Borget. He and his operation were simply too important. They were generating too much revenue. In the pre-disaster memo he'd received from the senior executive, the directive he had received was blindingly simple: "Please keep making us millions."²

The personalization of risk

This case, and hundreds of others like it, reflects a critical lesson of the financial crisis in respect to risk. In many leading firms, even those with dauntingly complex risk management procedures, objective analysis of risk surrendered to highly personalized judgments that gave more credence to individual perspectives and personal relationships than to data-based analysis. This pattern repeated itself in a dozen major investment firms, and throughout the market as a whole. Judgments were skewed by expectations, relationships, and beliefs that are a function of personal characteristics rather than meaningful analysis.

This effect makes itself felt in two primary ways. The first is that peoples' personality skews their perception and appreciation of risk. Personality is a major component of individual perspective, the lens through which people view events. Different people simply view risk in instinctively different ways, as a product of their personality. This is an obvious enough fact, but one that was routinely forgotten by institutions who put major risk management responsibilities into the hands of men and women who were dangerously aggressive risk-takers by nature. The basic pattern was the same, over and over again: risk procedures being subverted by supremely intelligent, hyper-confident gamblers, more than ready to roll the dice for a big payoff.

The second basic mechanism by which personality skews risk management is through relationships, and the way that peoples' perception of a risk choice depends on their view of the person taking it. Over and over again, both within firms and among them, very senior leaders judged risk, not so much by objective measures, but by their degree of faith in—and the perceived success of—the people managing the investments. When an investment manager gained a reputation as fiercely smart and highly aggressive, and when their bets appeared to be paying off handsomely for the firm, risk management all too often went out the window. They became darlings of senior leadership, and almost untouchable. Any doubts about the risk inherent in the positions they were building up were roundly dismissed.

Here again we see a hallmark of the constraints imposed by a context of uncertainty. It is tempting, and common, to think of risk as an objective function, a property that can be discovered and shared, the outcome of an equation. For some issues, as we have seen, this is indeed the case. But when

we are operating under uncertainty, risk becomes the plaything of personal interpretation—it is what we make of it. And one important result, a major lesson of the financial crisis, is that a leading tool or crutch that senior leaders use to manage uncertainty is to put store in personal factors, or to trust their own highly idiosyncratic perspective on events. We find this again and again in the cases of major financial firms affected by the financial crisis: the displacement of objective, procedural assessments with eccentric judgments about favored individuals.

The result is to create another set of human factors that conspire to undermine procedural risk management. As Andrew Ross Sorkin has written about the crisis, ultimately, “whether an institution—or the entire system—is too big to fail has as much to do with the people that run these firms and those that regulate them as it does any policy or written rules.”³ The effect of personality has the potential to overwhelm risk management processes: Many leaders trust their own instincts and relationships more than the findings of formal risk management mechanisms.

Personality as engine of interpretation

A major lesson of the crisis is that personal style and characteristics shape how leaders see the world, including risk.⁴ Aggressive people will be more risk-accepting; conservative leaders will be more attentive to risk. This is obvious enough, but a key lesson of recent experience is that, under uncertainty, personality can have decisive effects on how organizations—even those with elaborate risk management procedures—handle the issue.

As we have seen, decision-makers must impose meaning onto ambiguous events. One leading engine of this imposed meaning is personality. A paranoid tyrant like Josef Stalin will interpret any actions of a potential competitor in threatening terms. A relaxed and patient leader like Dwight Eisenhower will have more empathy and understanding for another side’s perspective. Intense rationalists, like John F. Kennedy or Richard Nixon, will be able to engage in sophisticated sophisticated signaling games with adversaries assumed to be playing a grand sort of chess match, and will interpret actions as abstract gestures. History is full of examples of decision-makers viewing reality through a refracted prism of their personality and its associated biases and expectations.

One of the challenges with assessing the role of personality is that decision-makers do engage in rational calculations that justify their interpretations—but this isn’t evidence that those interpretations are objective. The historian Orlando Figes, for example, describes the attitude of Russian Emperor Nicholas I in the run-up to the Crimean War. Nicholas was an impetuous figure, characterized by “impulsiveness and rash behavior and melancholic irritability,” who “behaved at times like a reckless gambler who overplays his hand.” And yet Figes asks whether Nicholas really thought, from the

constrained lens of his personality driven perspective, whether he was gambling at all. Russia's ability to fight off Napoleon had demonstrated that Moscow could fend for itself, making Nicholas overconfident. But this justificatory leap "was not a reasoned argument," Figes concludes. "It was not based on any calculation of the armed forces at his disposal or any careful thought about the practical difficulties the Russians would face" in a war against Europe's leading powers. "It was a purely emotional reaction, based on the Tsar's pride and arrogance, on his inflated sense of Russian power and prestige, and perhaps above all on his deeply held belief that he was engaged in a religious war to complete Russia's providential mission in the world."⁵

Personality, in other words, combined with visionary worldviews and mixed with a rich broth of cognitive bias, can lead decision-makers to interpret events in skewed and misleading ways. A thousand case studies of strategic judgment, as well as the anecdotal view of most people who have served at senior levels in government and business, would suggest that this conclusion is so obvious as to be self-evident. And yet the critical fact for our purposes is that formal risk management procedures assume that it is false—or assume, at least, that a relatively simple process for risk assessment and mitigation can overcome this fundamental influence on interpretation and judgment. The experience of the financial crisis suggests that the opposite is more nearly the case.

The effect of personality in fomenting a crisis

The history of the financial crisis is brimming with ambitious, confident, belligerent senior leaders in the financial sector whose personal style shaped their approach to risk. Based on their natural inclination, some leaders were simply more risk-accepting and less concerned about dangers of their favored courses of action. Lehman chief Richard Fuld was a leading example—brash, confident to the point of recklessness, dedicated to exceptionally loyal but not always completely competent underlings⁶—but there were dozens like him. Their default inclination was to push for profits without attending specifically enough to risk. They had the personality of risk-takers; the term "gambler" crops up again and again in histories of the crisis. Roger Lowenstein describes the aggressive Washington Mutual chief Kerry Killinger as aggressively pursuing risky subprime loans, someone who "pushed aside more cautious bankers, replacing them with gunslingers."⁷

When the determined John Mack took over Morgan Stanley in 2005, he was anxious to shake up what had become a somewhat relaxed culture and force the company to challenge its rivals in new areas of financial innovation. Mack "promised to bring back the old aggressive culture of Morgan," the journalist Scott Patterson has chronicled. Mack "had a taste for risk. Morgan, he believed, had lost it." In his inaugural talk to employees he emphasized that Morgan was being left behind by Goldman and Lehman, and that "the new

paradigm for investment banks on Wall Street was risk taking." In particular he had in mind banks that not only advised and managed money, but put their own capital into the market in service of profits. Mack laid out audacious targets for revenue growth; "Nice idea," another Morgan senior executive thought, but they needed a plan to do it. "The answer, he feared, was simply to take more risk."⁸

At Bear Sterns, a leadership group—including eventual CEO Jimmy Cayne—that was smart and savvy about risk was nonetheless also aggressive, ambitious, and determined to win in the hyper-competitive investment bank industry. As the financial industry began boasting mammoth revenues and profits in the 1980s and 1990s, Bear's leaders increasingly pressed their traders to accept more risk, and praised and promoted managers who generated the biggest profits. Men like Cayne and "Ace" Greenberg were supremely confident adventurers, combative executives who'd clawed their way to the top without exclusive pedigrees, and their corporate culture came to reflect these personalities. One example, chronicled by William Cohan, was Howie Rubin, an investor who ran into trouble at Merrill Lynch for exceeding his investment limits and taking too much risk. Embracing Rubin reflected Bear's "propensity to be comfortable taking risks even with people who might be colorful," one senior executive told Cohan.⁹

The role of specific personalities can become joined to more general, role-based perceptual lenses. Research suggests, for example, that the more senior an executive is, the more risk-accepting she or he is likely to become.¹⁰ This makes sense from a number of angles: More experienced executives are likely to trust their own judgments more readily and with less analytical support, for example. They may also believe that their continued ascent to the rarified senior-most positions demands bold, risk-taking ventures. The natural tendencies in organizations are thus likely to favor the selection of people like Cayne, Greenberg, and Rubin, and create a context full of aggressive risk-takers.

It should not come as a surprise that many such examples crop up in the Enron case. Jeffrey Skilling was a visionary with the blind faith of an ideologist, "enamored, always of the Big Idea, with surprisingly little appreciation for how things got done in the real world. He had zero interest in the nuts and bolts of operations." He "largely disregarded—indeed, he had an active distaste for—the messy details involved in executing a plan." He was supremely confident from the start and became more so with Enron's seeming success. All of this produced someone blithely ready to take massive risk: "For all his analytical abilities," Bethany McLean and Peter Elkind explain, "he was a gambler at heart. ... He always assumed that he could beat the odds." Another senior Enron executive, Rebecca Mark, had an "inherent optimism" that "led her to push forward where others might at least have hesitated"—and to sideline dissenters. She "trusted her gut far more than any spreadsheet," and her gut just about always told her to go for it.¹¹

Organizations need such risk-taking personalities to remain aggressive in seeking opportunities. Problems emerge, though, when this inclination is unchecked by institutional procedures—when risk-takers are allowed to bet the future of the organization on highly perilous and highly leveraged strategies. The crisis, in this sense, is partly about the intersection of personality and procedure, just as it is about the larger intersection of risk management processes and human factors more generally.

A typical aspect of aggressive risk-takers was the belief that they had figured out the problem, and anyone who doubted them was a fool or a dissenter. Before Bear's collapse, when some of its lenders questioned some of its practices and trades, a senior executive said, "You guys don't know what you're talking about," which Brian Burrough describes as a "classic display of Bear-style arrogance."¹² We see the same basic pattern in dozens of other cases, whether the role of financial leaders like Fuld or Cassano or national security cases involving people like Donald Rumsfeld and 1960s CIA leader Richard Bissell—an unquenchable sense that they had the answers, they were smarter than most people around them, and they could intuit the risks of most situations better than most processes.

A related factor is the desire of ambitious senior leaders to prove themselves, which can obstruct objective risk analysis. Many case studies describe aggressive senior leaders whose outward confidence and brashness conceals a powerful, sometimes desperate, inner desire to demonstrate their worth through accomplishments. Research shows that "people with a fixed mindset wish to broadcast their abilities," with the result being "biased attention toward events that confirm your superiority and away from events that do not. And this is how a sense of infallibility sets in."¹³ Repeatedly among financial leaders we see people determined to make a mark, to prove their superiority—again not an unimportant trait for competitive organizations, but one that must be channeled and carefully controlled to avoid problems.

Aggressive personalities tend in particular to be overconfident, displaying excessive degrees of a very typical cognitive bias affecting decision-makers. A clear lesson of the crisis is that certain personalities will be far more subject to wishful thinking, cognitive dissonance, and overconfidence, as will leaders with more experience and demonstrated track records of success. Such biases can make investors and CEOs sure that they are smarter than everyone else and getting a bargain even while driving prices to unbelievable heights in a bubble.¹⁴

A big part of the story behind the AIG FP collapse was the new leadership style of its chief Joe Cassano. Whereas the former CEO, Tom Savage, was a quant himself who grasped the risks they were running, "Cassano knew a lot less math and had much less interest in debate." Lewis describes him as "a guy with a crude feel for financial risk but a real talent for bullying people who doubted him," a man who reacted violently to the "faintest whiff of

insurrection." The "culture changed," an AIG employee told Lewis; the "fear level" kept people from speaking honestly about risk. Many of those willing to stand up against him left the firm. Lewis concludes that many who knew him said his main flaw was "a need for subservience in others and an unwillingness to acknowledge his own weaknesses." He believed his own claims, though, which made it tough to see through his bluster.¹⁵

In the case of the parallel financial crisis that unfolded in the tiny (but ultimately hugely leveraged) nation of Iceland, mostly young people without significant experience in investment blasted into the market and started making huge trades—and not only that, but traveling the world and instructing others how to do their business. They declined to do even basic due diligence about the companies they were buying. And the process was in part a product of personalities: Many of Iceland's young dealmakers believed in the "natural superiority of Icelanders." And yet, as one of the targets of their buying spree put it, "They were all highly educated people."¹⁶ This was exacerbated by a form of groupthink, with a profound homogeneity prevailing within the Icelandic financial industry. They thought similarly and became caught up in the same mania. And they come, Michael Lewis has argued, from an aggressive culture of risk-embracing fishing. "They are gamblers," one economics professor wrote in the 1950s about the mindset of fishermen, "and incurably optimistic."¹⁷ These same terms crop up again and again in histories of the financial crisis (and related crises, such as the Long-Term Capital Management or Enron collapses): gamblers, cowboys, risk-takers.

Overconfidence is a general problem of human decision-makers, and senior executives in particular. One of the biggest problems in risk-related calculations is that "people are typically very confident about judgments they make," in part because they do not want to believe how tentative their judgments have to be under uncertainty.¹⁸ Or as Daniel Kahneman puts it, there is a "pervasive optimistic bias" in human thinking. It is a form of motivated or wishful thinking, of a specific type: the tendency to be too confident in the quality of our judgments and plans. In surveys, over 80 percent of entrepreneurs believe they have better than a 70 percent chance of success; fully a third thought there was no chance at all that they would fail. Confidence, Kahneman explains, is more a product of the level of belief in "the story they can tell about what they see" than it is a result of objective assessment of facts.¹⁹ And a major lesson of the financial crisis and related business calamities, as well as of a raft of national security disasters, is that experienced, highly self-confident executives can tell themselves wondrous stories.

At Enron, overconfidence was baked into the very culture. As Bethany McLean and Peter Elkind have written, possibly the most fundamental attitude within the company was that "Enron people were simply better than everybody else. At conferences, [COE Jeffrey] Skilling would openly sneer at competitors." This degree of arrogance extended to risk analysis, thinking of

ways things could go wrong for the company. When an outside consultant “tried to ask Enron employees where they thought they were vulnerable,” all they “got in return were blank stares. People at Enron simply didn’t believe they were vulnerable.”²⁰

Overconfidence also tends to produce a willingness to double down on risk when things start sliding toward the abyss. In the LTCM case this tendency was pronounced. Leaders of the firm stubbornly held fast to their worldview despite evidence to the contrary, and bid up their leveraged positions in all manner of securities, even some far afield from their original expertise and focus on bonds. As things kept getting worse, the partners kept expecting things to return to “normal”—that is to say, to conform to their expectations—and this view caused them to believe that further investments along the same line could only reap more profits.²¹ The same tendency crops up in the Enron crisis and a number of others—a rigid conviction that a firm’s strategies have got to be correct, leading to throwing good money after bad once a position begins to turn bad. It has to turn around, they believe—they couldn’t be that wrong.

“When they come together,” Daniel Kahneman concludes, “the emotional, cognitive, and social factors that support exaggerated optimism are a heady brew, which sometimes leads people to take risks that they would avoid if they knew the odds.” Confidence, he explains, is a *feeling* rather than the outcome of analysis. When decision-makers indicate that they are very confident of some outcome, he writes, this primarily indicates that “an individual has constructed a coherent story in his mind, not necessarily that the story is true.”²² And the lessons of the crisis suggest that certain personalities are uniquely vulnerable to such skewed perceptions.

Finally, personalities are related to worldviews, which provide a powerful lens onto events. A powerful example from the crisis is Alan Greenspan, who depended heavily on statistical models built on the sand of his free-market ideology. His “empiricism” and use of modeling, John Cassidy has written, was married to a “fervent belief in the efficiency and morality of the free market system.”²³ Even Greenspan himself, in his post-crisis effort to take stock of where he went wrong, admitted that while modeling in nonfinancial sectors “has worked tolerably well,” in part because they rely on more linear, engineering, and physics models of causality, finance is different. “Our propensities related to fear, euphoria, herding, and culture, however, virtually define finance.”²⁴

Personality and relationships

A second and perhaps even more important aspect of personality in risk decisions is the role of relationships and the reputation of particular individuals in organizations. When a given manager or leader has a certain standing with the CEO, or when they are seen in particular ways throughout the organization—either greatly respected or generally feared—risk management

will become highly skewed. A major lesson of the financial crisis is that, all too often, judgments about risk become judgments about individuals. And when senior leaders have favorites, as is often the case, their appreciation for risk will become skewed by their view of key subordinates.

This tendency is partly the result, once again, of the looming context of uncertainty. Without an objective standard to use, senior leaders are tempted to outsource risk assessments—to rely on some trusted agent who can “assure them” about the status of risk on certain issues. “I trust [a given manager, risk analyst, or leader], and they are telling us X” is a common refrain from leaders in the run-up to crises.

At Merrill Lynch, the role of the trader Osman Semerci in creating the high risk levels was in part a function of those who had promoted him—the senior leaders of the firm, Stan O’Neal and others—and the perceived political risk of taking him on. Reporter Greg Farrell found that “Because Semerci was perceived as being [CEO Stan] O’Neal’s guy, his decisions and trading activities did not receive the pushback or scrutiny they should have.” Moreover his own aggressive style created an environment in which risks could balloon. Farrell has described him as “a salesman with the instincts of a riverboat gambler.” He arrived under the stress of a self-made challenge—promises to his boss that he could dramatically ramp up income from his division. More broadly, O’Neal was known as a gruffly confident leader who didn’t like to second guess his decisions. He “tuned out” executives who disagreed and ruled with an “imperial” style, according to Greg Farrell.²⁵

In the case of LTCM, the senior partners had close relationships with many other leaders on Wall Street. They golfed and drank and played cards with them. Perhaps even more important, the founding partner, John W. Meriwether, was widely respected and considered someone of good and cautious judgment—someone who knew how to manage risk.²⁶ In that case, it was the relationships between senior members of the firm and outside organizations that disrupted effective risk assessments.

Skewed risk assessment based on personal relationships is also a huge part of the Enron story. Jeffrey Skilling was viewed as a genius who was reinventing a whole industry. Other Enron staff who spoke of Skilling to Bethany McLean and Peter Elkind did not just call him smart—they used “phrases like ‘incandescently brilliant’ or ‘the smartest person I ever met.’” When things started to go wrong, people were more mystified than angry: They could not imagine that Skilling would have screwed up too badly. “I’m not necessarily long Enron,” one investor said at the time, “but I’m long Jeff Skilling.”²⁷ This was precisely the problem, in a nutshell: The perception of risk was bound up in the leaders taking it rather than a truly objective analysis.

One hallmark of this process was that, once really cynical analysts looked past the leaders involved and dug into the hard numbers, it wasn’t all that difficult to determine that Enron was essentially a highly ornamented Ponzi

scheme. McLean and Elkind tell the story of Jim Chanos, a young short-seller who began investigating Enron in 2000, specifically its broadband unit. He quickly decided that there was something deeply wrong. "They were chewing up capital," he said, and the business model didn't hold up to scrutiny. "As soon as anyone looked, they could see the stuff we saw."²⁸ But others weren't looking, and even when confronted with the analysis of Chanos and other skeptics, the market—and Enron leadership—dismissed the warnings. They did so in part for many of the reasons outlined in this study: cognitive bias, the incentive structure of the industry, a sense of urgent imperatives. But they also did so because their analysis of risk was blinded by their faith in the brilliant personalities at the firm.

Another example comes from Fannie Mae. Treasury Secretary Henry Paulson and Fannie Mae chief Jim Johnson had worked together before at Goldman Sachs. Johnson had been a sort of indirect supervisor of Paulson's.²⁹ This sort of relationship was not going to make for objective analysis. Within the industry, similar ties also discouraged objective analysis: Fannie Mae CEO Jim Johnson was reportedly close to Countrywide chief Angelo Mozilo, and the two hatched collaborate plans whereby each became complicit in the other's risk. Meantime Countrywide was later found to be granting "VIP Loans" at reduced rates to senior government officials and other leaders in the financial world.³⁰

Senior decision-makers should be thinking of risk objectively. But it should not be a surprise that at the highest levels it acquires a human face—*who* is asking you to take what risk. Senior leadership is as much about relationships and trust as it is about any objective criteria for judgments, and this general flavor comes to characterize the approach to risk. In fact a good deal of recent literature on business strategy and entrepreneurship magnifies the problem. What's important with a new company, dozens of business books suggest, is not so much any specific idea. It's getting the right management team in place. Investors often say they don't invest in a product or a business model—they invest in a team of people.

This makes sense, at one level. Talent is a good indication of future success, regardless of the momentary idea. Business strategy studies are filled with examples of brilliant innovators who failed repeatedly before their big success—but failed in interesting and creative ways that paved the way for their ultimate victory. All along, their talent was obvious, or should have been, to any potential investor. (Professional investors, of course, especially after a decade of recent experience with visionary but operationally suspect entrepreneurs, look well beyond creativity these days in their search for the right team.) But the dark side of this faith in the right people is its frequent inability to distinguish aggressive profit-seekers from risk-defying gamblers, and its recurring ability to blind a market to risks accumulating in a firm whose leaders are viewed as business savants.

Lessons for risk management in national security

The personalization of risk exacerbates the challenge, described in previous chapters, of the subjective character of risk and risk management. Partly the tendency to view risk through the character of the people taking it is a product of uncertainty and interpretation, but it is also an engine of such patterns. Thinking in relational terms is natural to human beings and magnified by the advice of current management theory, which is all about getting “the right people on the bus,” as Jack Welch puts it. The result is a substantial tendency to subordinate objective assessments to personal ones.

Most fundamentally, these lessons reinforce the importance of the role of individuals in shaping outcomes. As Daniel Byman and Kenneth Pollack have argued, traditional theories of world politics tend to downplay the influence of specific personalities in favor of abstract forces of history.³¹ They contend that this is a mistake, that we cannot understand most outcomes in world politics without comprehending the perspectives, personalities, and biases of specific leaders. History is made by particular decision-makers operating within the general context, not merely by the context itself. The role of personality in the making of the financial crisis strongly supports this contention. It would argue for approaches that bring the role of the individual back into national security theories.

An obvious corollary is that, if risk perception is governed by personality and relationship factors, risk management must make those variables work for it rather than against it. A single individual with moral authority inside an organization and in the eyes of senior leadership will be more effective in conveying risk than the most elaborate process. Because risk management is in the end a human process depending for its success on the conquest of human factors, someone who enjoys great respect in the eyes of the CEO and whole organization will be a more credible messenger of risk than a hyper-organized mechanism. The relationships between leaders and risk analysts are the key to effective warning, and that issue should be borne carefully in mind when designing a risk management process.

The same line of thought suggests strategies for dealing with the gray swan problem outlined in the previous chapter. Evidence from many strategic judgment settings, including the example of financial firms leading up to the crisis, suggests that getting warnings heard is far more about having the right people offer them than it is about getting the right data to support them. This again points to the critical importance of placing trusted, respected senior leaders into positions with clear responsibility—and accountability—for not only assessing potential risks but also generating effective mitigating actions.

It may be, in thinking of decision-making through personality lenses, that certain types of personalities are more disposed to make these sorts of errors. There is some evidence that thinkers who are “hedghogs”—who know “one

big thing,” people who view the world through generally ideological lenses—are less good at analyzing risk than “foxes,” people who analyze events through many different lenses. Hedghog thinkers get stuck on norms and values in opposition to evidence.³² There is some evidence that information-processing styles significantly affect risk-taking behavior, and this points to the importance of choosing people for risk management positions with great care.

Finally, the role of personal factors once again suggests the importance of a risk process focused on outcomes. If the financial crisis is any guide, the places where personality factors—both the biasing effects of personal traits and of interpersonal relationships—can have the most damaging effects is in their effect on an institution’s ability to assess clearly the potential consequences of its actions. Again and again in the lead-up to the crisis, financial organizations did not adequately assess and question the risks involved in particular strategies, either because the responsible division managers turned out to be highly aggressive gamblers or because corporate leadership put more faith in personalities than institutional processes of oversight. Targeting risk management especially at consequences would go some distance to counteracting such damaging effects of these most common human factors.

9

What You Don't Know Can Destroy You: Ignorance and Correlated Risk

The story of many of the firms that fell prey to the financial crisis reads very similarly in many respects. Institutions that experienced risk disasters were misguided by wishful thinking, overconfidence, groupthink, the skewing effect of personality, and much more. The case studies of these companies have at least one more important thing in common: Senior leadership lost touch with what was going on in one or more of their business units. Senior officials who ought to have been responsible for risk—whether unit heads, CEOs, or board members—simply did not pay enough consistent attention, or educate themselves sufficiently about the nature of the risks being taken, to exercise proper oversight. Once that happened, it was often just a matter of time before disaster struck.

Effective risk management presumes knowledge of the risks the firm is trying to manage. Yet all too often in the lead-up to 2007, senior leaders of giant financial firms allowed vast swaths of ignorance to arise about what was going on in their own institutions. Their awareness of the details of specific financial instruments was often slim to nonexistent. In many cases, board members and top officials in key financial firms were not even aware of the scale of investments, or degree of leverage, being pursued under their own roofs.

The story is much the same, as we will see, in national security. When senior leaders lose touch with facts and developments in their organizations, risk calamities occur. This is the pattern in the invasion of Iraq, the Bay of Pigs, the Soviet invasion of Afghanistan, and a hundred other case studies of risk failure. The lesson both reinforces and complicates the case for rigorous risk management procedures. They are essential to ensure that accurate and sufficient information reaches senior leaders. At the same time, the temptation for such leaders to distance themselves from the nitty-gritty of daily operations turns out to be very great, for reasons both sensible and troubling.

Ignorance and the financial crisis

The danger of ignorance of risk emerges when senior leaders do not pay enough attention to what is going on in their organizations. There can be various reasons—faith in a certain subordinate, loss of focus, simply feeling too busy. But the result is always the same: to leave some manager or division unchecked in its embrace of risk.

Time and again in these cases, the road to disaster was paved when an individual or business unit went off on their own and made fantastically risky decisions that the overall organization refused to take seriously. In a few cases leaders of the organization literally didn't know, but in many others they knew and used some excuse to allow for it: The trader is a genius, the specific market is safe, the division has been on a two-year winning streak.

This was certainly true in the case of Merrill Lynch, where a single, poorly supervised head of an investment unit—Osman Semerci—was allowed to create massive, risky positions in derivatives tied to bad mortgages. The information on their activities was kept so secret in the months before the fall that the portfolio of investments was called the “Voldemort book.” Only when curious managers from other areas of the firm started poking around the investments did they create a cascading effect of awareness and, finally, some degree of oversight.¹

Roger Lowenstein argues that many CEOs like Lehman chief Dick Fuld were “of too early a vintage to grasp the nuances of newer, exotic securities.” Robert Rubin at Citigroup learned more about what his traders were doing—but not all; just enough to reassure himself that it was defensible. “This half-knowledge was potentially lethal,” Lowenstein argues, because Rubin allowed himself to be mesmerized by the quantitative prowess of his modelers. As had happened in the Enron case, moreover, in some cases the CEOs were flabbergasted when they finally got the news. When Merrill Lynch CEO Stan O’Neal heard—once the market had already started to curdle—that his firm had amassed almost \$50 billion in credit default swaps, he was “stunned.”²

At Bear Sterns, lack of awareness of the company’s burgeoning positions in complex derivatives was built into a number of aspects of the company’s operations heading into the crisis. When a highly leveraged fund collapsed, senior leaders later claimed they didn’t know the positions the lead trader was taking. A senior Bear official told Brian Burrough that it was “one of those things where everyone thought someone else was paying attention.”³

As chronicled by William Cohan, one key manager of Bear’s risky CDOs, Ralph Cioffi, was a smart, well-respected sales leader who was simply not known for management skills or persistent attention to detail. As at so many firms, senior leaders, in part looking at Cioffi’s stellar track record, stayed mostly aloof from the details. “None of them really understood what he did,” another Bear executive told Cohan. The financial instruments he was

building were stunningly complicated, and it took substantial effort to grasp them. CEO Jimmy Cayne “had only a vague understanding” of the derivatives, Cohan explains—and a key manager named Warren Spector, one of the few really senior people who understood the new investments and had a penchant for watching the fine details, ran afoul of Cayne. He was sidelined, and his expertise and attention were as well.

This example illustrates another key similarity in so many cases of risk disasters: One or two critical people, senior enough to force attention onto risk and predisposed to do so, leave or are pushed aside. Sometimes the issue is simple turnover; in other cases, it is the perceived disloyalty of dissent that causes them to be ejected from the firm. Either way, in a situation of broad-based ignorance when only a few people really know what is going on, an organization is always sailing at the edge of disaster—and one or a few especially meaningful departures can dissolve the last vestiges of meaningful oversight.

Management’s lack of understanding of what was going on in the firm was then often magnified by another common symptom of risk crises: misrepresentations and outright lies as the firm tried desperately to reassure an increasingly skeptical market of its health.⁴ In a wide range of cases, certainly including Enron but also AIG and many others, senior officials in risk-taking divisions simply concealed what they were doing, in whole or part, from some combination of senior leaders inside the firm or independent oversight bodies outside it. When people are actively trying to hide the level of risk being taken, it will be difficult for any risk management process to do its job.

In some cases the issue was not intentional concealment but simply the malign effects of an extreme stove-piping of information. In the case of Merrill Lynch, for example, the risks accumulated by Osman Semerci’s financial products division were hidden because it simply didn’t share data with others in the firm. Risk management processes have to be diligent enough in sharing information to ensure that no element of an organization can conceal key data about its ongoing actions.

In the Enron case a consistent theme is the substantial ignorance that senior leaders developed with regard to the operations of the company. What is sometimes forgotten is that Enron was characterized by many of the same kind of incredibly obscure financial instruments that would help spur the financial crisis. Though their character may have been different, the use of complexity to achieve somewhat artificial profits without taking seriously underlying economic realities was very much the same. And in the process, top leaders became blissfully unaware of what was actually going on.

Enron chief Ken Lay, increasingly playing the role of a philanthropic titan, paid less and less attention to the firm and “lost touch with the company’s business.” His appreciation for how key programs made their money “was fuzzy at best. He usually seemed at one remove from the nitty-gritty of the business itself.”⁵ Right from the beginning, the oil trading unit within Enron

operated without much oversight, in part because senior executives didn't really understand what was going on. Sometimes the whole goal seemed to be to create projects of such byzantine complexity that no one could grasp their true nature. (Because some of them were essentially illegal, that's not necessarily surprising.) Right at the end—when the dissenter Sharron Watkins had forced Lay and Skilling to confront what was going on—Lay himself “seemed concerned and surprised at what he was hearing.”

Another lesson of the crisis is that oversight of risk-taking, and leadership messages on risk appetite, can become especially unglued during transitional periods. When Hank Greenberg was forced out of AIG, according to Michael Lewis, his successors wanted to prove they could generate big profits without him—“and turned a blind eye to all sorts of risks.” The AIG collapse begins, Lewis explains, “with a change in the way decisions were made, brought about by a change in its leadership.”⁶ Yet the change didn't spark enhanced oversight, because no one was paying close attention to the implications of the leadership transition.

The dangers of highly complex instruments of strategy

A primary cause of such ignorance was the intricacy of the financial instruments being developed in the financial industry, which had become so complex that senior leaders simply didn't understand them. Before the crisis the CEO of AIG insurance, Ed Liddy, had not deeply investigated the dense network of derivative trades being made by his finance arm; the post-Hank Greenberg leadership of AIG had “no idea” of the risks being run by their investment people.⁷ Fed chief Alan Greenspan would later admit to being “bewildered” by some of the complex derivatives that lay at the core of the crisis. “Even the CEOs of the firms that sold these products,” Andrew Ross Sorkin writes, “had no better comprehension of it all.” The two top Lehman Brothers executives, Sorkin reports, did not fully understand the instruments their own company was marketing, “and showed remarkably little interest in learning more.”⁸

Because of this complexity, most big risk decisions, or collections of them ended up amounting to one or two big judgments—guesses—rather than a true calculation of all risk elements. In the case of AIG FP's bets on mortgage-related derivatives, many of those involved told Michael Lewis they were “shocked at how little actual thought or analysis” underpinned the strategy. “It was simply a bet that U.S. some prices would never fall.”⁹ In the case of the financial bubble and subsequent collapse in Iceland, part of the problem was that the financial sector grew so quickly, it vacuumed up people with no background in the field—and was overseen by government officials whose careers spanned areas like veterinarian and philosopher, but not economist. They then turned around and rejected advice from trained specialists.¹⁰

Much the same was true with the earlier case of Long Term Capital Management. The trades they were involved with were just too complex for many people to understand, including institutional investors who pumped cash into the fund's coffers. They had to take them on faith, based on the reputation of the firm's leading traders. (Here again we see personality trumping objective risk assessment when risks simply can't be adequately assessed.) Each bank investing with LTCM knew only of its own positions and trades but not the whole picture. And the problem of incentives intersects here as well: "The bankers were too busy making money," Roger Lowenstein writes, "to bother about the risks or the shoddy disclosure in this fast-growing business."¹¹

Once a crisis hits, this uncertainty and lack of knowledge is magnified. Accounts of the Asian financial crisis, Paul Blustein has written, "do not adequately convey how frightening, disorderly, and confounding it all was, most notably for the people in charge of quelling it." Policy-makers in the affected countries and at the IMF "found themselves overwhelmed and chastened by the forces unleashed in today's world of globalized finance." Participants, he writes, "recall with anguish having been thrown into the midst of crises with bewildering origins and no obvious solutions."¹² The same could be said of the 2007–2008 crisis; Blustein's powerful phrasing nicely captures the feel of most inside accounts of the recent financial crisis.

Missing specific types of risk: accumulating and systemic risk

The crisis suggests that risk appreciation was especially lacking—and ignorance was especially devastating—in relation to two particular categories of risk. One is risk that emerges gradually, in accumulating positions; and the other is systemic or correlated risk. Both of these issues are stand-alone risk problems, but they caused damage particularly when combined with a lack of attention, oversight, and understanding from senior leaders.

Risk hides in accumulating positions

The lack of awareness or attention allowed massive, dangerous positions to accumulate without a clear sense of the growing risk. In AIG FP, for example, the initial investment positions were relatively sound. But over time the continuous application of the same strategies led to massive quantities of insurance on questionable mortgages, accompanied by huge leverage. Further deals were "simply rubber-stamped" up the line "on the theory that this was just more of the same."¹³

The general pattern would be repeated across the financial sector in a host of areas. An initial investment would be made, or new business started—complex derivatives, mortgage-backed securities built on subprime loans. An initial sense of risk around the decision would settle into place. And then, lured by profits or competitive pressure, the firm would continue building

its portfolio until it harbored massive amounts of risk. And along the way, there was never a point where a particular threshold had been crossed that made clear to the companies that they needed to fundamentally reevaluate their risky behavior. The assumption seemed to be that, if ten units of a risk rated at moderate seemed fine, a thousand units should not necessarily be any different—but of course, a different way to look at it is that risk was a function, not only of the estimated risk level of each investment, but also of the size of the position.¹⁴ If every unit involved a potential liability of a million dollars, then a thousand or a million examples would be far riskier. Risk processes in major companies were designed to catch accumulated dangers, but they failed to do so in a reliable manner.

The underlying question in estimating accumulating risk is one of thresholds or tipping points. At what level does the accumulation of risk become qualitatively more dangerous, and why? Instruments such as Value at Risk are designed in part to measure the total risk profile of a firm's investments, but in a number of firms the level of scrutiny given to more deals of a category that had already been evaluated was relatively poor—especially if they were generating significant revenues, and particularly if the leadership of the relevant division was well-respected. It can be difficult to create an objective measure for the point at which a firm crosses some line of risk. Absent a clear standard, the incentive to keep loading up on more revenue-generating investments can be overwhelming.

Systemic or correlated risk

But the most common form of risk that senior leaders failed to heed was that of correlated or systemic risks. All too often, corporations treated individual decisions as if they had no wider ramifications. They were one-off risk judgments, whose potential consequences were limited to the single investment or product. There was little attention to the systemic interdependencies involved, or awareness that the financial system—with so many firms in so many ways tied to the fate of mortgage-backed securities—had become like a giant Jenga tower, with every risk event amounting to the withdrawal of another block. Sooner or later, the interdependencies would bring about collapse.

The firms involved should have known better, because there had already been extensive experience with systemic risk in the financial sector. Before the ruin of Long-Term Capital Management, the architects of what was for a time a fantastically profitable firm believed that their models could accurately forecast aspects of the bond market. Risks were low, in part, because they believed they'd diversified: While the firm had large positions in a few countries, its risk was dispersed. If something went very wrong in one context, the firm would be hedged through its positions in others.

The problem was that when a financial crisis struck—built around the collapse of the Russian ruble—it wasn't just one or two countries or securities

that began to go haywire. It was all of them. A major assumption of the firm's models was that they were fully diversified—they'd invested in various bonds and related bets (and ultimately other securities) around the world. The problem was, they made essentially the same bet, on the ratio of bond "spreads" and their likely trajectory, in just about all of their trades. When a general financial crisis attendant to Russia widened the spreads, they were sunk.

John Meriwether, LTCM's founder, admitted in a note to investors that "our losses across strategies were correlated after the fact"—but in truth, as Roger Lowenstein has pointed out, the fact that they were tied together could easily have been understood from the beginning. A US Fed official named Peter Fisher saw this immediately once he got a true look into LTCM's books during the crisis. "Long-Term's trades were linked," Fisher realized (according to Roger Lowenstein), and "they had been correlated before the fact." A Treasury official thought, "During a crisis, *the correlations always go to one*. When a quake hits, all markets tremble. Why was Long-Term so surprised by that?"¹⁵

A decade later, this same phenomenon—individual risks quickly accruing into calamitously system-wide properties—proved the undoing of many financial firms that thought they were assessing their discrete risks with great diligence. The "new ultra-connectedness" among financial institutions, Andrew Ross Sorkin has written, created dramatic new risks. As the crisis broke, senior government officials saw right away that the systemic risks were the most dangerous ones.¹⁶ Firms should have been aware of the risk, not only from the LTCM case but also from the Asian financial crisis in the late 1990s—an example of a generalized contagion in which crises in some countries led to crises in others, without always a clear rationale.¹⁷

In the AIG FP case, for example, the firm absorbed thousands of global risks thought to be uncorrelated. As Michael Lewis put it, how could it happen that "all sorts of companies and banks all over the globe would go bust at the same time?" The idea was that their portfolio was intendedly diverse, and "it was unlikely to all go bad at once." Even when things started going downhill, AIG Financial Products head Joe Cassano doubled down, figuring diversification meant that not everything could continue downward.¹⁸

A key lesson of the crisis, therefore, is that keeping a handful of discrete risks at bay won't be sufficient if they can snowball into a larger one for which you are not prepared. System-wide risks can dwarf discrete, idiosyncratic ones, especially in a situation of high leverage. As the firms found out in the crisis, typical risk management tools target idiosyncratic risks not systemic ones. The problem is that leaders are often surprised by this factor—by the tendency, *in some circumstances*, for risks to be correlated.

Organizations should be increasingly aware of such dangers. Especially today, in an era of interdependent ecologies and economies and interconnected information networks, when risk arises in one place, it tends to cascade. A collapse of one financial market rapidly attacks others. Water

shortages in one country can easily become more widespread. The germ of extremism can spread across boundaries and become a global phenomenon. And it was a lesson, again, that should have been driven home by the LTCM crisis: Just after that event in 1998, a Federal Reserve board member gave a speech in which he suggested that the LTCM example “should cause all of us to reassess our practices and our views about the underlying nature of market risks,” by which he meant in particular the perils residing in correlations.¹⁹

Ian Goldin and Mike Mariathasan argue that the growth in complexity of global interactions has boosted the systemic risk in fundamental ways. There is more than we cannot know, and each action can have more wide-ranging and cascading consequences. They view the international system as a “living organism,” effectively built on interdependencies. The biggest risk, they suggest, is not the vulnerability of any one of these areas but “our lack of capacity to manage the growing complexities and interdependencies between them.”²⁰

Goldin and Mariathasan define systemic risk according to a few leading characteristics. It can be a major shock “triggered when relatively modest tipping points, breaking points, or regime shifts hit their thresholds and produce large, cascading failure.” Systemic risk can also be manifest in effects translated through a network through some form of contagion. And it can be a “common shock” to many elements of the system produced by some indirect cause that affects many elements of the system simultaneously.²¹

The *potential* for such risks is also an essential component of the sort of complex uncertainty described in the introductory section. The complexity of the system creates an inherent danger of cascades, because of the tendency of such systems to display characteristics like sensitive dependence on initial conditions and nonadditivity. Small events can produce big outcomes because of the dense set of interdependencies involved and because of the central aspect of “emergence” that characterizes complex systems. This is, once again, the “sandpile effect”: One or a few grains added to a seemingly stable pile will cause a cascading collapse and changed situation. And more than that, we can’t know, in advance, which grain will have the effect. We are waiting for cascades without being able to adequately anticipate them.²²

This reality is closely related to what is sometimes called the “normal accident” problem. As the Yale sociologist Charles Perrow argued in a landmark 1984 book, in complex systems you will inevitably have a cascading of small accidents that eventually produce large ones. Multiple intersecting errors generated in complex systems. Automatic trading systems that create automated response cascades; when a value begins to fall, others sell until a collapse occurs.²³ Dealing with normal accidents is more than just chasing after each discrete risk, because Perrow’s argument—and the lesson of the financial crisis—is that one or another of them will always emerge. Systemic stability demands a sort of correlated resilience above and beyond discrete risk assessment and mitigation.

The great challenge, particularly in national security, is trying to mitigate such potentials for cascading risk. In investments, it's somewhat easier: Firms can be much more fully hedged than was LTCM. There are limits to this strategy, though—often investors have to forgo potential profits in order to create powerful hedges—and in national security the dilemma is even more acute. No country has enough resources to prepare against potentially cascading risks. There are simply too many of them.

But another important lesson of the crisis is that it is easy to leap to the wrong conclusions about the connections between risks. As we will see, in the financial sector and in some other areas, risks are systemically correlated. Because of the essential design of the system and built-in interdependencies, events in one space will necessarily and inevitably cause an impact on other parts of the system. Those effects can be large or small, but the correlation is an inherent feature of the system. In most national security issues, the correlation will depend on peoples' reactions—what we can term perceptual correlation. This can have profound effects too, as in cases of herding and manias. But it is not an essential feature of the system, and it is not inevitable. Some risks that we typically assume might be, such as territorial aggression, are not usually cascading phenomena. They might even be self-correcting. Understanding whether a strategist faces correlated risks—and of what kind—is an essential step to comprehending the overall risk situation.

Goldin and Mariathasan also emphasize two other challenges that emerge from an environment of intense systemic risk.²⁴ One is problems with accountability. When decision-makers lose a sense that their actions will have definable results, in part because of the cascading outcomes of systemic processes, they can abandon responsibility for those actions. "The system did it" is a common refrain—implicit or explicit—from decision-makers who decline to be accountable for their actions. This has been a common refrain from senior financial managers after the crisis: Everyone was doing it; we had the best intentions; we were only trying to make a profit; we could never have imagined the degree of systemic risk.

A second challenge, one that will make the assessment of threats even more difficult, is difficulty in assigning attribution when systemic and cascading threats are pouring through the system. A highly interconnected system offers the opportunity for troublemakers to generate systemic risk through indirect effects and remain anonymous. The nature of systemic risk offers opportunities for destabilizing behavior without consequences, an emerging reality that poses significant threats of systemic instability.

Lessons for national security risk

There are a number of primary lessons for risk in national security from the problems of lack of awareness and accumulating and correlated risk that struck the financial sector.

An obvious one has to do with ignorance of detail. A dozen national security examples could be stacked up next to the financial ones cited above to make the same basic point: When senior leaders lose touch with the details of their organization's initiatives, disaster awaits. A prominent recent example is the invasion of Iraq. In 2002–2003, President George W. Bush and even senior officials like National Security Advisor Condoleezza Rice and Defense Secretary Donald Rumsfeld were as ignorant of the facts on the ground as well as many of the details of war planning—and especially post-war planning—as the financial titans would become of their complex derivatives. Like the financial leaders, President Bush had reasonable goals and was well-intentioned. But he did not educate himself sufficiently to be able to ask the right kinds of questions, and the result was a strategy grounded in lack of awareness. In such situations, it becomes impossible to manage risk effectively.

A related lesson is that entrepreneurs in organizations will frequently exacerbate this dynamic by building dangerous, and highly complex, schemes to achieve some organizational goal. In the case of Iraq, advocates of unseating Saddam Hussein in the Defense Department created intricate proposals for dealing with Iraq, with elements ranging from support for exiles to elaborate political objectives to a proposal to hand power to a provisional government. It was filled with dangerous assumptions and bad analysis, but no one ever took the time to find this out. Policy entrepreneurs can undermine risk management by enticing organizations into quagmires, whether they are aggressive traders or government officials. The solution is oversight, which demands awareness of the details by senior leaders.

Correlated risk can also be a significant problem in the national security realm. Events on one issue or in one area can end up being linked, perceptually or systemically, to others. It can be difficult to recognize when managing risk on a given issue demands consideration to such networked effects.

Yet as important as cascading risks can be under complex systems, another lesson is that the prevalence of correlated risks can also be exaggerated. Every risk, if it comes true, will have second-order effects; that doesn't make it a systemic or correlated risk as we are discussing them. In some cases, certain national security risks are treated as correlated when they might not be. The best example is the concept of precedent: Many believe that if an aggressor "gets away with" an attack on another, for example, that it will have systemic effects by weakening the barrier to aggression. If hostile acts are not counteracted, the argument goes, they will necessarily be correlated with further instability.

But this claim is hotly contested, and in fact the evidence for precedent, and such related concepts as credibility, in international relations is surprisingly mixed. It's not at all clear that one aggressor's actions would inspire others, in part because the specific contexts are so unique. In one case the risks of aggression might be far lower, and the aggressor's urgency far higher.

Ultimately behavior will be mostly a product of an intuitive judgment about the interests at stake, and because no two cases will be the same, the most that a precedent could do would be to address one of possibly many variables.

This perhaps points in the direction of thinking of correlated risks in two ways: Those that are *systemically* correlated, and those that are *perceptually* correlated. In the first case—as with bond markets—changes in one area are systemically linked to others, because of the design of interdependencies built into the network. In the second category, risks can become correlated and spread because of peoples' reactions to them. This sort of correlation will be much more chancy, of course, because different people will react very differently to the same event. Some will be inspired by terrorism, and some will be horrified by it. If one lesson of the recent crisis is for national security leaders to be very attuned to correlated risk, then, a second and equally powerful one is that they should be very careful about identifying what is and what is not a correlated risk.

Ultimately, the biggest lesson has to do with simple awareness. Managing risk demands intense efforts by senior leaders to remain as aware of the design and details of their organization's initiatives as they can. This sort of awareness is the fundamental response to the ignorance that can lead to risk disasters, and the foundation for effective appreciation of risk.

10

Risk, Incentives, and Culture

Risk management procedures exist in part to create objective, abstract answers without bias. This is an important objective, but as we have seen it runs up against a number of powerful realities of human organizations, from personalities to cognitive biases. Behavior is a function of many variables; avoiding risk is one, but there are many others that can be equally powerful. In many ways the essential story of the financial crisis can be boiled down to economic agents acting according to perverse incentives. In particular, judgments do not take place in a vacuum; the human factors that influence the perception of risk take root in a social and institutional context. Peoples' perception of risk is also a function of these broader environmental factors. A critical lesson of the crisis is that different environments can generate different levels of incentive to take risk.

These factors turned out to be especially important in the financial crisis, and on closer inspection seem common to many risk events. The first is the role of incentives in shaping the perceptions and preferences of decision-makers. Most leaders most of the time are operating under some sort of incentive structure that guides their thinking. In the case of the financial crisis, access to easy money, competitive pressures, and aggressive norms of risk-taking in the investment world conspired to generate a dangerous incentive set.

A second contextual factor is institutional culture. This turns out to be one of the dominant foundations for the financial crisis, because it is a variable that can clearly delineate the firms that got into the most trouble from those that partly escaped disaster. Elements of institutional culture are crucial in shaping perceptions, expectations, and behavior, and too many firms had by 2007 allowed their cultures to become engines of disastrously risky behavior.

As it turned out, in many large investment banks, the incentives and corporate culture all pointed in the same direction: taking excessive amounts of risk without adequate analysis or safeguards. There is a good case to be made that these effects were not arbitrary: Operationally oriented organizations

that prize decisive action, have a can-do culture, and operate under the pressure of short-term imperatives will generate this result more often than not. In national security as well as business (and related arenas like space flight), incentive structures and corporate culture carry an ever-present potential to skew risk analysis in dangerous ways. And the result, like so many other factors surveyed here, is to obscure decision-makers' view of the future, and the consequences of their actions.

Incentives in the Financial Crisis

In psychology and economics, incentives are defined as stimuli that motivate behavior. They can be intrinsic (internal) or extrinsic (in the surrounding environment), but either way they create a context in which people make choices. Individual or intrinsic incentives can include desire, fear, and beliefs; extrinsic ones can include social pressure, the potential for gain, threat of punishment, and more. They can bias choice in various directions, sometimes intentionally and sometimes not. And ultimately it was the incentive structure surrounding investment bank decisions that played a dominant role in fostering the crisis.

In a field like investment banking, senior decision-makers will always be motivated by earnings, to beat the market and their competitors. To some extent, these incentives run in the direction of risk-taking, but they are usually counteracted by opposing factors: the fear of disaster and embarrassment, for example, or most fundamentally, regulations that establish various mechanisms for mitigating risk. But in the period before the financial crisis, in specific firms, incentives for risky profit-seeking were allowed to run free. And the result was to magnify other human factors that tend to embrace risk and discount consequences.

The role of institutional culture is grounded in the essential fact that the very concept of risk is socially constructed, based upon prevailing cultural norms. This is the main thesis of Mary Douglas and Aaron Wildavsky in their classic work *Risk and Culture*. "In risk perception," they continue, "humans act less as individuals and more as social beings who have internalized social pressures and delegated their decision-making processes to institutions. They manage as well as they do, without knowing the risks they face, by following social rules on what to ignore: institutions are their problem-simplifying devices." The "latent goals" of an organization, they contend, unspoken norms built into the social fabric of a place, assumed and taken for granted, may be "more influential than those more openly acknowledged."¹

Such socialized norms and conformity strongly affected risk judgments leading to the financial crisis. Established and internalized values recommended—even demanded—the urgent, passionate, sometimes desperate search for short-term profits. In a ruthless competitive environment in which CEOs felt

their positions constantly at risk depending on the last year's performance, keeping up with the results of competitors was an obvious requirement. Sometimes this meant simply adopting their strategies, including complex financial derivatives. These influences suggest yet again that decision-makers' view of the future can be disastrously clouded: The ingrained habits and mindsets of an organization or industry can place all the emphasis on the immediate, the current, and (as we shall see) the perceived imperatives. Worries about consequences are pushed to the side, to become largely irrelevant.

By the mid-2000s, the structured finance vehicles that had been created in the 1990s and 2000s had become ludicrously profitable, offering a range of income streams to the companies that embraced them. Firms simply could not avoid them, and when they tried and others raced ahead and showed astonishing results, the pressure from shareholders and the market to match the innovations became intense. And when some firms appeared to be deriving huge cash flows from complex and risky ventures, it proved impossible for others to exercise restraint. They rushed to the trough, not taking the time to realize it was toxic.

Andrew Ross Sorkin catalogues numerous cases in which firms, even with some awareness of the dangers involved, simply couldn't pass up the opportunity for profits. Not only was it their fiduciary responsibility, the chance for leaders to earn fame and respect—added to the competitive pressure that everyone else would take advantage of it if they didn't—made the temptation irresistible despite the risks. It was almost as if these factors blinded them to a serious risk analysis—it was avoided, brushed off, actively suppressed.

"Intellectually," Sorkin writes, Lehman Brothers head Richard Fuld "understood the risks associated with cheap credit and borrowing money to increase the wallop of your bet. ... But, like everyone else on Wall Street, he couldn't pass up the opportunities. The rewards of placing aggressively optimistic bets on the future were just too great." The massive returns of some of the risky instruments placed even "traditionally conservative institutional investors ... under pressure to chase higher returns" with riskier investments. At one especially telling moment, in a reflection of the underlying biases at work at the time, a senior Lehman official pressed an investment manager to be more aggressive toward risk. "You're holding back," the senior official worried, "and we're missing deals." Merrill Lynch chief Stan O'Neal "saw how firms like Lehman were minting money on investments tied to mortgages, and he wanted some of that action" for his firm.² In the Merrill case, another author has explained, "The failure ... resulted from a single-minded focus on profits, regardless of the mortal risk those profits created for the firm. ... Like every other Wall Street bank, Merrill Lynch under Stan O'Neal chased profits wherever they were to be found from 2004 to 2007."³

This is very much the story of John Mack's tenure at Morgan Stanley, beginning in 2005. As we have seen, Mack's primary concern was that Morgan was

being left in the dust by competitors engaging in large, complex, and risky investments of their own. The classic Morgan image was of a staid, objective, risk-averse financial advisor. Mack saw a need to transform the personality of the company into something else—a daring, risk-accepting firm that played the markets to win. He had this perception in large measure because of the primary incentive faced by all senior leaders in the industry: Keep up with your competition. The result was a plunge into derivatives and, ultimately, subprime mortgages that left Morgan hugely vulnerable when the music stopped in 2007–2008.⁴

The same basic perceptual shift unfolded at the mortgage giant Washington Mutual, which in the early 2000s was gobbling up high-risk subprime loans as fast as it could find them and created a perplexing array of complicated mortgage offers. Profits rolled in, and “With success coming so easily,” Roger Lowenstein has written, “WaMu was deaf to the need to monitor its risk.” The implicit imperative was to close loans, regardless of quality or risk. At one point WaMU’s chief risk officer sent out a remarkable memo suggesting that “risk management functions were being adapted to a ‘cultural change’ and that, in the future, the risk department would play a ‘customer service’ role and avoid imposing a ‘burden’ on loan officers.”⁵

Firms that displayed a willingness to hold back from risky new bets eventually found themselves without any choice, confronting a series of incentives that added up to the sort of imperative-driven thinking examined earlier. Fannie Mae found itself in a more competitive landscape dotted with hard-charging mortgage companies like Washington Mutual. In order to keep up, even to survive, it had to “innovate”—which meant take higher risk. And when worrisome signs emerged about some of the new mortgages they were assuming, Bethany McLean has written, “senior executives disregarded internal warnings, because the lure of big profits was too great.”⁶

In service of these incentives, leaders of financial institutions knowingly took highly dangerous actions that put their firms at risk. Jim Johnson at Fannie Mae pressed hard for lax regulations that would allow his firm to become more leveraged—a route to profits but also greater risk, an odd bargain for a publicly held organization. Meantime later Fannie Mae chief Franklin Raines took the goal a step further by setting ambitious—many thought ridiculous—goals for doubling share prices. An “unhealthy pressure throughout the company” emerged, journalist Bethany McLean has documented, to meet the targets. One senior executive gave an internal speech in which he fulminated, “You must have a raging fire in your belly that burns away all doubts” about the share price goal.⁷

Once such incentives become firmly established, they can spread throughout an industry like a fire through a dry forest. Competitive pressures tamp down any resistance and produce self-destructive behavior. Fannie Mae and other firms lobbied for reductions in the amount of capital they were required to keep

on hand for risk purposes—the most basic, crude and, in many ways, effective way of building resilience against threats. Once one firm can devote more of its resources to profit-making ventures, it gains a competitive advantage—and others are sure to follow. Risk-taking became a slippery slope, with others racing to slide down the abyss behind Fannie.⁸ “With each firm under pressure to perform,” the financial journalist Martin Wolf has concluded, “management found it difficult to avoid entering the race for ever more generous rewards.”⁹

The institutional incentives for profit were matched at the level of individual leaders and traders, whose salaries and bonuses were closely related to firm-level productivity.¹⁰ Fannie Mae CEO James Johnson intensified the industry-wide and firm-level incentives at work by turbocharging the company’s compensation structure, offering bigger rewards to executives who drove higher profits. “Compensation,” Gretchen Morgenson and Joshua Rosner explain, “became tied almost solely to earnings growth.”¹¹

Stories became legendary of massive bonuses earned by people operating in the new arenas. Individual incentives mixed with institutional ones to create huge momentum behind the complex derivatives that helped spark the crisis. Even efforts to structure the character of pay packages to reduce short-term incentives wilted in the face of this avalanche of cash. At AIG FP, traders had to defer half of pay packages for several years to reduce the incentive for get-rich-quick plans. And yet somehow it didn’t prevent the overall incentive of maximizing profits from skewing decision-making in very risky ways. People didn’t speak up against CEO Joe Cassano because “the money was simply too good.”¹²

These incentives, and the drive for an unhealthy degree of profit, did not emerge out of thin air. In the wake of the crisis, many critics pointed to the Federal Reserve, and in particular its Chairman Alan Greenspan, as leading culprits. By keeping interest rates near zero for a long period, the Fed makes money appear essentially free. “The very accessibility of credit,” Roger Lowenstein has argued, “made it appear less menacing.”¹³ Debt and leverage take on a different perceptual meaning when interest rates are roughly nothing. Lowenstein cites a specific moment when Lehman chief Dick Fuld felt able to boost his leverage because, if anything unfortunate happened, credit would be easy to come by. Over and over, when asked about the dangers of complex financial instruments that few people understood, Greenspan calmly assured the world that they were brilliant innovations whose primary effect was to create greater efficiency in the markets. He did little or nothing to use the powers of the Fed to tamp down risk. In effect, the Federal Reserve during this period was an active conspirator in the effort to cloud financial decision-makers’ view of the future, and the consequences of their actions.

At the time, as he would later admit, Greenspan’s rationale was clear enough. He believed in rational decision-making and the efficient market hypothesis. There was absolutely no reason, within this worldview, for him to think that

investors would not be thinking rigorously about risk, and outcomes. Once the crisis struck, he admitted that he had been wrong about one leading factor: Human beings do not make risk choices as his theories suggested.

Effects on group behavior

Once these sorts of incentives become established, they magnify the social and group factors that obstruct objective consideration of risk. Powerful incentives to drive in an aggressive, risk-taking direction will leave very little enthusiasm for taking dissent seriously. Risk-magnifying incentives create an environment in which those who worry about risk are branded as naysayers and even disloyal. This is one reason, perhaps, why short sellers are so reviled by aggressive financial leaders: They are professional dissenters whose job it is to puncture the myths of corporate strategies.¹⁴

Such unwillingness to question a course of action that serves overriding incentives can quickly descend into groupthink, herding, and punishing of dissent. When the environmental incentives are all running in one direction, naysayers often gain little traction. In a different way, similar forms of incentives affected the supposedly independent agencies overseeing markets. Paul Blustein's account of the Asian financial crisis of the late 1990s describes how the IMF and the US Treasury "should have been far more vocal in warning developing countries against the risks of welcoming foreign funds before their banking systems had matured" sufficiently to handle these vast capital flows. But the capital was viewed by everyone as such a critical support for development that "the admonitions were muted."¹⁵ We find the same pattern over and over again—those who warn of risks are ignored, or actively suppressed, by others in thrall of the incentive structure.

Groupthink, loyalty-based inner circles, and the punishing of dissent are common mile-markers on the road to a disaster. Lehman chief Richard Fuld gave rise to such an environment despite the fact that he was reportedly advised to "break out of the inner circle that was buffering him from the firm." Fuld was a fire-breathing, fear-inspiring senior executive who cowed other leaders. "To criticize the firm's direction," Sorkin has explained, "was to be branded a traitor and tossed out the door." Top Merrill Lynch executives similarly came down hard on naysayers: Merrill chief Stan O'Neal was infamous for "fighting back" against people who stood up to him—even having people thrown out of the building over disagreements. "Some employees began referring to O'Neal's top management team as 'the Taliban' and calling O'Neal 'Mullah Omar.'"¹⁶ Such patterns prevented risk from being adequately assessed, because the real value in the organization ceased to be about objective analysis—it became about loyalty and "getting with the program."

Enron is a depressing story of the collapse of institutional responsibility under the weight of malign incentives. The overwhelming incentive structure inside Enron was to value short-term profits—or apparent, reportable profits,

even if largely mythical—over all else, including sensible risk analysis. The business units and their leaders that generated profits and helped the company meet its earnings targets were handsomely rewarded, regardless of how well those investments played out over the long-term. Those who did the troublesome, thankless business of running physical investments with solid collateral and long-term staying power were often overlooked. Enron CEO Jeffrey Skilling would sometimes say, “I don’t like this deal. I *hate* this deal”—as a prelude to approving it, on the rationale that they had to keep trying new things.¹⁷

The effect of such thinking furnished one reason why, even when someone has passionate interest in safety of the firm, it won’t matter. One of the reasons more of the “highly educated, highly intelligent” AIG people didn’t confront Cassano was because he put so much of his personal identity into AIG—and “he was the last person, they assumed, who would blow the place up.”¹⁸ Keeping the money spigot open will always take priority over true risk caution; those around the person taking the risk, if they are sharing in the benefit or goals, will have no incentive to shut things down. When policies to achieve some end—say profit—generate new risk, as Michael Lewis has written, they beg to be “either honestly accounted for or disguised.”¹⁹ All of these incentive structures tended to favor concealment over honesty, if being clear about risks would threaten short-term profits. In such a situation, procedural risk management will be crippled.

The time orientation of senior leaders exacerbates the effects of such incentives. Corporate executives are often worried about the next quarterly report (or even, in the case of many financial firms and Enron, obsessed about daily swings in stock price). In the national security area, senior leaders are constantly focused on looming, immediate crises. Issues that are longer-term, more complex, and more interrelated will tend to be deferred. One result, according to a good deal of empirical research, is that decision-makers are commonly very poor at calculating significant risks if they emerge over the long-term—for example, in the sort of accumulating positions examined in the previous chapter. Leaders are focused on near-term costs more than long-term risks, and decision-makers of all kinds, from CEOs to voters, routinely refuse to mitigate serious long-term risks when there are significant short-term costs.²⁰

Incentives and preference construct in group settings

Nate Silver makes a critical point about the measurement of success and failure for risk management positions like leadership jobs. If you fail in adequately judging risk—but everyone else in your industry fails too, and you all take a big loss, you can escape blame. “Everyone missed it.” But if you *fail to succeed* when others do—that’s when you’ll be edged aside. “They took a risk and made a killing and we got left in the dust.”²¹ All of this creates a huge premium on several dangerous behaviors—embracing extra risk, avoiding a serious engagement with risk analysis, wishful and motivated thinking.²²

The phenomenon could perhaps be called the problem of “accountability in an environment of collective responsibility.” For senior leaders, it’s perfectly fine to sink your firm if you’re making the same mistake everyone else is making. But to veer away from the herd in a way that sacrifices profit—that is almost unheard of, and will get a CEO or whole leadership group fired. “The whole market was pressuring us,” one senior Merrill executive told Roger Lowenstein. “To suffer the organization telling you that you are losing business—it takes a tremendous amount [of courage] to stand up and say, ‘I’m not going to do it.’”²³

The financial journalist John Cassidy argues that the crisis was not a product of stupidity or pure greed, but rather, skewed incentives. The “competitive environment they operated in provided them with no incentive to pull back.” Imagine the results, Cassidy suggests, if the CEO of a leading financial firm had pulled out of subprime investments because of the risk. Competitors would have “rushed in” and grabbed the resulting business. The prudent firm’s short-term earnings would have taken a hit, and the CEO would have been “written off as a fuddy-duddy.” The result is “purposeful but self-defeating behavior.”²⁴

The context of uncertainty—and the limitations of risk instruments in such an environment—midwived this sort of group-influenced risk-taking. Joe Nocera has argued that

The fact that VaR didn’t measure the possibility of an extreme event was a blessing to the executives. It made black swans all the easier to ignore. All the incentives—profits, compensation, glory, even job security—went in the direction of taking on more and more risk, even if you half suspected it would end badly. After all, it would end badly for everyone else too. As the former Citigroup chief executive Charles Prince famously put it, “As long as the music is playing, you’ve got to get up and dance.” Or, as John Maynard Keynes once wrote, a “sound banker” is one who, “when he is ruined, is ruined in a conventional and orthodox way.”²⁵

As the scholars Stephen C. Nelson and Peter J. Katzenstein point out, it isn’t as if most actors in the financial crisis knew they were making horrible bets. The existing conventions, and the incentive structures created by competition, set a *context for understanding* in which they thought their actions made sense. They were unconsciously self-deluding. The two scholars quote former Fed governor Lawrence Lindsey as saying, “we had convinced ourselves that we were in a less risky world.”²⁶

Corporate culture is decisive in shaping risk behavior

The Long-Term Capital Management debacle strongly emphasizes the importance of culture in determining how an organization will respond to risk.

LTCM was founded by a number of supremely self-assured number-crunchers who believed passionately—in some cases ideologically—in the effectiveness of their models. It was a combustible mix of academics (Merton and Scholes of Nobel fame) and almost ridiculously optimistic traders. (Roger Lowenstein's brilliant case history refers to one partner's "supreme conviction in his own rightness.") The partners had the "arrogance of people who had been to Harvard and MIT—of people who really believed that they were more intelligent than others."²⁷ LTCM was also secretive to the point of obsession. Senior partners refused to share detailed information about their trades with their investors—or even more junior people at the firm. This lack of transparency became a significant factor in the disaster.²⁸

Culture and risk

Organizational culture, writes the scholar Edgar Schein, "is a property of a group. Wherever a group has enough common experience, a culture begins to form." This process of social construction imbeds the assumptions of the culture deeply into the minds of the participants. "Culture is so stable and difficult to change because it represents the accumulated learning in a group." Culture matters because "it is a powerful, latent, and often unconscious set of forces that determine both our individual and collective behavior, ways of perceiving, thought patterns, and values." An organization's culture can include everything from its values to its workplace design to its promotion systems to things like dress code and level of formality. Jargon, uniforms (if any), ways in which decisions are made, working hours, rites and rituals of the place—all of these play into the total sense of an organization's culture.

Culture is both important and stubborn because it addresses human needs. It is, Edgar Schein writes, "a mechanism for making the world meaningful and predictable," functioning like a theory, or an ideology, or a belief system. Culture allows us to internalize values about work and leadership without having to figure out all the messy details for ourselves.

A critical thing to keep in mind about culture is that, as Schein explains, it is "the residue of success." An organization develops a culture and either fails or succeeds. If it succeeds, the organization comes to view its culture as part of the reason. If the Navy has a culture of independence, emphasis on time at sea and a conviction that its capabilities are essential to the nation, it is because these things have made it, or been seen as, successful over the years. Microsoft's culture of tough and uncompromising focus on numbers is, senior executives believe, intimately connected with its success.

The scholar Donald C. Langevoort has described the role of organizational culture in shaping views of risk. Overconfidence in general, he argues, can drive excessively risky behavior. Institutions can then reflect a reverence for highly confident people who take risks that pay off. The result can be a snowball of risk-taking and those who have had some initial success are encouraged

to take more risk. The same effect can emerge in organizational culture, Langevoort writes:

An overly optimistic or confident internal corporate culture should have precisely the same effect. It is easy to see how the common perceptions that mildly excessive optimism and confidence generate are organizationally adaptive. They reduce worry and anxiety, which facilitates focused work. They prompt thoughts of a bountiful future, a larger pie to share, which facilitates intrafirm trust and cooperation, while a depressed culture does just the opposite: it frames the situation as a last-period risk that generates selfishness and defection. Thus, integrative solutions to the endless negotiations that make up the day-to-day work inside the firm are easier to find. And that, of course, is also the problem. If those deeply enmeshed in this culture are responsible, say, for accurately disclosing the level of risk that the firm faces, they will often get it wrong—albeit often in good faith—because they are too optimistic.²⁹

Wishful thinking and overconfidence can become baked into the corporate culture of many competitive organizations, whose cultures then come to reflect and reinforce these habits—and to forcibly dismiss doubters who raise worries about risk.

Organizational culture and the financial crisis

Aspects of the culture of specific institutions played a critical role in creating the financial crisis. Simply put, many leading financial institutions had become dissent-quashing engines of risk-taking by 2007. Controlling risk was not a priority. The reigning institutional cultures at many firms valued, almost worshipped, bold risk-taking in service of huge and rising revenues. The ultimate effect was to create incentive structures that made it almost impossible to control risk.

In the case of Merrill Lynch, once some people started cluing in to the risks built up in the firm's CDO positions, a few canaries in the coal mine began warning about them. They were typically dismissed as naysayers.³⁰ A sure tell-tale sign of a culture in crisis, one where dissent will not be rewarded, is for doubters to repeatedly be told, "You don't get it."³¹

William Cohan has chronicled various aspects of the Bear Sterns culture that contributed to its swallowing of billions worth of high-risk CDOs without understanding what it was doing. It was a culture of brutal competition and regular, vicious criticisms of senior managers. The chief value—after making money—was loyalty. One finance person compared Bear and Goldman Sachs in their basic personalities: The difference could be seen in the language they used. Whereas a hugely successful trade at Goldman would be described in technical, almost boilerplate terms—"That was a very attractive

and commercial price you purchased those securities at and I think we'll have a very interesting economic opportunity in the near future"—at Bear the same trade would prompt a boast that "I just ripped that fucker's head off." Bear Sterns, this person told Cohan, "was never as aggressive as its reputation. But its language, its culture, and its swagger put it more at risk than its actual actions."³²

A particular aspect of the culture of many of these firms was the emergence of a centralized, self-righteous "in crowd" of partisans who saw serious questions or dissent as disloyalty. One example is LTCM. It was "dominated by the two senior traders," Roger Lowenstein has argued, "a partnership only in name." Two leading players dominated the discussions, and because they were aggressive, their role defeated formal risk assessments.³³ In the case of Merrill Lynch, CEO Stan O'Neal created an inner circle of a few key advisors who dominated his perceptions and appeared in some cases to shut out those with contrary views. The disaster at Merrill was ironic given the culture with which it had been founded—a family-oriented "affiliative" environment that revered individual investors and eschewed Wall Street bling. Stan O'Neal felt this culture was not tough enough for the new competitive landscape, but what he "did not realize was the danger of eviscerating one corporate culture without establishing another one in its place."³⁴

These hallmarks of a risk-blind institutional environment could be found across the industry. Freddie Mac had "a culture of arrogance" stemming from its perceived ability to dominate Washington political debates. "That overconfidence led both companies eventually to move into derivatives and to employ aggressive accounting measures." Freddie Mac CEO Jim Johnson refashioned the personality of the formerly sleepy purveyor of loans into an aggressive, profit-seeking, industry-leading shark—and one more than willing to spend huge amounts of lobbying money to make sure it got the treatment it desired. The cultural ethic was to make loans, whatever the risk. Meantime at AIG "executives stubbornly clung to the belief that their firm was invulnerable," in part because they thought their structure and low debt ratios insulated them from risk.³⁵

The story of Enron can in many ways be boiled down to an institutional culture gone haywire, poisoning any sort of legitimate or accurate risk assessment. Put simply, Enron was a deal-making engine. There was little emphasis on how the new investments or business would actually be run afterwards. "The mentality on most of this stuff was they did deals and moved on," one Enron executive told McLean and Elkind. In service of this goal, the firm built up a secretive culture of mutually suspicious cliques, a dog-eat-dog competitive environment, people intensely loyal to their unit heads, and a general sense of arrogance that they were working for the best company in the world. Enron traders had an especially toxic subculture, "insular" and "self-righteous," which "allowed the traders to justify making money in ways

companies should never countenance.” They thought they were “creating a new world”—and such visionaries needn’t adhere to the lame rules that applied to the rest of an industry. Most of all, it was a culture of risk-taking and gambling: “Betting was a way of life” at Enron.³⁶

A notable counter-example in culture was Goldman Sachs. It was hardly a perfect culture. Top Goldman officials were every bit as aggressive and profit-seeking as their colleagues in the industry, and the firm has generated its share of mistakes and follies. Nonetheless, the institutional regard for risk management and rigorous decision-making appears to have been just strong enough to avoid the consequence-free thinking of its competitors. It retained, alongside aggressive profit-seeking, a mostly sober and serious attention to reality. “Of the great investment banks,” Greg Farrell has written, “only Goldman Sachs heeded the warnings of its own internal risk managers, who began sounding the alarm in late 2006.”³⁷ Goldman actually started shorting the mortgage industry thereafter, and got off best in part because its culture was the most conservative and risk-averse. But it was more than willing to allow others to wreck themselves on the same shoals: Goldman continued selling the instruments long after it had determined that they embodied unacceptable risk for its own investments.³⁸

Organizational culture and the “normalization of deviance”

Another way in which culture and related variables can exercise an effect on behavior is through the gradual accumulation of small decisions that add up to highly risky actions that an organization would not have endorsed at the beginning of the process. The sort of accumulating risk described in the previous chapter can change the perceptual lenses of an organization over time: Having weathered a hundred decisions on risk, institutional leaders can be lulled into a sense of false confidence about the next three or five risk-oriented choices. We tempted fate a hundred times and got away with it, they will unconsciously believe—why place special emphasis on the hundred and first?

The main thesis of Diane Vaughan’s study of the Challenger space shuttle launch decision is that social norms and values were dominant causes of the disaster. She dismisses the conventional wisdom view of the accident—that NASA leadership was affected by pressures to launch to save the space program, and consciously overrode safety warnings to do so. Instead, she suggests, a long series of specific decisions summed up to one large risk choice—but that larger choice was never really made on its own terms. It is a form of the boiling frog metaphor, except that it is the frog itself that is gradually turning up the temperature of risk, one notch at a time, until it has boiled itself.

The space shuttle was a risky proposition to begin with, as any space flight will be. But it was an especially perilous venture because of the design parameters required of a vehicle designed to return to Earth and be reused.

From the beginning, too, designers were forced into a series of compromises on risk: Weight constraints, for example, caused them to abandon the idea of any sort of escape pod. During the flight history of the shuttle, dozens of risks had cropped up and been dealt with, including problems of temperature and potential dangers of the infamous O-ring seals on the engines. By the time of the Challenger launch, NASA had made a long series of choices that accustomed engineers and senior flight managers alike to the idea that they were taking a substantial new risk. Warnings about the O-rings were met with indifference because the program was grounded in the simple reality that risk was ever-present. That very same risk of O-ring vulnerability had been discussed, and seemingly conquered, a number of times. The result was to create an unconscious disregard for new warnings: The process of incremental risk acceptance dissipated the perceived danger of any future choice, even if the risks were, in fact, much higher.

The decision to launch the Challenger in borderline weather conditions is often presented as a one-off choice with tragic consequences, one that violated many standard procedures and ignored obvious evidence—and clear dissenting opinions—about the vulnerability of the O-rings. But Vaughan demonstrates that the whole shuttle program was based on the *acceptance* of risk, rather than its unambiguous avoidance. In order to justify the series of incremental choices, flight managers had to create a shared view of risk that caused them, mostly unconsciously, to disregard future warnings. The key lesson, for Vaughan, is to remind us about “how environmental and organizational contingencies create pre-rational forces that shape worldview, normalizing signals of potential danger, resulting in mistakes with harmful human consequences.” The Challenger disaster, she argues:

is a story of how people who worked together developed patterns that blinded them to the consequences of their actions. It is not only about the development of norms but about the incremental expansion of normative boundaries: how small changes—new behaviors that were slight deviations from the normal course of events—gradually became the norm, providing a basis for accepting additional deviance.³⁹

This process can become an element of institutional culture because the “weighing of costs and benefits” occurs under the shadow of “institutional and organizational forces.” Norms related to them “create unreflective, routine, taken-for-granted scripts that become part of individual worldview.” Despite NASA’s intense safety culture and elaborate risk assessment processes, “social, organizational, and institutional influences on risk assessment” skewed the decision. The case “illuminates culture as supremely important in shaping risk assessments in the workplace.”⁴⁰ The result, of course, is to erect more roadblocks in the way of the rigorous assessment of consequences.

We find exactly this sort of process leading up to various financial and business risk disasters. An archetypical case is Enron, which, in its risk-taking culture and ultimate embrace of illegality, represented a perfect case study of the normalization of deviance. The extreme risk-taking and financial mischief that would ultimately bring down the company started relatively small, with garden-variety excessive risk-taking in investments and the tamest rule-bending of the sort that goes on in most organizations. But soon, in part out of a desperate need to fashion revenue-generating schemes to feed the beast of Enron's growth, executives began taking more and more risk, and hiding more and more truth. Enron's collapse didn't stem from a single decision to do something blatantly illegal. Instead "it grew out of a steady accumulation of habits and values and actions that began years before and finally spiraled out of control."⁴¹

Fannie Mae provides another example of a long stream of choices that gradually changed the culture of a once conservative organization. CEO James Johnson brought an aggressive, risk-taking approach to the firm and slowly, through dozens of specific decisions, carried it to a place where consuming thousands of shockingly low-quality loans—behavior that would have been out of bounds years earlier—seemed perfectly acceptable. Fannie Mae "morphed" over time, a former executive of the organization told Gretchen Morgenson and Joshua Rosner, into "an institution that [was] corrupt or corrupting at more than the margins."⁴²

Vaughan's concept is related to the political scientist Charles Lindblom's argument about incrementalism. Linblom contends that, under uncertainty, decision-makers often resort to incremental actions: Because they can't be sure about outcomes, and guided by an unconscious sense of risk avoidance, they tend to prefer a series of modest actions over big leaps. The emphasis is always on the immediate future, the constraints of the current context, and what minor action makes sense to take a halting next step. Vaughan describes a situation in which such a line of action can "normalize deviance," but the effect could be broader: Dozens of large-scale outcomes, whether risk disasters or not, can be characterized as the sum of a hundred incremental choices that added up to a larger strategic reality no one had necessarily intended at any one step along the way. The dominance of incrementalism would also tend to dim decision-makers' perception of consequences, because all they are really thinking about is the next, minor step. The true consequences of the whole series are invisible to them.

Implications for national security

One fundamental lesson of the crisis, then, for national security as well as other realms, is that institutional culture is critical to effective risk management and uncertainty tolerance, and that intrinsic and extrinsic incentives

can be strong enough to override even the most elaborate risk management process. Faith cannot be put in processes, because they will often give way in the face of intense pressures from incentives and skewed culture. Effective risk management demands, first and foremost, attention to human factors: shaping the incentive structures in organizations to favor objective analysis and constructive feedback; creating institutional cultures where a prudent balancing of risk and reward is the highest value.

In particular, under the influence of skewed incentives and broken cultures, the quashing of dissent, imposed loyalty, groupthink, and other barriers to rigorous risk analysis become disturbingly commonplace. The potential for such outcomes makes it especially critical that organizations nurture cultures that are tolerant of candor, warning, dissent, and qualitative arguments. They must entertain broad statements of risk without dismissing anything that doesn't have a recommendation attached (the "fine but what do I do about it" syndrome). Despite the pressures and urgencies of the incentive structures surrounding competitive organizations, it is absolutely critical that people feel able to raise doubts and risks. The lack of candor, effective dissent, and other forms of open dialogue was a major contributor to the crisis. Nassim Nicholas Taleb complains that business executives enthuse about his discussion of black swans—and then go right back to their offices to play around with their Gaussian distributions.⁴³ Organizations need to create an environment in which people get more credit for being a contrarian and an accurate forecaster.

A structural choice within organizations that can have a similarly chilling effect on rigorous discussions of risk is to exile risk functions into distinct areas of the company, separate and discrete from the operational heart of organizations, the profit-making or mission-completing elements. Those offices tend to view risk specialists as naysayers and worriers who don't contribute to the bottom line. Even success in risk management isn't seen as contributing to the bottom line, only averting disasters that weren't assumed in the profit projections.⁴⁴

One of the most important overall implications of these issues, then, is that it becomes critical to integrate and involve senior leadership in the process. Winning their respect, and willingness to accept the risk advice, is essential, and extremely difficult. In part this is true because the whole judgment process is so emergent and intuitive, and senior leaders believe *their* feeling-of-knowing to be superior to anyone else's.

Part III
Toward Improved Risk Practices

11

The Role of Risk in Strategy

The previous chapters have examined the ways in which risk management can go wrong, often under the influence of human factors that overwhelm procedural approaches to risk. What remains is to discuss potential solutions—ways of approaching risk that offer better opportunities to deal with the potential dangers.

As argued in Chapter 1, the problems with risk in the financial crisis—and in parallel national security cases—represent an issue of a lack of warning and institutional awareness as much as a failure of risk management. It wasn't that the risks weren't identified, but they were not adequately understood or communicated to and among senior leaders, and warnings were downgraded or ignored. Improving risk management therefore must revolve around a program to enhance the institutional culture supportive of warning and analysis. More risk procedures, or more elaborate ones, will not make a significant difference.

Part III of this book makes a three-part argument about the shape of effective risk management. First, in this chapter, it contends that the specific role of risk in strategy must be carefully defined and understood by all the participants. As suggested earlier in the book, I will make an argument for an outcome-oriented concept of risk that leaves considerations of threat and sufficiency to other steps in the strategy process. Chapter 12 will then illustrate the concept of outcome-oriented risk assessments by offering a sample assessment of the current US national security strategy. Chapter 13 lays out the second part of the argument—avenues to improve institutional culture and habits to improve the consideration of risk and warning. And third and finally, Chapter 14 will make an argument for a broader and more encompassing approach of “managing uncertainty,” in which outcome-oriented risk management is one important component.

As noted in Chapter 1, my focus is the way risk analysis can support what I am calling “complex strategic judgments.” These are strategic choices under

uncertainty that offer none of the deterministic predictability of problems that can be addressed with quantitative measures. As defined earlier, they are choices in which the outcomes of current trends or prospective behavior cannot be forecast; based on incomplete information; and which involve inherently subjective issues, complex, nonlinear dynamics, and contested values. As a result they offer no optimizable solutions, instead requiring unconscious interpretation rather than objective decision analysis. These are only one subset of the kind of choices faced by enterprises or senior leaders, of course. But they also represent the most consequential national security decisions. In such cases, I am arguing, the evaluation of the possible consequences of strategic choices often gets limited attention and should be the focus of risk analysis.

Effective risk management in such complex strategic judgments is about managing human factors, personalities, culture, and incentives. It is not about assessing data or establishing checklists or frameworks. Organizations that are serious about managing risk must first decide what it is and what role it has in strategy, and then work to ensure that their culture reflects the right mindset and habits. This is much harder work than simply creating a chief risk officer, giving him or her a staff and a budget, and beginning to “measure” risk.

The objective: driving the necessary conversations

Organizations need a very specific concept of the role risk and risk management functions will play in strategy. All too often, that is conceived as delivering accurate *data* about the current risk profile of a firm or institution. Part of the goal of any risk process will be to inform leaders about as many details of risk as possible. But this data is only a means to an end—driving the right kinds of conversations, in order to inform complex strategic judgment.

A leading characteristic of risk management in the financial industry before 2007 was the widespread conviction that complex and sophisticated models had made sense of markets well enough to allow a very fine-tuned estimate of the risks undertaken with various investment strategies. This was, in a sense, the foundation of the whole emphasis on procedural risk management—the idea that formal, quantified models could make a critical difference. The implied message was that there were reliable models that could offer a clear answer about a firm or organization’s risk level. As long as the organization’s leadership paid attention to these static reports, they could think of themselves as “managing risk.”

Such tools allowed firms to take on substantial risk, in the form of leverage as well as the character of the investment instruments, with sometimes blithe confidence. The edifice of risk management in national security has also developed along these lines, in ways that are methodologically sophisticated and yet also potentially misleading. “Today, DoD, the services, and combatant commands are involved in a bewildering array of risk-assessment processes,”

several RAND risk experts concluded, “involving thousands of man-hours. Many of these activities may be of great value, but at least some are such black boxes that they obscure rather than illuminate risk and uncertainty for senior decisionmakers.”¹ Challenges with such models arose in both the terminology and the reliance on probabilistic analysis to tame uncertainty.

The issue isn’t that data are unimportant or meaningless. It is that a key lesson of the crisis is that any risk process must be a spur to further analysis and dialogue, rather than become perceived as a self-contained and objective assessment of risk. Risk at the strategic level—the focus of this study—is generally qualitative and nonlinear. Senior leaders will never obtain a comprehensive or reliable portrait of it from data or models alone. All too often in the financial sector, however, highly structured risk processes and analytical frameworks concealed as much as they elucidated before 2007. Areas of potential danger were underappreciated or ignored altogether.

A risk process cannot be perceived as a series of quantitative models that dispense unambiguous and objective assessments of a firm’s risk. It should instead be used to illustrate more qualitative, outcome-oriented issues—to highlight potential dangers and force leaders to address things that could go wrong with their strategies. Indeed, too much focus on data, as we have seen, can have a counterproductive result: It can shut down dialogue rather than encourage it, and generate false confidence in the degree of certainty a firm or organization has about its level of risk.

The goal of risk management in this context is to equip senior leaders with the best available information and insights about the possible outcomes of their choices. In complex strategic judgments, the ultimate decision will rely on a leader’s sense or instinct—the meaning they impose onto the choice. There will seldom if ever be a single optimal solution, but risk processes can enhance the quality of judgments in a number of ways: Making sure senior leaders have all of the available facts and information about what could go wrong; ensuring that they do not miss a major issue or category of risk; and forcing in-depth dialogues about the tradeoffs between risk and opportunity and how to mitigate likely dangers of various options. The more objective, rigorous, and widely appreciated a risk process is, the more it will be able to pursue these goals in ways that help correct for the human factors catalogued in earlier chapters.

Ultimately, then, the most important standard for a successful risk process is that it has encouraged decision-makers to interrogate the potential consequences of their actions seriously, in depth, without being able to brush aside unwelcome consequences. And the most important conversation has to do with the outcomes of proposed actions. This chapter argues that the purpose of risk management in national security strategy ought to be narrowly targeted on assessing outcomes of strategic choices. It lays out a framework for doing so—a first draft of possible criteria and categories to employ in order to identify outcome-oriented risk.

Risk as outcome analysis

Whatever goes by the name of risk management should aim first and foremost to generate a disciplined, rigorous, and objective look at the potential dangers embedded in current policies or proposed options. The focus should be on what could go wrong with policy and strategy—things that could either undermine the effectiveness of the program or else generate larger malign consequences.

Chapters 4 through 10 surveyed a number of the most important lessons of the financial crisis for risk management, and in particular the threats to effective management of risks that became apparent in that experience. Taken together, these factors point to one leading area for risk failure—insufficient consideration, and mitigation, of potential negative consequences of a proposed strategy. My research suggests that the most profound risk disasters in finance and national security come from insufficient attention to and awareness of the potential risky consequences of intended or favored strategies. Simply put, organizations and leaders confront sometimes intense cognitive, cultural, and institutional barriers to adequate consideration of outcomes. The incentive structure for senior leaders, the default decision process under uncertainty, and the leading implications of a number of cognitive biases all tend in the direction of muting outcome-oriented risk analysis. When catastrophes strike, whether a financial crisis or the collapse of a company like Enron or a foreign policy debacle such as the Bay of Pigs or chaos in post-invasion Iraq, the culprit is often the same: Decision-makers refused to take seriously the consequences of the course they came to advocate.

The lessons of the financial crisis, as outlined in the preceding chapters, support this conclusion for a number of reasons. To begin with, the presence of uncertainty points to the need for disciplined thinking about potential outcomes and scenarios rather than quantified risk metrics or efforts to conduct valid foresight. It also magnifies the danger that decision-makers will rely on largely intuitive guesswork because of their appreciation for the implications of uncertainty. When leaders appreciate a situation to be complex, ambiguous, and nonlinear, they are more likely to default to a gut-feeling approach to judgment. This opens the door to all manner of biases and framing errors. The subjective, meaning-making character of the confrontation with risk points to the ever-present danger that decision-makers can make of risk whatever they want—a tendency especially on display in the repeated ability to downplay consequences.

The character of institutional culture as well as the personalization of risk exacerbates these problems. This seems especially and dangerously true with regard to the consideration of consequences, which can be deeply submerged under layers of relationships and constraints on open debate of issues. Decision-makers' ignorance about the details of strategy becomes a particular

problem when it hampers their ability to understand the potential consequences of the actions they are endorsing.

The result is that leaders are all too often blind to outcomes even as they make decisions with a supposedly future orientation. Driven by imperatives, values, personal stakes or other factors, decision-makers in finance as well as national security commonly engage in nonconsequentialist thinking. Once in this mode, risk analysis becomes almost impossible. Decision-makers are simply not in a state of mind to care much about what might go wrong as a result of a given strategy. This lack of outcome orientation on the part of many senior decision-makers demands a process that forces attention on potential consequences.

Existing approaches to risk management in national security do not, by and large, deal decisively with these dangers. As we have seen, they tend to employ one of two alternative approaches: equating risk with threat, and defining risk as the gap between requirements and capabilities. Both of these methods have severe handicaps, and neither makes a decisive contribution to the making of strategy. Risk assessment should focus on an analysis of potential outcomes of offered or chosen strategies—things that could go wrong in their implementation, whether internal to the organization choosing the strategy or as a result of reactions it provokes. The result would be a risk-informed decision about alternatives, within the context of a general process of managing uncertainty.²

Interestingly, the statute that mandated a risk assessment as part of the Quadrennial Defense Reviews appeared to have this as its intention. It mandated an assessment of the risks *associated with the proposed strategy*, rather than in the environment. The statute called for a risk assessment that would “define the nature and magnitude of the political, strategic, and military risks associated with executing the missions called for under the national defense strategy.”³ This could refer in part to ends-means analyses, which would fall under a broad interpretation of the language. But the focus seems clearly to be on the outcomes of the strategy: If adopted, what risks would the national defense strategy entail?

In practice, this has not been the emphasis of the risk sections in successive QDRs and other risk management processes. They have focused much more on sufficiency issues. These are arguably a form of outcome risk—a consequence of the strategy is that we will not have sufficient resources to do the job. As part of a broader outcome assessment, this could be perfectly appropriate. But it is generally not conceived that way. The risk analysis aims more at the consequences of programming and resourcing decisions rather than at the big choices inherent in the strategy. And when risk is defined *as* sufficiency, the focus is on managing the ends-means relationship, not on considering outcomes from a wide range of scenarios.

Again, this recommendation does not presume that existing national security processes never capture potential consequences. In some cases, elaborate

options memos spell out pros and cons of specific courses of action, and by implication deal with risk. On balance, however, a review of available information about risk processes and approaches suggests that such assessments are not consistent. In some cases, as case histories of key national security choices make clear, debates or options analyses of courses of action will bypass risks. Seldom are they comparative—a default option will be presented without comparative risks relative to others, or at least to valid and plausible options. The assessments will not be based on clear criteria or established frameworks, nor will they be in depth—risks will often be presented as a single line or simple statement.

In order to be disciplined, such an effort must focus on what could go wrong that is unique to the proposed course of action or strategic concept. Only then will the analysis be a unique outcome analysis rather than a general discussion of dangers. Organizations must consider not only the risks of a generic future world, but the dangers of the *specific strategy* on which they embark.

A potential source of confusion arises if an organization is trying to measure outcome-related risks relative to a strategy that is already in place. For existing strategies or policies, trends or events or factors in the environment that would generally come under the heading of contextual risk (or, in the preferred conception, simply threats) become risks to the ongoing strategy. This confusion will be especially evident when organizations have persistent, long-term elements of their strategies (such as investments in a volatile marketplace) that can be upset by contextual factors. The United States, for example, has a default strategy of deploying forces in South Korea to help deter attack. This is not a choice that is made again every year (or even every decade)—but it is a choice, backed up by a detailed strategy, nonetheless. Is instability in the North Korean regime then a threat, or an outcome-oriented strategic risk?

At a certain point, of course, this discussion becomes semantic. The important thing is to identify strategies—whether current or prospective—and assess specific risks in relation to them. If aimed at an existing strategy, the risk process should focus carefully on the specific outcomes or consequences of that choice, rather than anything that could go wrong in the environment relative to US interests. Whether dealing with current or prospective strategies, the purpose of risk should be to interrogate their consequences.

Risk exists relative to a strategic concept

Very often in national security, officials and analysts refer to risk as if it were an objective condition, measurable against a generic set of capabilities or qualities. But an important lesson of the financial crisis is that risks only emerge in relation to specific strategic concepts. One lesson for the current US posture is that a variety of factors often cited as generating risk—growing Russian

belligerence, global extremism, declining US defense resources, the overextension of the US global posture—are not in and of themselves risks. The relative decline of US hegemony is a *fact* of strategic life, not a risk. The relevant risk comes from the potential consequences of the strategy the United States chooses to deal with any of these situations. This seems obvious enough, even self-evident—but many discussions of risk in the national security context rush past it in order to begin talking of threat-based risks in the context.

Competitive advantage comes from a very specific idea of what role a country is trying to play in the world—from the coherence of its strategic posture. If that concept is too ambiguous or undefined, effort will be expended on a random set of initiatives that do not maximize a country's resources. The United States needs a clear understanding of the specific objectives it hopes to achieve, and the theories of success that allow it to apply constrained means to achieve them. So far there is little hint of such a defined role in US national security strategy documents, which remain vague and all-encompassing.

The risks that financial firms faced as a result of the financial crisis were the direct result of strategies they adopted. It is true, as we saw in Chapter 10 on incentives, that some aspects of the situation influenced the behavior of those firms: Easy credit, for example, and lax regulation both helped to lay the groundwork for disaster—and as the next section will suggest, an important task today is using norms and standards to incentivize stabilizing behavior. But corporations were still free to adopt a range of strategies to deal with that situation. Some chose concepts that were far riskier, less well understood, and less proven than others—mostly, as we saw, because they had become engines of profit such as had never been seen before on Wall Street.

This lesson suggests that the nation has spent far too much time debating and worrying about the “risk” to its existing posture, and far too little productive time discussing feasible alternative courses that allow it to deal with that changed circumstance in the best possible way. The United States can influence events to a degree, but the large trends of history are beyond its control. The national debate should define these aspects of the context but then shift quickly to debating alternative strategic postures, and use that dialogue to lead into an intelligent risk discussion.

The experience of the Eisenhower administration demonstrates how difficult it is to generate a true comparison of alternatives at the strategic level. The Solarium Project demanded that the system lay out and assess three alternative approaches to grand strategy during the Cold War. And yet it was perhaps the only such formal exercise that has taken place since World War II—or perhaps ever, in US foreign policy history. US national security policy is typically sluggish in making changes as a result of potential alternative approaches. A good example is the emphasis on non-military instruments of statecraft. For well over a decade now, most national security leaders and experts have agreed that the strategic demand set calls for more fully developed non-military

tools, more evenly balanced with the military instrument. But because there is no formal strategic concept that embraces them in a doctrinal and policy way, the investments continue to lag the recognition.

One of the challenges in assessing risk, in fact, is that US national security policy lately has been weak on general concepts. National security strategies have been general statements of intent and aspirational wish lists accompanied by laundry lists of possible means—without a basic core notion of how the strategies expect the means to achieve the goals. In the absence of clearly articulated concepts at the strategic level, it is impossible to conduct outcome-oriented risk analysis: There are simply no strategies whose outcomes can be analyzed.

The importance of alternative views

The concept of alternative interpretations is central to any outcome-oriented analysis of risk. For one thing, in assessing what could go wrong with regard to a specific strategy, decision-makers need a rigorous sense of different conceptions of the *current* situation. A given context—the behavior of a state in international politics, a particular set of economic and market conditions—could have varying explanations, and disputes among alternative ways of understanding the current context are central to strategy and to risk. The market may have reached certain heights, for example, because of objective factors, government policy, an irrational bubble, or other reasons. Put simply, taking seriously various reasons why the present situation—the grounding context for risk assessment—has arisen is central to examining possible emergent risks.

It is the mindset of alternative thinking that is as important as any specific ideas it generates. As we have seen, risk, and strategy more broadly, most often goes wrong when decision-makers get locked into a narrow, singular view of the world. A process of alternative interpretations helps to fracture this mindset by forcing them to consider different possible explanations for the current context and alternative scenarios of what could happen once a strategy is adopted. A major risk of the sort of intuitive, naturalistic, feeling-of-knowing approach to decision under uncertainty is that it encourages decision-makers to grasp onto the first persuasive narrative, and then defend it against contrary evidence. Creating a formal and extensive alternative interpretations process as part of both strategy and risk assessment offers a partial answer to this tendency.

In the process, the alternative analyses should focus heavily on the causalities involved. As we have seen, ambiguous or nonlinear causalities are common under uncertainty. In complex strategic judgments, mistaking the causalities is common and can be dangerous. Discussing detailed alternative interpretations can help leaders interrogate the causes of current situations

and potential causalities operating in the future. The idea is not to develop an accurate, objective understanding of precise causalities, which is impossible under true uncertainty, but to encourage a rigorous and effective dialogue that improves the understanding of senior leaders.

This focus on alternative interpretations should then extend to the future, and embrace alternative scenario analysis of how proposed strategies could unfold, with a particular emphasis on dangers. This is classic scenario planning but focused on what could go wrong with strategies. It can build on alternative analyses of the current situation, pushing those forward and playing with the various causalities identified during that part of the process. The idea, again, is to encourage rigorous, open-minded thinking about different possible futures, and to pull decision-makers out of a single, default vision of what might happen.

Clarifying theories of success

As part of this process, risk management should aim at a critical related objective: forcing decision-makers to be explicit and rigorous about their theory of success. This is arguably the most common gap in strategy—a rigorous assessment of the concept by which a strategic actor's various means will add up to the effects they intend to achieve. Such an assessment of strategic concepts should be *the* central part of the process: The characteristic that distinguishes strategy is the employment of a carefully selected group of measures in order to have a desired effect. Understanding the mechanism or theory of success at work is the *sine qua non* of effective strategy.

And yet this crucial phase of strategy-making often gets brushed quickly aside in favor of laundry lists of means. Concepts offered as “strategies” are typically nothing more than lists of actions that the authors hope will result in certain outcomes. But the *mechanism* involved—the specific process *by which the means work together to achieve the end*—is often left to the imagination. On closer examination, the presumed mechanism of most strategies has no firm basis—perhaps some very rough commonsensical appeal, but no foundation in rigorous analysis.

In Afghanistan, for example, a core principle of the US strategy was that the accumulation of dozens of tactical wins, a thousand aid projects, a hundred capacity-building efforts, and overall pressure on the Taliban would naturally generate a political settlement. Exactly *how* they were meant to do so, however, was never clear, apart from the generic belief that a political actor under pressure is more likely to deal. When a laundry list of means substitute for a true strategic concept, the result will always be that many critical questions get begged, and assumptions passed over. In this case, the questions of how many wins, what degree of development, or what level of political legitimacy would be sufficient to achieve the objective were never defined. The Taliban's

breaking point was never known with any degree of fidelity; the very fact that they would be amenable to a settlement short of their goals was assumed but never proven. The result was a list of actions rather than a strategy.

The fundamental answer to this problem is not, strictly speaking, risk assessment. Laying out a strategic concept and proving why its mechanism will have the desired effect—or at least making a strong case for it (which is often the best that can be done in complex, uncertain, nonlinear contexts)—is the work of the strategy process itself. But risk assessments focused on outcomes offer a critical testing ground for the assumed outcomes of strategic concepts, asking decision-makers to discuss and debate whether their proposed strategic concepts will generate the consequences they intend.

Challenges with outcome analysis

Two complications immediately present themselves in focusing risk management on outcomes. One is that it can be difficult to separate environmental threats from risk to specific strategies, because risk emerges from the *intersection* of proposed strategies with aspects of the environment.⁴ Factors identified as threats can be the engines of risk when they affect a proposed strategy. For example, Russian belligerence in the Putin era is a threat—but in relation to a proposed US strategy to arm Ukraine, it also becomes, or generates, a risk to the strategy. What this suggests is that there will be some overlap, and should be a close analytical connection, between the environmental analysis and the risk assessment. It still does not, however, indicate that the two are equivalent: Identifying and assessing the *threat* of Chinese coercive pressure is one thing; understanding the *risks* of a strategy adopted to counter it is something else.

A second challenge is that outcome-oriented risk analysis will have to take into account—indeed often be built largely around—the reactions of others. Outcome-oriented risks are easier to calculate when based on linear and endogenous factors—the risks of flying aircraft beyond their service life, for example. Things become far more complex and nonlinear when a strategy merely initiates a series of highly unpredictable interactions. And yet there is no way around the requirement for interactive analyses. This is one of the leading factors in the framework, offered below, for capturing key criteria for outcome-related risk.

Criteria for judging outcome-oriented risk

Once the focus is on outcome, a risk process should create an objective set of criteria that various parts of the organization can use for consistent evaluations of strategy consequences. Part of the discipline of a risk process, in fact, is that it will force leaders to confront a number of standardized questions and thereby test their assumptions about a given strategy. The purpose, again, is to

ensure a more informed strategic judgment along at least three measures: The quality and comprehensiveness of the information made available to senior leaders; the comprehensiveness of the set of risks considered; and the depth of discussion about risks and their potential mitigation.

The solution is not a common stoplight chart, so often employed in national security risk analyses, to grab a few basic categories and offer broad judgments about risk levels—“low, medium or severe” risk, for example. Such ordinal scoring methods are frequently counterproductive because they tend to oversimplify what they are trying to portray.⁵ Ordinal representations of risk—based on a few simple, subjectively defined levels without clear distinctions—often magnify the impact of individual perspectives, exacerbating the problem of risk subjectivity and personalization. They can introduce a false sense of rigor and specificity when ambiguity remains under the surface. Just as frequently, the conclusions represented by such charts are methodologically unsound: The factors that generate final risk numbers of representations reflect dangerous assumptions about cause-effect dynamics. The result is that, “In practice, they may make many decisions *far* worse than they would have been using merely unaided judgments.”⁶

In developing a better, more general and qualitative framework for organizing thinking about risk, the first step is to define the risks themselves according to various broad categories. These provide ways to thinking about things that could go wrong with a strategy in a number of areas. The framework suggested here is only a rough approximation, an invitation to further refinement, but it could include the issues outlined in Box 11.1. These factors could of course be broken down into dozens more: The variables that could impede success, for example, could be cultural, economic, political, and otherwise. But they provide a strong starting point for the analysis of specific outcomes that should be borne in mind when assessing risk.

From a structural standpoint, a framework like this could be included in formal policy processes, either at the National Security Council or Defense policy level. The process would mandate a “risk-based consequence assessment,” and an independent assessor—someone not directly engaged in the development of the strategy or policy options—would use categories like these to outline potential risks.

Over time, such a framework could be modified based on experience. New categories could be added. Risk analyses for specific issues could be updated and deepened with the benefit of events. A major purpose of such a coherent risk framework, in fact, would be to provide a platform for learning, both about the character and elements of outcome-oriented risk and the specific risks involved with particular issues.

The whole point of this process would be to generate critical discussions among senior decision-makers, to inform and deepen their strategic judgment. As we have seen, ultimately the sorts of issues under consideration in

Box 11.1 Categories of outcome risks

1. ***The reactions of other actors***, based on their interests, preferences, threat perceptions, and other factors. Examples could include:
 - ✓ Customers' reactions.
 - ✓ Investors' or capital market reactions.
 - ✓ Opportunities that will be created for other actors on the basis of the strategy—for good or ill, competitors and allies.
 - ✓ Reactions of partners or allies.
 - ✓ Reactions of competitors, rivals, or adversaries.
 - ✓ Employee, staff, or public reactions, especially in terms of credibility or reputation.
2. ***Possible dangerous outcomes from trade-offs or opportunity costs*** in the form of initiatives or requirements that will have to be downgraded or abandoned because of the resources devoted to the proposed strategy.
3. ***Factors that could interrupt the intended effects*** or objectives of the strategy, including things which might render it counterproductive (lack of feasibility is one form of risk). This category considers sub-optimized gains from strategies as well as potential losses.
 - ✓ Opponent/target capabilities.
 - ✓ Reactions different from those expected.
 - ✓ Institutional/operational ineffectiveness.
 - ✓ Legal, ethical, or regulatory compliance issues.
4. ***Ways in which the strategy could trigger or exacerbate threats in strategic environment*** that are dormant or not currently a problem. These could include stubborn cultural, social, or political dynamics that the strategy would run across.
5. ***Possible indirect or second-order effects*** triggered by the strategy.
 - ✓ Negative effects on other strategies: Reputational, resource or other impacts that undermine other efforts.
 - ✓ "Reactions to reactions": When one actor responds to your strategy, it may trigger secondary, indirect reactions by others.
 - ✓ Cascading effects through systems.
6. ***Comparative dangers: Risk relative to other options***. Evaluating the risk of a specific strategic choice never occurs in a vacuum; it must be undertaken relative to other options. Therefore a necessary element in outcome-oriented risk assessment is to compare and contrast the risks of available options to achieve US strategic objectives.

this study, complex strategic judgments, can only be resolved by qualitative, and often largely intuitive, discernment. The decision will amount to a sort of educated guess based on an emergent, feeling-of-knowing sense of rightness. But such emergent judgments can arise flippantly or with rigor and serious debate, especially with regard to analysis of consequences. An effective risk management process, as described here, could be judged on its ability to leave senior decision-makers without significant regrets in the list of possible consequences they considered and the manner in which they did so.

One implication of this analysis is that risk assessments should deemphasize specific levels of risk in favor of more general phrases that convey the kinds of risks associated with a strategic move, and qualitative statements (including evidence) to judge the likelihood and severity of those risks. The more that outcome-oriented risk processes can encourage a focus on broad themes and general categories of risk, the more they are likely to spark the right kinds of conversations. The goal of the process should not be to provide senior decision-makers with an objective estimate of likely risk, but to force them to confront things that could go wrong with their intended plans.

12

Outcome Assessment of the Emerging US National Security Strategy

This chapter will offer an example of the sort of outcome-based risk assessment highlighted in the previous section by examining the emerging US grand strategic approach, which could be called “selective engagement.” It first describes this strategic posture, and then informs its analysis with two previous case studies of similar strategies: the gradual British recognition of a need to wind down their empire, and the Nixon Doctrine. The discussion then undertakes a brief risk assessment of the emerging US approach along the lines suggested in the previous chapter.

The purpose of this chapter is not to offer a comprehensive risk assessment of the emerging US strategy. It is primarily an illustration of the sort of outcome analysis that ought to inform strategic choices. As suggested previously, most national security documents and strategies tend to deal with risk either in the form of threat or as a sufficiency analysis of the balance between ends and means. This chapter uses the two case studies to illustrate ways in which states manage risk during a period of strategic restraint, and then concentrates on three major themes of outcome-oriented risk relative to the current strategy. The goal in part is to demonstrate how elements of the framework offered in the previous chapter can be used to inform risk analyses.

In this analysis, the categories of possible outcome risk suggested in the previous chapter—reactions of others, dangerous outcomes or trade-offs, factors that could impede effectiveness, potential to exacerbate other threats, indirect or second-order effects, and comparative risk—are employed as a general guide to analysis. Each of the proposed risks below falls into one of those categories, even if the assessment is not grouped in such a way. Those five factors serve as a spur to thinking rather than a rigid framework.

The emerging default strategy

Various scholars have suggested a range of grand strategic concepts that could be adopted to guide US policy over the coming decades.¹ From among a

range of possible options—such as isolationism, offshore balancing, unipolar concert² and primacy—the current US national security posture is characterized by a different concept, and seems likely to continue to be guided by it. From the last years of the Bush administration through the Obama foreign policy, the United States has pursued some variant of what has been called “selective engagement”: an effort to remain engaged in the world as a leader and *primus inter pares*, but doing so with greater selectivity and discrimination. Selective engagement is a catch-all term for a whole range of strategies that basically amount to ways of preserving a forward-postured US global role on the cheap, with efficiencies in various aspects of the strategy.³

The over-arching strategic challenge for the United States is how it proposes to sustain substantial, and indeed largely unchanging, ends with fewer means. This demands a broad approach of selectivity, efficiency, and innovation in the “ways” by which the ends are served. Some forms of selectivity could be very limited, others more elaborate. There are many potential ways to more efficiently match ends and means. Versions of this approach might reflect slightly different answers to the key choices confronting the architects of US national security strategy: whether to focus on current technologies or future ones, whether to deploy forward or rely on strategic strike and rapid reaction forces, and to what degree to rely on allied and partner efforts to handle local security issues. Some would be very explicit about the commitments the United States would reject, others would be more circumspect.

A recent study group summarized the basic approaches of a more selective and discriminate approach as including three basic approaches. Such a concept would pursue “targeted, catalytic areas of competitive advantage, both capabilities and practices, that the United States can uniquely bring to the table to help meet challenges, while retaining a broad base of capabilities at lower levels to act alone when necessary.” It would defend “existing goals and interests in more innovative, selective, and asymmetrical ways, which involves becoming more discriminate in how our concepts envision the application of power.” Finally, a more discriminate approach would sustain stability by “Enabling and spurring others to do more in the combined, multilateral approach to crises, conflicts and persistent challenges.”⁴

It appears that these theoretical concepts have begun to work their way into US policy. Official national security strategy documents have tended to broadcast messages of continued US global engagement, and have downplayed the more discriminate aspects of the strategy. This is understandable from a political and even strategic standpoint, but it leaves us with few official characterizations of the actual strategic concept that is emerging. Analysts and observers have had to infer the real elements of the strategy, which in all likelihood has not been written down in comprehensive form anywhere, from US actions. And those actions point to a complex balancing act between continued global leadership and intensifying limits.

It is therefore possible to identify enough elements of the emerging strategic approach to define it for analysis. It is an approach that:

- Continues to insist upon a substantial and engaged US global role, not a policy of complete disengagement or “over the horizon” balancing.
- Recognizes the limits inherent to US power and embraces a more multi-lateral and shared vision of global governance.
- Invests significantly in the capacity of friends, allies, and partners.
- Attempts to find discrimination in the US global role in terms of the promises and threats covered by US strategy.
- Accepts smaller standing armed forces in exchange for other values such as fiscal prudence and investments in modernization.
- Tries to achieve greater leverage and close the ends-means gap with innovative ways—operational concepts, doctrines, and approaches that get the biggest bang for the buck out of US strategic effort.

This is not the first time that a global leader confronted emerging constraints on its power and influence. This chapter now turns to two cases in which previous leaders—in Great Britain and the United States—had to develop their own versions of discriminate strategies. The lessons of these cases can inform how we think about the risks involved in the emerging US approach.

Case studies: Great Britain contends with relative decline

Great Britain was the predecessor to the United States as the leader of a global system of liberalism and free trade and sponsor of modern institutions across the globe. By many measures, it enjoyed global dominance for the better part of a century—until London faced challenges of declining relative power and rising peer competitors, and had to manage a gradual decline of its role over the course of a century. The scholar Walter Russell Mead has argued that “Britain’s world role in the 19th and first half of the 20th Century, does, in fact, remain the sort of only really helpful comparison for Americans trying to get some perspective on what we’re doing, or not doing, in the world.”⁵

Even by the middle of the nineteenth century, cracks had begun to appear in the façade of British dominance. By 1830 Britain generated almost 10 percent of world manufacturing on its own, and two-thirds of world trade in manufactured goods; by the last decades of the nineteenth century, however, British economic growth slowed to just about 1.5 percent a year.⁶ From that point to the end of the century, Britain would consistently slip in relative global dominance. In 1870 the British economy accounted for a quarter of world totals; by the end of World War I it was still a less remarkable but robust 15 percent.⁷ Geopolitically, Britain confronted a number of emerging competitors, led by Germany and the United States. Others claimed the right

to shape and lead the international system and set out to acquire local spheres of influence.

“What was different now” compared to prior eras of challenge, Paul Kennedy has written, “was that the relative power of the various challenger states was much greater, while the threats seemed to be developing almost simultaneously.” And the scale of the challenge was such that simply throwing money at the problem could no longer suffice: “No matter how regularly the Royal Navy’s budget was increased, it could no longer ‘rule the waves’ in the face of the five or six foreign fleets which were building in the 1890s.”⁸ These trends were sufficient to demand a strategic readjustment, as Britain confronted a global posture that was increasingly insolvent to meet its needs.

A critical parallel with the United States is that, in Britain’s liberal political culture, the resources devoted to national security were relatively small even at the apex of empire, and far less than more militaristic societies routinely devote to the military. On the eve of World War I, the role of government in Britain was so modest that the sum of national and local government spending totaled just over 12 percent of GNP. By 1933, even as tensions soared in Europe, London devoted just over 10 percent of its national budget to defense.⁹

These and related trends confronted British leadership with an increasingly untenable strategic posture. Britain’s former dominance was being eroded in ways that could not be countered without self-destructive levels of exertion. As Aaron Friedberg had described the situation, “The nation appeared to have its neck in a gradually tightening noose from which no easy escape was possible.”¹⁰ The question was what strategic concept Britain adopted to deal with this changing strategic reality. “Either it had to come to terms with its new rivals or it had to fight them. In the end it did both.”¹¹

British responses: a strategy for strategic solvency

The British strategy for dealing with relative decline involved off-loading burdens in three particular ways. The goal remained the maintenance of a world order favorable to British interests, one that would keep Britain safe.

First, London transferred some degree of the burden to broad systemic norms and institutions. Even if British power were to decline, British interests would thrive in a world where its values continued to thrive. In a manner of speaking—and reflecting some of the explicit debate that went on at the time—it traded a formal empire for an informal one. Jean-Marie Guehenno has described the globalizing set of exchanges, values, and expectations as an “empire without an emperor”¹²—which is what the architects of British strategy had in mind.

Second, London engaged in the selective shedding of commitments, along with a series of concessions and appeasements designed to ease its global

burden. London increasingly focused the attention of its Navy and other military power closer to home, on European concerns. But this disengagement was selective; England remained committed to imperial control in Egypt and, for a time, India as well. "It would fight for certain obvious aims," Paul Kennedy has written—"the defense of India, the maintenance of naval superiority especially in home waters, probably also the preservation of the European balance of power—but each issue had to be set in its larger context and measured against Britain's other interests." The result was a constant series of balances and compromises that "made Britain's policy frustratingly ambiguous and uncertain" to other actors on the scene.¹³

Third, British policy-makers set about transferring responsibilities to the rising world power most in line with its own values: the United States. In order to make this possible, Britain had to concede a sphere of influence to a rising competitor, and gradually step aside as the newer global player assumed more and more influence. This was hardly an easy decision for the leaders of a global empire who viewed their own judgment as nonpareil, and it took place in a context—easy to forget today—in which Washington and London still viewed each other warily. But broad factors in the strategic environment pushed the two toward rapprochement, and by the mid-1890s, London was signaling that it would tolerate US interventions in Latin America.¹⁴

The gradual retreat from empire was not simple or unidirectional, and it did not affect all issues equally. Where Britain saw continued essential interests in slices of its empire—in the Middle East, for example, with its oil supplies, or parts of Southeast Asia—it remained fully committed to the principle of British domination and paternal rule. When it saw the need it would intervene casually and sometimes ruthlessly to put down revolts, from Africa to Afghanistan. British thinking did not culminate in a single moment of realization, and global preeminence and empire remained central to the self-conception of the British ruling class. Its selectivity and concessions aimed to preserve the most sustainable global role, not abandon it.

Nor did the implicit agreement to share global responsibilities proceed in a straight line on the American side. Well into the twentieth century, the United States—accustomed to a comfortable stance as an isolationist power protected by two flanking oceans—remained uninterested in the role London had mapped out for it. After a brief flirtation with a more muscular global role during World War I (energized in part by Woodrow Wilson's crusading internationalist ideas), the United States pulled back into its self-protective shell. For historical as well as temperamental reasons, the British empire did not fill all American policy-makers with awe. When Americans hear complaints today about rising powers hesitant to assume responsibility or leadership in the international system, they should keep in mind that the United States took a good half-century to become accustomed to its growing role—and made a full transition only when forced to do so by an attack from a revisionist state determined to upset the global order.

But the basic choice for Great Britain was clear and fundamental, and dictated by the strategic terms of the day. By the 1930s, simply escalating defense expenditures was no longer an option: Massive growth in defense budgets put the British economy under severe strain, affecting currency reserves and its balance of payments. Meantime a series of hurried military commitments to place deterrent checks around Hitler's Germany promise far more than the British military could deliver; overstretch was a significant reason for the early setbacks of British forces in World War II, spread so thinly around the globe as they were.¹⁵ And eventually the United States, under Franklin Roosevelt, became a co-conspirator in the dismantling of the British empire, a mindset that perhaps reached its apogee in the administration of Britain's former wartime partner, Dwight Eisenhower, whose instinctive sense of anticolonialism was powerfully strong and whose reaction to the British Suez adventure put the final nails in the coffin of British aspirations for persistent imperial activities.

Eventually the hand-off became formalized in the post-World War II series of agreements, built around the Atlantic Charter, in which US and British leaders worked together to design a series of institutions that would help to govern the post-war world, institutions that the United States had by now admitted that it would have to lead. A shared political culture was critical in allowing London and Washington to build a sense of implicit trust necessary for the relationship both eventually desired.

Britain manages risk in strategic choices

In retrospect Britain's basic choice—to step gradually, even haltingly aside while its strategic protégé, the United States, came to dominate free-world political leadership—seems obvious, even foreordained. But it was hardly so at the time, in part because the risks of the course were substantial. It reflected a clear choice to manage risk in part by positioning Britain as the benefactor of a hopefully stable renewed world order built in support of liberal economic and political principles. In this regard it was highly successful.

One lesson of the case is that managing risk at the strategic level is partly about defining a specific role for the state, one that allows it to exercise disproportionate influence while mitigating the negative consequences involved in the chosen strategy. Even without its former preeminence, Britain could play a critical balancing role in Europe and, through its "special relationship" with the United States, shape the views and behavior of the rising power. It would be wrong to portray this concept as fully thought-out from the beginning. British leaders went through many sometimes grasping stages of redefining their role, not always willingly, and when they began the process relations with the United States were anything but special. But a clear element in the advantage sought from the new arrangement was a position of continued influence for Britain beyond what its objective power would suggest—an

effort to structure the system to that London could regularly “punch above its weight.”

A second lesson is that the foundation for the management of risk under uncertainty is a sustainable foreign policy grounded in vibrant domestic social, economic, and political realities, and this sometimes demands accepting inexorable facts of strategic life. The most obvious seeming risk of the approach was that Britain would lose substantial global influence as it vacated its position as hegemon. In fact, though, this risk was really more of a simple reality: Britain’s influence was destined to decline based on absolute trends in the system. The case suggests that strategic actors should distinguish risks over which they have some influence from the inexorable tides of history.

An important contrary lesson of the case is the powerful difficulty of effective geopolitical signaling as part of a strategy of managing risk. Even carefully designed efforts to prioritize will not necessarily send the desired signals. Britain in the two decades before 1914 clearly chose to focus on European interests and deployments, for example, and establish itself as the linchpin of European security balances. And yet the war it confronted was not a product of failing to emphasize the periphery—it was an aggressor in Europe. Managing risk in security affairs is a process rife with potential misunderstandings and misperceptions.

A related lesson of this case is that steps that might be perfectly safe in one context can be hazardous in another. The aftermaths of economic crises, in particular, can be perilous times. As Charles Kupchan has argued, the Great Depression had a hand in sparking both a virulent nationalism in Germany and “underbalancing” by an economically constrained Great Britain, a combination that proved dangerous.¹⁶ The most effective strategists begin all their thinking with a clear sense of what sort of age they are living in.

Indeed, a less charitable interpretation could argue that Britain’s more discriminate approach contributed significantly to both world wars. Its gradual step-back from preeminence and empire led British leaders before 1914 and 1939 to decline the role of large-scale balancer in Europe. British governments repeatedly failed to commit large ground forces to the continent, and thereby allowed two generations of German leadership to believe that it could succeed with offensive campaigns. But this critique overstates what could ever have been expected from British military power. Even at the height of empire, Britain remained a primarily maritime, constabulary, and commercial power; its political culture never endorsed massive standing land armies. As late as 1883 the British army only numbered about 125,000 men, and by 1901, as ranks grew to fight the Boer War and for other purposes, it swelled to just over 400,000. These numbers compared with later, hastily assembled armies of almost 4 million to fight World War I and just over 3 million during World War II. Britain was simply never in a position to offer forward deployed forces that would have made a decisive difference in European balances of

power. That constraint was a built-in limitation of Britain's grand strategy, not a strategic error of weakness or unwillingness to spend more on defense.

Britain's experience suggests another lesson: Managing risk in national security demands attention to global norms and institutions. Just as individuals deal with uncertainty in part by relying on shared understandings and norms, so do states and other collective actors. Expectations, rules, and institutions can affect behavior and create patterns that determine how stable or threatening a context will be. At various stages, whether working within the European balance of power or in the post-war international environment, British leaders sought to manipulate institutions and regimes to create competitive advantage and promote the safety of British interests.

Another lesson is simply that Britain got extraordinarily lucky, specifically in the potential attitude and policies of its chosen hegemonic successor—the United States. Of all the potential successors, its values and interests were obviously most in accord with Britain's. London enjoyed substantial trade relations with the United States; neither had any territorial claims affecting the other; and most importantly, the two political cultures grew from the same heritage. By 1918, the relationship had been forged in conflict, as the United States entered World War I. The risk posed by rising American power was relatively minor, it seemed—and certainly posed fewer risks than any other course.¹⁷

Just how lucky Great Britain was in the character of its successor becomes obvious from the experience of another relationship Britain decided to cultivate even as it was expanding its ties to the United States, its episodic alliance with Japan. In World War I, British-Japanese ties were well-enough established that Japan agreed to attack German positions in the Far East and provide other support to the Allied war effort. But the relationship was always uncertain, in part because the two sides simply did not understand each other well. And of course Japan ultimately rose into a modern, belligerent power that turned against its one-time ally and launched a massive war of aggression, in large measure against British territories. The American relationship is the more widely told story, but Britain's efforts to embrace Japan as another potential successor in the international community show the limits and dangers of the approach.

But that was the essential challenge of the situation British policy-makers found themselves in—the limits of choice. Britain's efforts to manage uncertainty for competitive advantage was significantly a story of the risks it faced if it did *not* attempt to attenuate its global posture. This is why the British choice seems so self-evident in retrospect: Trying to sustain global dominance would have been foolhardy.

This may be a profound parallel for the US case today. Britain's example suggests that the corresponding US effort to manage risk with a strategy of discriminate power promises neither in a new rush at dominance or a selfish

isolation, but rather a constant, dilemma-strewn embrace of restraint and limits, and the development of clever approaches and strategies that bridge ends and means. The problem for the United States today, of course, is that there is no obvious successor to whom it can safely pass a substantial degree of global responsibility. The United States appears to be trying to solve this problem by passing some responsibility, not to a single successor, but to a large number of allies and partners. The difficulty then becomes bequeathing not just responsibility but also authority—something that Britain had eventually to do, and that the United States has so far been reluctant to try.

Managing uncertainty in the Nixon era

The situation confronted by the Nixon administration as it moved toward its second term had many similarities to the present moment. The United States was in the process of trying to wind down an exhausting, extended conflict. US leaders felt pressed on many fronts around the globe. And at home the administration faced economic challenges that led to fears of a growing mismatch between foreign policy obligations and the resources available to meet them. In essence, the United States confronted large and growing gaps between ends and means in strategy, and needed to take some action to reverse a widespread loss of credibility.¹⁸ The United States, in the words of one commentator on the doctrine, was “no longer in a dominant position, but still expected to lead.”¹⁹

To deal with its basic strategic situation, the United States had a number of options—accelerated global confrontation, generalized withdrawal, or a more selective approach that attempted to sustain existing ends with different “ways.” Richard Nixon, as Henry Kissinger would later explain, attempted to balance far-flung global commitments with the realities of world politics and limits on American resources. He “saw it as his task to define a *sustainable* role for an idealistic America in an unprecedentedly complex international environment.”²⁰ The Nixon administration selected the latter approach; as Great Britain had done, it emphasized the role of allies and partners (as well as persistent emphasis on nuclear deterrence) in closing the ends-means gap in US strategy. Because of the similarity to our own strategic context, the Nixon case can perhaps offer some useful lessons about the risks involved in a broad strategy of selective engagement.

The doctrine and its context

President Richard Nixon first announced the principles of the doctrine that would bear his name during a press conference on the island of Guam on July 25, 1969, and later described it more elaborately in a speech on November 3 of that year. The speech in November was also known for announcing the administration’s policy of “Vietnamization,” and the doctrine itself reflected

not so much a generally developed grand strategy as a technique to come to grips with the specific challenges the war had produced. A key lesson of the war for frustrated US strategists was that the United States could not make itself responsible for other nation's problems. If a partner, especially one engaged in a counterinsurgency or state-building process, would not commit itself to needed reforms, no amount of US power could close that gap of commitment. In trying to do so, Vietnam had taught, the United States risked exhausting itself, in financial as well as domestic political terms. Washington needed a doctrine to promise a leading global role without taking on burdens it could not sustain, either politically or economically. It is therefore no surprise that the essential concept in the Nixon Doctrine is a demand for US friends, allies, and partners to take responsibility for their own security.

The Nixon administration had considered a number of strategic options to deal with this situation, at least in passing. The idea of simply getting out was appealing on some levels, especially as Washington decision-makers became increasingly frustrated with the conduct of a string of South Vietnamese regimes. But Washington rejected withdrawal for obvious Cold War reasons of credibility: Lyndon Johnson had internalized the logic behind the "domino theory," and he and Richard Nixon believed that "cutting and running" would only encourage more aggression.²¹ At the other end of the spectrum, escalatory options were also on the table, particularly for the Nixon administration. But more dramatic possibilities risked conflict with China, and ultimately were rejected because the interests at stake did not justify a potential nuclear confrontation or land war with Asia's geopolitical giant.

It was in this context that Nixon announced his doctrine to enunciate "the principles that would guide his country's new approach to international relations" in the post-Vietnam era, which would spell out the "criteria for involvement abroad."²² In the November speech, Nixon described it as a policy which "not only will help end the war in Vietnam, but which is an essential element of our program to prevent future Vietnams."²³

The doctrine had three core principles. First, Nixon reassured the world that the United States would preserve its treaty commitments. Second, he promised that US extended deterrent pledges remained in force. Washington would continue to threaten to escalate in service of commitments to NATO or Japan or Korea. But third and finally, non-vital threats would increasingly become the responsibility of US friends. "In cases involving other types of aggression," Nixon's speech went on, "we shall furnish military and economic assistance when requested in accordance with our treaty commitments. But we shall look to the nation directly threatened to assume the primary responsibility of providing the manpower for its defense."

The core choice in the doctrine, then, was to distinguish the interests involved in various types of contingencies. The United States would remain involved in such conflicts in the name of global anti-Communism, but it

would not assume the dominant burden as it had in Vietnam. The doctrine was thus trying to walk a fine line—to reaffirm US commitments to fulfill true treaty obligations while backing away from a promise to fight more Vietnams. It reflected a rebalancing of ends, ways, and means to create a more sustainable foreign policy. As Nixon would put it in February 1970, “our interests, our foreign policy objectives, our strategies and our defense budgets are being brought into balance—with each other and with our overall national priorities.”²⁴

Henry Kissinger has readily admitted that there were holes in the doctrine’s criteria. What really counted as a “non-vital” interest? “Reality,” he wrote simply, “showed itself resistant to being thus encapsulated in formal criteria.” The pledge of respecting commitments was “boilerplate,” he explained, and really pointless since everyone knew the United States still enjoyed the flexibility to abandon any commitment it wished. The doctrine made a somewhat tortured distinction between the US response to internal versus external aggression, when the whole problem in Vietnam had been the difficulty of telling the difference. Because it reflected no real alteration of leading commitments, then, the new, more discriminate approach mostly applied to secondary challenges. “In its attempt to devise a ‘doctrine’ for avoiding another conflict like Vietnam,” Kissinger has written of the administration, it “developed a doctrine which applied primarily to situations like Vietnam which it was determined not to repeat.”²⁵

The doctrine emerged in a larger strategic context: The Nixon administration’s growing view that the world was becoming more multipolar, and trying to sustain US primacy was a losing gamble. Nixon announced the doctrine alongside a number of contemporaneous speeches and statements making clear his view that the future of grand strategy lay in a balance of power among a number of rising power centers.²⁶ The practical result was the beginning of a transition to Vietnamese forces of the burden of the fight, with concomitant reductions in US combat operations.

Nixon manages risk

The Nixon Doctrine and associated efforts—including the famous opening to China, which emerged from the same broad assessment of the demands of a changed context—reflected another sophisticated example to adjust to changing global circumstances. The Nixon administration perceived an increasingly complex, potentially threatening international situation alongside a necessary and inevitable contraction of US means and global influence. This was not an outcome that could be contested: The historical record, in fact, suggests that the administration felt it had no other meaningful choice. Neither withdrawal nor escalation were palatable, and yet the course they were on could not be sustained. Much as today, then, the administration backed into a new strategic doctrine more out of necessity than choice. Conducting an elaborate

risk assessment may have seemed somewhat beside the point. The question was then how to create the basis for enhanced competitive positioning under uncertainty, and to manage the risks of the resulting strategies.

At the time the Nixon Doctrine confronted withering critiques over its sufficiency and feasibility. Critics charged that the doctrine did little to revise the ends or commitments of US national security strategy; it merely cut back on the means able to sustain them. The strategist and scholar Earl Ravenal, writing in *Foreign Affairs*, argued that “the Administration’s new policies and decision processes do not bring about the proposed balance; in fact, they create a more serious imbalance. Essentially we are to support the same level of potential involvement with smaller conventional forces. The specter of intervention will remain, but the risk of defeat or stalemate will be greater; or the nuclear threshold will be lower.”²⁷

One lesson of this case, however, may be that the relationships and causalities involved in issues such as deterrent relationships are far from linear or predictable. More discriminate strategies can sometimes manage risk by stretching the necessary claims a good deal without sparking deadly challenges. The result of the Nixon recalibration was not a global tide of Communist aggression, in part because, as Nixon and Kissinger well knew, the incentives and motives involved in each case were highly idiosyncratic. Managing uncertainty in a situation of necessary strategic restraint, then, is about playing to the nuances and complexities of issues like deterrence and credibility.²⁸ A possible lesson of the Nixon experience, then, is that US credibility is less fragile than commonly assumed: The United States can scale back its commitments and security posture to some degree without breeding a collapse of the global system.²⁹

Managing the risks of a more selective global posture can also, however, involve bold and dramatic moves to change the overall balance of power in favorable ways. This was where the opening to China came in, a masterstroke of managing uncertainty for competitive advantage, tilting the geostrategic situation heavily in the US favor without any new expenditure of resources. In this respect it has much in common with Britain’s strategy of passing preeminence on to the United States: In both cases, a state sought to shape the global balance of power with large, transformative gambits, in part to make up for and camouflage encroaching strategic insolvency.

But the Nixon approach involved a number of problems and dilemmas as well. One lesson for the future may be that any attempt to rely on the efforts of others, even if close allies, will always be unsatisfying. In the United States, support for aid and engagement of allies will seldom prove a sustainable approach. It became difficult to budget for the partner development mechanisms essential to the Doctrine’s success. Meantime relying on efforts of partners reduced US influence and leverage, but not US commitment if something

went truly wrong. The Doctrine embodied a constant risk of being sucked into conflicts of allies' making.

Another problem was that the Nixon approach remained firmly committed to the global contest with the Soviet Union, and to the full scope of America's global interests. In this regard it offers another lesson for the present context, one having to do with the inconsistencies and dilemmas of a concept that both constrains the US role and is unwilling to step back from its logical consequences. Earl Ravenal saw the problem in 1971. The Nixon Doctrine's idea that Washington would rely more on partners to fight its wars, and thus avoid undesirable commitments, was really nothing more than "a postulation that the unwished contingency will not arise," he argued. Because at the end of the day the same question would press itself on US strategists: If the issue at stake was important, what would the United States do if its partners could not do the job?³⁰ It would either be forced to abandon the effort, or to double down.

The only way out of the dilemma, Ravenal pointed out, would be for the United States to clearly articulate the interests at stake, and reconsidered the scope of US Asian commitments—that is, as Britain had done, to formally shed commitments. This the Nixon Doctrine refused to do. The one basis for restraint built into the concept—the claim that the United States would only intervene in a big way when internal subversion metastasized into "external aggression"—was too ambiguous to provide clear guidance. As long as generalized containment of China remained the goal, US strategists would *inevitably* view all local subversions as forms of Communist "aggression." This was, after all, what had brought about US involvement in Vietnam in the first place. "In this respect," he wrote, "the Nixon Doctrine does not improve on the policy that led to Vietnam." That same policy of global commitment would draw the United States into more conflicts in service of half-hearted interests. In the most fundamental terms, this dilemma remains unresolved to the present day.

Assessment of risk: outcome-oriented risks of selective engagement

In this analysis, the categories of possible outcome risk suggested in the previous chapter—the reactions of others, dangerous outcomes or trade-offs, factors that could impede effectiveness, potential to exacerbate other threats, indirect or second-order effects, and comparative risk—are employed as a general guide to analysis. Each of the proposed risks below falls into one of those categories, even if the assessment is not grouped in such a way. Those five factors serve as a spur to thinking rather than a rigid framework.

The following analysis describes a number of sample risks, assesses them using elements of evidence, and evaluates possible mitigation strategies for each. It does so as an example of the sort of analysis encouraged by outcome-oriented risk assessments. It does not attempt to be comprehensive, though

the assessment does point to some essential findings about the emerging US strategic posture.

Some general risks to US national security strategy are not covered below because they are not risks that apply *specifically* to this choice. As an example, the United States is amassing hundreds of billions of dollars in the deferred costs of hugely expensive US modernization programs.³¹ Meantime unfunded liabilities in Pentagon personnel accounts have reached several trillion dollars and by about 2040 could balloon to levels that would submerge the entire defense budget.³² Left unchecked, these bow-wave liabilities will crowd out every other category of national security expenditure in coming years. Another broad danger is the loss of technological overmatch, which would endanger US military dominance. But these dangers apply no matter what approach the United States takes.

One of the most profound lessons of the British and Nixon cases outlined above is that outcome-oriented risks in national security are very difficult to anticipate. As suggested in earlier chapters, the causalities in complex strategic judgments are so ambiguous, the nonlinear dynamics and transmutability of the contexts so profound, that objectively “pricing” outcome risks becomes effectively impossible. The purpose of risk assessments then becomes to inform strategic judgment by providing essential information, ensuring that senior leaders do not ignore key consequences, and serving as a forcing function for a rigorous dialogue about the risk-reward calculus of various strategic options. The brief and suggestive survey of risks below hints at the sorts of issues that would be evaluated in such an assessment. These categories, though, only offer hints. A full-blown risk assessment would go into much more depth.

Potential risk: inability to deal with multiple contingencies leading to opportunistic aggression by revisionist actors

With smaller armed forces and persistent global responsibilities, a more discriminate strategy could create situations in which the United States cannot avoid engaging in a number of conflicts or interventions at the same time. These could overstretch the force and produce longer wars and even threaten defeat. The risk of the emerging posture is not that the United States could not undertake one or more missions, or credibly deter in multiple regions. The problem is that it could not do many things at the same time,³³ which then points to a potential risk that aggressors would be tempted by growing insolvency in the US posture to take advantage of US distraction if it were engaged in a contingency elsewhere.

Yet the United States has confronted this risk many times since 1945 without bringing about the feared opportunistic aggression. During the first Gulf War, for example, virtually the entire US active Army was deployed into the Middle East, and it did not tempt others to aggression. During the Nixon

administration, the United States quite explicitly stated that it did not have the forces for a full two-war criterion. Nor is the relationship between capabilities and contingencies linear: US war planning models may dictate a certain number of forces as “required” for a given contingency, but in the breach the United States could, and would, do what it could with what it had. An aggressor would rarely confront a situation in which it could be confident that the United States would be simply unable to respond.

Potential adversary interests and preferences will exercise a decisive role in determining the true degree of risk on this issue. States and non-state actors will have many motives for considering aggression; the degree of perceived US strategic solvency will be only one. Many actors have staged attacks or adventurist gambits when the United States had more than enough capacity to respond—they merely engaged in wishful thinking about the likely consequences of their actions. Most states have powerful reasons to avoid large-scale aggression; the potential for a decisive US military response is an important one, but it is only one of these reasons. Exactly what effect a slightly lower profile would have on such multifaceted calculations is very difficult to assess. Given the multiple costs and risks of such adventures, in fact, there is strong reason to believe that the default tool for revisionists to claw their way toward their goals will be the use of gray-area, hybrid, asymmetric, and salami-slicing tactics that do not represent decisive steps and remain under the threshold of US and allied red lines and tripwires.

A number of mitigation steps seem readily available to the architects of US strategy to dampen the risk of opportunistic aggression from a more selective strategy. One example is efforts to enhance allied and partner contributions and pledges to help fill the gaps in deterrent power. In all cases of potential aggression, the United States is not the only one being threatened—indeed, it is usually in the background. There are, as in the Nixon case, real limits to what allies and friends seem willing to do, and their interests do not always accord with those of the United States. Nonetheless, when compared to any potential adversary, the United States enjoys the huge advantage of a network of like-minded friends also threatened by potential aggression. The more direct the threats become, the more these others might contribute to common defense. A shared responsibility for global leadership, under the guiding influence of agreed norms and institutions, is the fundamental mitigation response to dangers of aggression posed by a more discriminate US strategy.

Potential risk: injury to US credibility

A second potential risk of a more discriminate approach is that others could come to doubt the US willingness or ability to continue to lead globally or to enforce key norms of international order. In outlining this risk we must distinguish the general trend of such thinking, which is already well underway, from the risk to a strategy of selective engagement. A reduction in the

perceived credibility of US commitments is *inherent* to the current situation. The question is the relative risk part of any given strategy and what the United States can do to mitigate it

Again there is historical precedent for a more limited US posture, and it does not suggest that credibility is likely to collapse overnight. As the Nixon experience suggests, the United States can partly counteract credibility risk in areas of selectivity with continued commitments and even bold initiatives in other areas. A major initiative to shift global power balances, some analogy to the China opening, could create a transformed situation in which doubts about US power would fade into the background. That case also suggests that leadership counts—the person of the US president will make a difference in the calculations of others, and may either exacerbate or compensate for perceived selectivity in the strategy itself.

Nonetheless, evidence from the last several years suggests that even modest US steps to walk back from a dominant global posture could produce exaggerated concerns on the part of friends and allies. Especially in Asia but also in parts of Europe, key regional partners have increasingly questioned the US willingness to sustain decades-old commitments. There is evidence that faith in US power and reliability is subject to significant and sometimes seemingly over-sensitive swings. The risk of a continued slide in this direction must be taken seriously in contemplating the consequences of a more discriminate posture.

In terms of the threat to US credibility, here again the United States has a number of potential mitigating actions. It can, for example, undertake significant diplomatic, economic, and informational initiatives to reinforce US leadership and influence: A more selective approach does not rule out vigorous leadership on the issues of priority. In strictly military terms, moreover, the United States could counteract any sense of unreliability by reinforcing its commitment to selected alliances and partnerships through restated or enhanced forward presence, security assistance, exercises, military-to-military contacts, training and advising missions, and other efforts.

Potential risk: cascading doubts in the rules-based order

To the extent that the United States is seen to be stepping back from decisive capabilities to confront both outright aggression and gray-area troublemaking, it could begin a process of unraveling of the existing order. More aggressors could be tempted to challenge elements of the order, and more potential advocates of norms and rules could be discouraged from supporting the system by a sense of US retrenchment. The risk is of a gradual, cascading collapse of international order.

Any more discriminate US strategy carries the danger of subtracting more elements of reassurance and enforcement from the system, and exacerbating the risk of an unraveling. This is, however, a highly qualitative danger, nearly

impossible to anticipate with any model or analysis. Historical experience suggests that the United States can step back from previous commitments, force structures or interventions without doing fatal damage to the system. The key question is the degree to which the present moment is more unstable, and poses greater dangers of such effects, than prior ones; and again, it is simply impossible to know for sure.

Mitigating actions could be focused on efforts to bolster the effectiveness of global norms and institutions—to employ the strategy of Great Britain and deepen the global commitment to a rules-based order to compensate for a relative decline in the hegemon's global role. The United States could, for example, undertake a number of significant new global initiatives to indicate continued leadership in non-military ways, including perhaps fissile material control and climate initiatives. More fundamentally, it could seek to multilateralize enforcement to a greater degree by offering leadership on various initiatives to others, to create a sense of shared responsibility—even, perhaps especially, when the outcomes do not match precisely the US goals or objectives.

Summing up this risk assessment of the current strategy, a United States committed to more selective and discriminate ways of pursuing its global interests should prioritize a number of specific categories of actions to help manage the risks of the approach. It must, first of all, address the long-term budgetary imbalances that could overwhelm the US national security posture. Second, it must work much more deliberately to build a broad range of non-military tools, and use them to generate global influence, credibility and deterrent power. And third, it must continue and build upon efforts to recruit energetic partners for key tasks from regional deterrence to addressing global issues like climate and fissile material control. These three efforts would help make a strategy of selective engagement more sustainable and secure, and would contribute to continuing US efforts to restore and sustain a global rules-based order.

An important message of this analysis is that, after more than a decade of focus on—even obsession with—the military instrument of power, the most fundamental answer to emerging risks with a more discriminate strategy is to maximize the non-military aspects of its power, and work to bring diplomatic, economic, informational, and social initiatives to the forefront of US strategy. This is where the United States has its primary competitive advantage, and despite the rhetoric of “whole of government” solutions, non-military instruments have been badly underutilized over the last decade. The risks of selective engagement will be exacerbated to the extent that the United States insists on focusing on military action to deal with the range of challenges it confronts, and will be eased to the degree that the United States can rediscover its global image and posture as a fundamentally non-military leader. Competitive advantage, as the last chapter concluded, more often derives from powerful geopolitical maneuvers than from simple military strength.

13

Principles of Effective Risk Management

Chapter 11 argued for an approach that placed risk management in service of a specific goal in the development of strategy: outcome assessment. This recommendation stems from two of the dominant lessons of the financial crisis: that the concept of risk had become fragmented and incoherent; and that many of the most common cognitive and framing errors push decision-makers toward urgent, nonconsequentialist thinking. Developing a single coherent approach to risk, and focusing risk analysis on outcomes, are the first steps toward the effective employment of risk as a component of a larger strategy process.

The next step in building an effective practice has to do with the actual practice of risk management in organizations—the habits, mindsets, and approaches brought to the task. The financial crisis suggests that this is the foundation of effective risk management: institutional, human, and cultural. This chapter surveys the basic qualities required.

Those qualities are the most essential response to risk, because this study's central theme has been that case studies of risk failures in both finance and national security show that risk is a human and organizational, not procedural or structural, challenge. When risk management fails it is usually a product of specific and identifiable human factors and cultural habits, which overwhelm even the most elaborate risk procedures. The fundamental response to risk, then, must not be expanded procedures or mechanisms, but efforts to insulate an organizational culture against the human hallmarks of risk failure. The most important lesson of recent experience is that institutions determined to manage risk well must pursue a comprehensive solution, shaping their institutional culture partly around the task. Managing risk well demands a complex combination of tools, judgment, culture, and resilience—what Deloitte has called a “risk intelligent enterprise.”¹ Specifically the lessons of recent experience point to at least five necessary qualities.

An analytically demanding organizational culture

As noted in the last chapter, if an organizational ethos isn't serious about risk, a firm can have elaborate processes and even a few key players who offer dire warnings, and it won't be effective. The challenge of effective risk management is in many ways a challenge of useful warning, especially with regard to widely appreciated but underestimated "gray swans." Creating a culture that tolerates warning and demands candor, as well as encouraging an in-depth discussion of risk, is the first and essential step to effective risk management.

The problem has been obvious in dozens of financial and national security cases. In the infamous case of Enron, to take a leading example, the company spent literally tens of millions of dollars on an elaborate risk architecture. It bragged to outsiders that the firm took no significant actions without a detailed assessment from the risk unit. And yet in practice, the culture of the organization overwhelmed the risk management structures and procedures with countervailing incentives. These included—as is common in such cases—an overpowering pressure for short-term profits (especially given the firm's Ponzi-scheme like financials), a reverence for deal-making over follow-up, hoarding information, unrestricted risk-taking, and more. Again, the fundamental lesson of this study is that the character of an organization will shape its risk profile, not the formal results of its risk analyses or the advice of its risk managers.

The first requirement, then, is for a corporate culture that helps understand and mitigate risk rather than exacerbate it. One challenge, of course, is that risk is the other side of the coin from reward. A culture should not be unduly risk-averse any more than it is unnecessarily risk-accepting. The requirement, as suggested by the case studies surveyed here, is for a culture of open dialogue, rigorous analysis, and support for dissent. It is when cultures and habits tend to ignore information, punish alternative views, and rush into ill-considered decisions that risk—and warnings—become discounted. Risk analyses should be designed to inform complex strategic judgment with enhanced information, clear identification of risks, and the requirement for in-depth debates about the variables and causalities involved.

In order to achieve this goal, it is critical to cultivate a culture of rigorous analysis and candor. Robert Kaplan and Anette Mikes suggest that the sort of hard-boiled confrontation so essential to real risk discussions are rare, and in fact an "unnatural act" for most human beings. They point to organizations that create rough-and-tumble dialogues of intellectual combat designed to ensure that risks are adequately identified and assessed. These can involve outside experts, internal review teams or other mechanisms, but the goal is always to generate rigor, candor, and well-established procedures for analysis.² The result ought to be a culture, set of habits and actual procedures to institutionalize what Jonathan Baron has called "actively open-minded thinking".³

The basic principles include a thorough search for information and true open-mindedness to any possibility, and avoiding self-deception through rigorous consideration of alternatives.

The famed social scientist Philip Tetlock has spent years examining the habits of thought and analytical practices that produce the most rigorous and accurate thinking. He has described the thinking process of the “superforecasters” in his good judgment tournaments: They are open-minded, critical analysts of information. They are always ready to challenge their own views and to change their minds when the facts dictate. They tend to break questions down into subsidiary components. They apply rigorous analytical techniques where appropriate, and update their conclusions with changing information.⁴ In brief, they take seriously the process of analysis, and demand close attention to every step in the thinking process.

Generating these sorts of mindsets throughout an organization means creating a supple and flexible risk process, one built around the consideration of many alternative scenarios. It means senior leaders who become extremely sensitive to a combination of surprising data and gut feel. When Goldman Sachs saw some risk numbers and losses in mortgage equities in 2007 that worried them, for example, the firm did a quick deep dive to examine the meaning. They decided they were worried that something was amiss, and began to take less risky positions. Risk-aware leaders should be extremely cautious in assessing reasons for events. In a world of uncertainty, as we have seen, causality is a very fickle thing.

Perhaps the single most important quality of organizational culture that senior officials must take determined steps to create is a “culture of dissent.” A common way station on the road to risk disasters is punishing alternative views, and sidelining or actively undermining those who raise potential problems with a course of action that seems urgently necessary or rewarding. Senior leaders must exhibit, through word and deed, their commitment to accept, value, and investigate serious, data-based warnings.

In the process, organizations must manage the personality dynamics that can be so destructive of risk management. The lesson of the chapter on personalities is very simple: Risk disasters are as much or more about people and their relationships as it is about any formal procedure. The same thing is true about the other side of the equation—effective risk management is in large measure about getting the right people into positions that have to do with risk. Because the issue is so personalized in its execution, if the right people are not installed in the right places, the best structural mechanisms in the world will not overcome the flaws outlined in the previous chapters. Risk management is ultimately about people and relationships more than processes.

The “right” people would share several characteristics. It goes without saying that they ought to be pragmatic and, for lack of a better phrase, risk-aware. But the leaders on risk in any firm must also be stubborn and

courageous, willing to put their very jobs on the line when powerful, aggressive gamblers in the institution rail against the risk assessments—and when CEOs salivate at the potential for profits, or senior national security leaders longingly consider the opportunities in some dangerous scheme.

The question then—as it will be for all of these habits of effective risk management—is how, specifically, to build these practices into the structure or procedures of an organization. Organizational culture and personality is a delicate quality and can seldom be created with a simple policy or intention. I would suggest three lines of approach to help create the right environment.

Communicating the intent from senior leaders

If an organization is to develop a culture of rigorous thinking and encouraging dissent, the message has to come from the top—and be repeated almost endlessly. It must become a mantra throughout an organization and be reinforced with both personnel decisions and decision-making procedures. But the initial message must be seen as a priority of the most senior leaders, and be reemphasized at every opportunity. The contrary incentives in organizations can develop very quickly, and the tendency will be for staff to think senior leaders want immediate action and short-term results. Creating a sense that the expectation is for rigorous analysis and dissent takes huge effort.

Personnel actions and policies that reinforce the message

Merely saying that an organization demands rigorous analysis and dissent is not enough. Senior leaders and managers throughout the ranks must reinforce the message through a hundred specific policies and decisions. These can include anything from rewarding staff who raise concerns with proposed strategies to making risk-oriented staff positions valued as part of a career trajectory to paying for extended graduate education. People can usually see very quickly what an organization values from what it rewards, and from the jobs and habits that get people promoted to senior levels. An organization concerned to maximize its risk tolerance must therefore engage in a daily campaign to convey analytical seriousness and openness to a range of opinions.

As we have seen, the departments of offices devoted to risk within large financial firms were often kept firmly leashed by leaders anxious to be allowed to take as much risk as they wanted. In firm after firm, from Enron to Countrywide, risk officers were either ignored or actively suppressed. In some cases senior leaders sent powerful messages that the goal of the enterprise was to do deals, and find ways around any risk manager who tried to stop one. The “chief danger from embedding risk managers within the line organization,” Robert Kaplan and Anette Mikes have argued, “is that they ‘go native,’ aligning themselves with the inner circle of the business unit’s leadership team—becoming deal makers rather than deal questioners.”⁵

Another example of such policies is that those with the responsibility for conveying messages of risk and warning must be assigned to those posts in part because they are personally close to the top leaders, and deeply respected by them. There is no more important truth about senior leadership in organizations than the fact that it is very often relational rather than analytical: It is about who the CEO trusts. If the chief risk officer is largely unknown to the CEO and the board, no matter how talented they may be, in the end they are likely to be overruled by leaders who go with their gut, and their trust of some adventuristic manager. The same is true in national security: A president will be hard-pressed to turn down the opportunistic scheme of a favored cabinet member because of the warnings of someone who is essentially a stranger.

A related lesson of the crisis was that many offices or institutions charged with the identification and assessment of risk were far less independent than they seemed. Many organizations that ought to have been playing a strict oversight role instead became enablers of excessive risk-taking, or looked the other way when firms swallowed risky derivatives. This was true both within firms and outside them.

Government and private ratings agencies, from Brooksley Born's Commodity Futures Trading Commission to Standard and Poor's, were viciously criticized when they challenged the risk-taking behavior of major firms. In many cases the companies recruited powerful government officials to weigh in on their behalf; Born, for example, was effectively marginalized during the Clinton administration by a cabal of senior economic policy-makers. The grandfather of all federal economic institutions, the supposedly independent Federal Reserve, was really more of a conspirator in the explosion of risk than an objective monitor, setting interest rates close to zero and broadcasting the message that risk had been conquered.⁶ Conflicts of interest were shot through the whole system, from accountants rendering judgments on firms that were paying them to senior government policy-makers regulating firms they had helped to run for decades.

A clear lesson of the crisis, then, is that it is critical to guarantee the independence of those assessing risk. If they are under the influence of the profit-makers, or if their judgments can be quashed by powerful institutions determined to keep taking risk, they will have little effect. Critical to this goal is a clear message sent from the top of an organization that risk management is a serious business, that the risk function is among the most valued in the firm, and that its opinions will be valued and taken seriously—especially when they conflict with the proposed direction of the firm.

Decision-making procedures and frameworks

Finally, organizations can build specific analytical stages and elements into their standard decision processes. Alone, as we have seen, such procedures will

not manage risk. But combined with the steps outlined above to create the right environment, a set of detailed procedures can be highly effective in disciplining the thought process and decision-making of an organization. At a minimum, as suggested in the previous chapter, an organization should plan to consider outcome-oriented dangers, the things that could go wrong with proposed courses of action. Its procedures should force it to confront specific categories of such risk. In the process, as will be emphasized below, the goal is not to designate a reliable forecast of risk levels (a classic “risk assessment”) as much as to force senior leaders to have the right conversations about key risk-oriented issues.

Beyond the general framework for outcome analysis offered in the previous chapter, another specific tool that can encourage organizations to think this way is an institutionalized “pre-mortem” or a “future regrets” discussion. The idea is to ask the question, “Five years from now we failed in our goals. Why?” The purpose is to get those involved thinking seriously about negative potential outcomes that could undermine the objectives of the organization, and doing so in a very realistic way, “living” the problems as if they had already become real. These and related mechanisms can help organizations create the mental space for rigorous, honest outcome-oriented risk analysis.

Engaged senior leaders

One of the commonest flaws in risk procedures is that risk management is not viewed as a priority—and senior leaders are not sufficiently aware of risk issues—because the function is disconnected from the top leaders in an organization. A critical principle for strong risk functions, then, is to ensure senior policy-maker engagement from the beginning, and to have the CEO and other top officials set a clear tone of risk appreciation from the top.

This problem cropped up throughout the pre-2007 financial industry, and it has characterized many risk and warning failures in national security as well. Too many CEOs of financial institutions were simply unaware of the character of the bets their companies were making before 2007, and therefore the degree of risk they were absorbing. Partly this was a product of some leaders’ ignorance of the complex new instruments; partly it was a result of their simple laziness, or unwillingness to dive deeply into the analysis. In national security cases like Iraq in 2003, once again we find a problem of senior leaders unable to comprehend risk—or understand the seriousness of warnings—in part because they simply declined to familiarize themselves with the details of what they proposed to do.

It is not enough, moreover, for mid-level officials to have this information. Only if the most senior leaders are fully engaged can potential disasters

be headed off. Once again, the principles guiding key decisions flow from the tone set at the top, and that tone must include a familiarity with the details of proposed policies. Senior leaders, of course, make the final call on such policies, and they will not be in a position to judge risk if they have not invested the time to understand the details.

A related problem pre-2007 was the compartmentalization of risk management. Partly with the goal of distributing the risk function throughout the organization, firms ended up creating stovepipes of risk reporting. Specific risk functions or officers would be aware of the dangers inherent in some strategies, but this information did not always flow upwards to the people who needed it—or when it did, it was so heavily censored or counteracted by investment leaders that it had little effect. Robert Kaplan and Anette Mikes studied the risk management procedures of dozens of companies. “Nurturing a close relationship with senior leadership,” they concluded:

will arguably be its most critical task; a company’s ability to weather storms depends very much on how seriously executives take their risk-management function when the sun is shining and no clouds are on the horizon. That was what separated the banks that failed in the financial crisis from those that survived. The failed companies had relegated risk management to a compliance function; their risk managers had limited access to senior management and their boards of directors. Further, executives routinely ignored risk managers’ warnings about highly leveraged and concentrated positions. By contrast, Goldman Sachs and JP Morgan Chase, two firms that weathered the financial crisis well, had strong internal risk-management functions and leadership teams that understood and managed the companies’ multiple risk exposures.⁷

Risk needs to be closely tied into the senior leadership of a firm or else it won’t be taken seriously. There is something of a myth of senior decision-making that tends to minimize detail. Effective leaders ultimately make gut calls, the legend goes, especially on major strategic choices with high degrees of uncertainty. The key to effective senior leadership is getting just *enough* detail, this approach suggests, and filling in the rest of the picture with experience, perspective, and judgment.

There is some truth to this perspective—especially because, as we have seen, deep uncertainty does call for subjective judgment. There would seem to be some tension between the finding here that, for major strategic choices, there will be no objective risk assessment available. But the argument here is not that the fine details of a risk assessment will provide the needed answer. It is that, if senior leaders are unacquainted with the details of their proposed policies, they will be unable to process even subjective arguments about the

potential risks adequately. Well-informed senior leaders are thus a prerequisite for effective discussion and mitigation of strategic risk.

Transparency in the Assessment of Risk

Another barrier to controlling risk before the crisis was the obscurity involved in many of the financial instruments being built, often in intentionally shadowy operations buried deep within companies. Very few outside analysts could grasp what was really going on. This lack of transparency prevented the market from exercising its checking function.

The challenge of transparency is related to, but distinct from, the problem of senior leader engagement. A non-transparent risk process will indeed create a situation in which leaders at all levels will have difficulty being well enough informed to make choices. Senior leaders can remain disengaged even in an environment that slows sharing information; lack of transparency is a structural rather than a leadership issue.

In theory—a theory deeply held by such regulators as Alan Greenspan and Ben Bernanke—the market should have seen through risky ventures before the financial crisis and chopped them down to size well before they had metastasized into economy-threatening perils. The market did not do this for a variety of reasons: the incentives involved, personal relationships that overrode objective analysts, and cognitive biases. But these were all abetted by the simple lack of awareness of what was actually going on. A related phenomenon, which we find at firm after firm, was the hoarding of decisions to a tiny in-group, obsession with loyalty over rigorous analysis, and a refusal to share the basis for corporate risk decisions. All of these together proved, unsurprisingly, a recipe for disaster. Hidden, obscured risk is far more dangerous than openly debated risk, just as a general principle, and effective risk management is partly about the courage to discuss the issues openly.

Lack of transparency can emerge in several forms, and for a number of reasons. One of the commonest on display in the financial crisis was intentional hoarding of information: Specific investment managers or departments wanted to keep secret the degree of risk they were taking, or the nature of their investments. Both within and among firms, data was withheld, obscured, fudged, or otherwise manipulated to keep an accurate picture from emerging.

Lack of transparency, however, can also emerge unintentionally, from the way in which the risk assessments are presented. This is especially common on national security risk presentations. Stoplight charts and other formats display a final risk analysis without any sense of the methodology used to generate it—in particular, the judgments and assumptions that lay behind the colors or numbers. These are exactly the key points on which senior leaders need to have in-depth dialogues, but they are frequently camouflaged by

generic risk assessments that ask the audience to take for granted the process of generating the findings. Sometimes the engine for generating the assessment, whether a framework or process or algorithm, is not clear from the assessment itself.

The result can be a product that embodies significant amounts of false precision, which also obscures transparent risk assessment. Various subjective judgments and significant assumptions can be built into what ultimately appears to be an objective, reliable estimate of very specific risk. The danger is especially great in situations of high uncertainty. Again this is a common challenge with national security risk assessments, which sometimes generate specific estimates in categories or figures that reflect far more arbitrariness than the presentation suggests.

The goal, then, is to convey as much about the elements of the risk assessment as possible—the assumptions made, the mechanism for generating the assessment, and whole analytical chain. Within organizations, information cannot be hoarded. But more than that, the format and content of the risk assessment itself must be as explicit as possible about the basis for its judgments, the specific decisions made in the process. Indeed, as we will see, this is the route to the main purpose of risk processes—generating dialogue among senior leaders about the most critical leverage points that determine risk. Indeed as I will suggest, risk assessments can be most transparent if they dispense with a focus on the final number and emphasize instead the key determinants of risk.

Organizations could institutionalize such transparency in a number of ways. First is to create the right mechanisms to share key risk information to needed offices across an organization. This can be built into normal processes but must also be emphasized as part of the senior leaders' message on risk. There can be regular sessions and mechanisms to ensure that various risk assessments and dialogues are widely shared. Transparency also meets inviting contrary opinions, both openly and anonymously.

Develop a revealing language and framework to drive the right conversations

As suggested above and in prior chapters, a significant problem with existing approaches to risk is that they often end up emphasizing static, simple, seemingly objective determinations of the level of risk. They can be expressed in numbers, degrees on a scale, or color judgments in a stoplight chart. But in any of these approaches, the goal is to convey a forecast of the risk an organization is undertaking, and do so in the most objective means possible.

Yet as suggested in Chapter 3, such processes continue to lack a coherent language to characterize such risk assessments. Terms like “high” and

“moderate” contain too much ambiguity and are not always used consistently from one presentation to another. Various risk documents attempt to create clear definitions, but these are seldom appreciated by the senior leaders who are the target of the briefs.

A version of the human-factors issue also overwhelms the terminology of risk: People tend to hear risk levels according to their own individual understanding of various terms. An infamous example of such problems came in the Bay of Pigs: A Joint Staff assessment of the plan declared that it had a “fair” chance of success, by which the military officers meant “fair” as in “poor”—not very good. Civilian officials, including President Kennedy, took the term in more common parlance: “Fair chance” meaning not too bad. Meanwhile, as researchers such as Daniel Kahneman and Philip Tetlock have discovered, people are routinely poor in attempting to characterize statistical relationships like probabilities, in part because they are interpreting such chances in the light of their own perspective. Even when risk assessments use a seemingly common-sense term like “low,” then, different decision-makers can interpret it in radically distinct ways: Some might see it as equivalent to failure, others as a risk in the 25th or higher percentile.

One potential answer to this problem is that we need a better and more shared taxonomy for risk in the national security realm. This is true enough, but such an initiative would also be partly beside the point: Merely publishing a list of terms will not change the fundamental human factors involved in the subjective interpretation of risk assessments, any more than establishing elaborate procedures for risk management will change the human factors that undermine such processes. A more promising avenue of response is to recognize that the basic assessment of risk is not as important as the discussions produced by the process, which can help senior leaders process dangers in ways that will ultimately mitigate risk. It’s not about telling leaders that risk is “orange” or “high,” in other words: It is about a process that identifies key outcome-oriented risk and exposes senior leaders to the most important variables that will determine it.

In this sense, the most important issue in regard to a taxonomy for risk is to find an approach, including the use of terms, that produces the kind of dialogues that will help mitigate risk. Again, singular judgments such as “high” or “moderate” may not be as useful as a range of less generic terms: True risk management is about senior leaders confronting the most critical dangers in direct and rigorous terms, but all too often risk processes become fixated on generating specific facts or levels. Indeed, this analysis would potentially suggest the radical step of abandoning general risk assessment levels altogether. A revised approach would replace terms like “moderate” or “extreme,” or even percentage or numerical estimates, with terms or short phrases that hint at the actual form of risk that the assessment has discovered. It could then perhaps, as a secondary function, add some judgment as to how severe the risk seems

to be, though this would pose the danger of distracting attention from the new primary assessment.

As an example, suppose the United States is considering a revised composition for its Army, one that creates highly specialized units for several leading forms of conflict. A typical current approach would be to represent under various conflict assumptions in classical levels of risk—"high, low, medium"—or as specific values on some sort of spectrum. Color-coded boxes or circles would be used to simplify the judgment even further.

Instead, such a risk assessment could break down its conclusions into *categories of risk* and then *evidence for risk judgment*. The categories would be a way of signaling the most important issues for consideration, drawn from the broad areas of outcome-oriented risk spelled out in the previous chapter. Each category would then offer a number of sub-categories of evidence for judgment—reasons to believe the particular risk is high or low, likely or unlikely.

In the example of Army composition, a category of risk could be "inability to respond to major regional contingency"—the danger that an Army spread around a number of boutique capabilities might not have a critical mass of heavy forces to take on big fights. (This is a form of sufficiency, of course, demonstrating again that sufficiency considerations can be a form of outcome analysis.) The lines of evidence for judgment could then offer specific facts and arguments helping senior leaders to determine if this risk is high or low. In this example, they could include things like estimated danger of major contingencies in key regions, specific gaps in capabilities, ability to regenerate necessary forces, role of allies in filling the gap, capability of other Joint Force capabilities to do the job, and more.

The goal, again, is to provide senior leaders with two critical sets of information: What *are* the leading potential risks of the proposed course of action? And what evidence do we have to *gauge the severity* of these risks? This approach is designed to make the rationale for risk judgments as transparent as possible, laying out all the reasons why risks might be high or low. It also leaves the final judgment to the senior leaders rather than giving them a pre-cooked risk judgment based on allegedly objective measures. The assessment could list potential mitigation strategies, and offer a very brief assessment of each of them. It could conclude with a provisional summary risk judgment, but one that then becomes an invitation to dialogue about the categories, evidence, and possible mitigation steps. Box 13.1 lays out one category of such a possible assessment. The goal is to help shape necessary conversations by defining the right categories, and then conveying risk as specific ideas of potential outcomes. Risk management processes in support of strategic choice should abandon the goal of conveying risk in particular levels or categories, because these more commonly obscure the needed conversations rather than spark them.

Box 13.1 Revised approach to categorizing risk

Sample issue: Proposal to reconfigure the composition of the United States Army.

The analysis reflects an assessment of the *types of risk* that could be created by the strategic choice, in specific categories; and *evidence for judgment* about the severity of those potential risks. Rough examples are offered below.

Category 1: Inability to meet large regional contingency demands.***Evidence for judgment***

- Landpower demands in potential regional contingencies are significant (Korea and Baltics).
- Proposed Army composition would fall X heavy brigades short of potential demand for single contingency, and Y brigades short of requirement for simultaneous contingencies.
- Risk of war in Korea judged low based on current indicators.
- Risk of war in Europe/Baltics is low/moderate but highly uncertain given potential misperception and unpredictability of Russian actors.
- Best evidence suggests Russian political and military leaders are not decisively affected by perception of gaps in US landpower response capabilities.

Potential mitigation steps

- Air- and maritime-based Joint Force CONOPS to enhance deterrence: *Limited ability to make a decisive difference beyond current approaches.*
- Invest in additional strategic strike assets to counteract landpower shortfalls: *Would demand new resource trade-offs; strike assets can achieve some goals but cannot decisively prevail, and their deterrent effect may be insufficient relative to ground forces in place.*
- Employ more direct form of nuclear deterrence: *Politically difficult, introduces intense new strategic risk.*
- Invest in partner capacities: *Difficult to do more than we are doing.*

Summary judgment: *Gap in capabilities would be real, and difficult to make up with mitigation steps. Role of landpower in decisive outcomes means that gap would have significant effects in case of war. However effect on deterrence is unclear. Risk judged **acceptable**.*

Pursue resilience for when risk management fails

A final principle of effective risk management is to expect it to fail. As we have seen, human factors are likely to undermine even elaborate procedures and processes. Gray swans will be missed. Warnings will be ignored. And at such moments, an organization will need a backup plan for risk management.

The substantive implication is to treat resilience as more important than foresight: Organizations prepared for risk under uncertainty make themselves robust against a wide range of potential dangers to their programs. Another implication is the subject of the following chapter: the need to complement risk management with a more encompassing process to guiding an organization and its strategies through the dangers of its strategic environment. The following chapter will therefore lay out an approach to managing uncertainty, one that makes different assumptions and uses different criteria for success from risk management and discusses the role of resilience in that process. The two approaches can work together: Outcome-oriented risk management can offer critical insights to the potential dangers lurking in an uncertain future. But the over-arching method is the management of uncertainty, a task to which we now turn.

This chapter has argued that the purpose of risk management in national security strategy ought to be narrowly targeted on assessing outcomes of strategic choices. It lays out a framework for doing so—a first draft of possible criteria and categories to employ in order to identify outcome-oriented risk. It then broadens the aperture by suggesting that such outcome-oriented risk assessments are only one part of a broader required approach to complex strategic judgments, an approach I will call “managing uncertainty for strategic advantage.”

14

Managing Uncertainty

If there is one consensus among senior US officials and defense experts today, it is that the strategic environment confronting US defense policy is one of radical uncertainty. It is an environment with a dozen major potential threats, challenges, opportunities, and demand signals, arrayed in a complicated mosaic without clear implications for the size, structure, or technology of a modern military. It is a context of profound ambiguity whose future is less predictable than ever, in which events across the world are connected in denser and more explosive ways.

Uncertainty has become a dominant theme in national security dialogues over the last few years. Former Chairman of the Joint Chiefs Martin Dempsey has argued that earlier in his career, “issues were a little clearer because they were mostly mechanical. For example, the confrontation that we avoided with the Soviet Union was extraordinarily predictable. We had a set of norms that had evolved over time and we both lived by those norms.” Today, on the other hand, “issues are complex. What I mean by that is that even after you unpack them, when you put them back together, the interaction that results will result in ways that are entirely unpredictable.”¹ Retired Army general Gordon Sullivan has written that the Army’s biggest challenge today is “unprecedented levels of uncertainty.”² The 2014 Quadrennial Defense Review summarized the emerging security environment as “uncertain and complicated.”³ The Army’s new “Capstone Complex” is self-consciously about “Operating Under Conditions of Uncertainty and Complexity.”

This sense is not restricted to national security issues. One study based on extensive interviews with senior business leaders concluded that they “have been thrown into a world of uncertainty and ambiguity. Any sense of stability in the present or confidence in their ability to predict reliably the future has disappeared ... A wave of anxiety has been unleashed.”⁴

Despite these recognitions, no well-established framework has emerged for managing uncertainty to parallel the management of risk, and defense

institutions have not evolved coherent processes for this broader requirement. There is a substantial academic literature on uncertainty management, and there have been few efforts to generate policy-relevant resources.⁵ But this has not coalesced into a coherent approach, and there are few conscious institutional efforts to manage uncertainty as a distinct strategic challenge. This chapter makes the case that the explicit management of uncertainty for competitive advantage ought to take center stage in US national security strategy.⁶

Managing risk versus managing uncertainty

One of the challenges with taking uncertainty more seriously is that, like risk, it is seldom defined, and different people mean different things by it. From the standpoint of national security planning, uncertainty is generally used to refer to a future with many possible outcomes, whose course cannot be reliably forecast. Under true uncertainty, the variables are numerous, the interdependencies intense, the lack of an ability to test options under real-world conditions even more obvious.

As argued in Chapter 4, uncertainty stems from at least two sources. It is partly a product of the fact that national security planners have too little reliable information about the present situation, and so they cannot know the full range of trends and variables shaping the future. But it is also, and more fundamentally, about the fact that those trends and variables interact in unpredictable and nonlinear ways to shape the future, and so even perfect knowledge of the present would not solve the problem. The future under uncertainty should be understood as an unfolding reality that emerges through the complex interaction of variables and human agency, rather than a linear progression from the current context. The challenge it poses comes from a combination of missing information and nonlinear dynamics. It is a situation in which future scenarios, threats and opportunities, and demands on an organization cannot be reliably anticipated. Probabilities cannot be judged because there are insufficient predictable or standard relationships in the system.

Despite the tendency of the defense establishment to predict the future, Richard Danzig concludes that “the unpredictability of long-term national security challenges will always confound the irresistible forces that drive prediction,” in part because “the number of variables that influence the national security environment” render accurate forecasting difficult if not impossible.⁷ The fundamental policy challenge is how to shape an institution (or a nation) for competitive safety and success in such an uncertain environment. Risk management can be part of the answer, but a subordinate one.

Three scholars have defined the essence of uncertainty from a strategic perspective in terms of the predictability of the range of scenarios.⁸ They distinguish four levels of strategic environment: a “clear enough” future, dominated by a single, relatively unchanging reality or scenario; “alternate futures,” a

context in which organizations can plan against a relatively small number of discrete and bounded scenarios; “a range of futures,” the potential for a large number of scenarios anywhere along a spectrum; and “true ambiguity,” a situation in which “multiple dimensions of uncertainty interact to create an environment that is virtually impossible to predict.” The current national security environment is closest to their fourth concept, in that it involves a large number of interacting variables that defy categorization even in a specific, defined range of scenarios.

In traditional risk management, an organization works to identify potential risks, to measure them and create frameworks or processes to track them. Managing uncertainty, on the other hand, involves sketching out possible futures relevant to specific options or alternatives, outlining possible risks in each future, assessing the similarities and parallels among futures and risks, and then identifying the capabilities or actions necessary to hedge against the widest array of possible threats. Its essence is building resilience against a range of possible outcomes. As two scholars have argued:⁹

Uncertainty management is not just about managing perceived threats, opportunities and their implications. It is about identifying and managing all the many sources of uncertainty which give rise to and shape our perceptions of threats and opportunities. It implies exploring and understanding the origins of project uncertainty before seeking to manage it, with no preconceptions about what is desirable or undesirable. Key concerns are understanding where and why uncertainty is important in a given project context, and where it is not.

“Replacing ‘risk’ with ‘uncertainty’ as a starting point,” they argue, would have the effect of broadening the perspective of the endeavor.¹⁰ It would encourage a deep and comprehensive dialogue about the many possible futures and uncertain possible outcomes that are related to a particular strategy. Rather than simply trying to identify gaps between means and ends, for example, and quantify them, a process of uncertainty management would ask decision-makers to consider the uncertainties associated with the necessary resources.¹¹ Table 14.1 outlines some of the major differences between the two mechanisms.

More specifically, organizations should seek to manage uncertainty *for competitive advantage*. The goal of such an effort should be to produce a lasting theory of competitive advantage *vis-à-vis* specific other actors in the environment. These strategies can use cooperative ventures—alliances and partnerships—to achieve this goal; the presumption need not be that the environment is universally hostile. The critical distinction is between a passive and defensive process of managing uncertainty (or risk), and an active and opportunistic one. Effective strategy-making under uncertainty is as much

Table 14.1 Managing risk vs. managing uncertainty

Risk management	Managing uncertainty
<ul style="list-style-type: none"> • Identify discrete risks • Measure and quantify risk factors • Comprehend and intervene in causal dynamics • Produce deterministic risk assessments • Predict and forecast likely outcomes • Develop mitigation strategies 	<ul style="list-style-type: none"> • Build complex picture of environment • Develop scenarios of possible futures • Understand potential sources of uncertainty—unknowns, nonlinear dynamics, and assumptions • Avoid predictive models • Focus on principles of robustness against many potential futures

about enhancing competitive advantage as it is about avoiding risk, in part because risk is a function of that position.

Toward a concept for managing uncertainty

The US national security enterprise needs principles to keep an organization (or nation) safe as it passes through a highly uncertain future, and to do so with constant attention to relative advantage in a complex environment.¹² The goal is not merely to manage uncertainty but to do so in ways that leaves the nation relatively stronger and more secure.

The concepts of *robustness* and *resilience*—and the distinction between them—play important roles in understanding mechanisms to manage uncertainty. A strategy for managing uncertainty should aim at a robust posture, one that, as the synonyms of that term suggest, is stout and tough, a system that can sustain damage and keep functioning. It achieves this result in part through its resilience—its flexibility, pliability, and capacity to adapt itself to changing demands. A robust energy network, as an example, is one that is physically sturdy, whose information systems are armored against hacking, and which has redundancies to allow continued operations after some damage has been sustained. It is resilient to such damage in the degree to which it can heal itself through flexibility—using various sources to bring energy to bear in distinct channels and ways. There is obviously some overlap between the two concepts; qualities that make a system robust can also contribute to resilience.

A process for promoting competitive advantage under uncertainty need not abandon all effort to assess and anticipate dangers. In fact, managing uncertainty recognizes that the *diagnosis of the situation* remains the foundation for all judgment. Rigorous analysis of existing realities, relationships, and other efforts to understand the strategic context does not disappear in uncertainty management; it merely takes on a more constrained role.¹³ But the goal should not be to develop a singular portrait of the “correct” current reality or likely future as much as to build a comprehensive mosaic of the potential realities.

The most important element of the necessary mindset is the concept of *alternative perspectives or lenses*. The single most critical requirement for judgment under uncertainty is the willingness and commitment to see everything—present, past, and future—through a range of different perspectives. The most dangerous habit under uncertainty is adopting a limited, narrow view that adopts a single explanation for events. Yet this is the default approach, both in terms of individual judgment and organizational planning. It is critical that the diagnosis set the right mindset for the analysis and strategy-making process under uncertainty, and that mindset is one that actively considers alternative ways of viewing a given situation.¹⁴ Managing uncertainty is first and most fundamentally about building alternative, competing lenses with which to view and analyze both present and future contexts.

One example is a technique of intelligence analysis called the assessment of “alternative hypotheses.”¹⁵ Decision-makers should make efforts to understand various interpretations of the current context, as well as a range of alternative futures that could emerge from it. Today’s growing Russian assertiveness, for example, could be the result of a number of factors; whether the behavior stems from defensiveness or aggression, whether Russia views itself as under siege or confidently claiming a sphere of influence, could generate different explanations for events. The emphasis should be on imagining the range of plausible futures that could emerge, in order to empower the creation of postures that offer the greatest robustness against that range. It is an open-ended probing of the future rather than a deterministic effort to identify specific, likely outcomes.¹⁶

With a vision of alternative explanations for the present context and a range of future possibilities in mind, a process for managing uncertainty could then assess alternatives for a defense or national security postures according to a number of criteria for robustness and resilience. These criteria would be designed to provide a sense of how well the strategy can sustain itself against the range of possible dangers. The assessment should be made against *types* of dangers in addition to, or even instead of, specific scenarios: Limiting the analysis to a handful of expected contingencies would undermine the whole purpose of managing uncertainty, which is to build resilience against a range of possible outcomes. By these criteria, risk becomes a function of the degree of ability to respond and adapt to classes of dangers in the future environment.

This approach has much in common with the “hedging” strategy recommended in a recent essay by Ben Fitzgerald and Scott Cheney-Peters.¹⁷ They contrast such a hedging approach to recently proposed “offset” strategies, which attempt to identify and target a specific competitive advantage on the part of a particular adversary. A hedging strategy, by contrast, “would emphasize an ‘eggs in many baskets’ approach, with a spread of investments over a portfolio of concepts and capabilities useful against a range of threats.” It

would invest in institutional adaptiveness, tailorable force packages, and capabilities that embody modularity and variety.¹⁸ The goal of the process would be to determine ways for a defense and national security policy to reflect goals such as adaptability, innovation, and institutional health, to create an overall posture that is robust against a range of alternative outcomes.

Nassim Nicholas Taleb prefers the concept of “antifragile” systems to terms or ideas of responsiveness or resilience. By antifragile he means “things [that] benefit from shocks,” which “thrive and grow when exposed to volatility, randomness, disorder, and stressors.” Resilient systems stand up against disorder; antifragile ones grow stronger by interacting with shocks. This is an important concept and I will try to build some of his resulting principles into my own framework below. He worries that concepts of resilience are “timid” and reactive rather than capable of sustaining institutions that truly thrive under uncertainty. But Taleb’s basic goal is the same: to “build a systematic and broad guide to *nonpredictive* decision making under uncertainty,” in situations in which “the unknown preponderates, any situation in which there is randomness, unpredictability, opacity, or incomplete understanding of things.”¹⁹

Today the policy focus is inverted. The emphasis, in procurement, force structure, and other decisions, is on building capabilities against a range of specified threats, and measuring risk in terms of the gap between those sets of means and ends. That would not be abandoned in an adaptation strategy, but it would be complemented and in some ways superseded by a planning architecture and process that prioritizes adaptability and resilience. Part of the goal would be to make the alternatives between the two approaches, which now often remain hidden, more explicit. Choosing more current capacity against immediate threats in order to under-invest in institutional health and innovation for the future may make sense—but today, that particular decision is almost preordained by the process, and it is seldom debated openly.

One important rule in managing uncertainty is to rely on a number of simple rules of thumb. Bank of England economist Andrew Haldane has argued that one lesson of the crisis is that complex situations call for simple practices. Over-complexity is a source of failure under nonlinear contexts. In a situation of probabilistic risk, “policy should respond to every rain-drop; it is fine-tuned. Under uncertainty, that logic is reversed. Complex environments instead call for simple decision rules. That is because these rules are more robust to ignorance. Under uncertainty, policy may only respond to every thunderstorm; it is coarse-tuned.”²⁰

Elements of managing uncertainty

Managing uncertainty is much more of an ecological process to manage complexity than an analogue to an engineering effort to structure a building. Senior officials should think of themselves as an actor in an adaptive network,

building robustness against potential dangers and developing sources of competitive advantage in the overall environment.²¹ In developing an overall strategy focused on managing uncertainty and dealing with complex and nonlinear dynamics, the leaders of the US national security enterprise could pursue a number of specific principles or objectives (see Box 14.1).

Box 14.1 Principles for managing uncertainty

- **Diversity:** Cultivate a wide range of capabilities to hedge against many potential futures.
- **Redundancy:** Build multiple capabilities able to tackle the same challenge.
- **Modularity:** Create a robust, redundant, and self-reliant system in which the loss of individual nodes will not doom the whole network.
- **Stabilizing feedback loops:** Create institutions, norms, procedures, and mechanisms that tend to dampen extreme swings and return systems to stable centers.
- **Leverage points:** Find the capabilities, principles, or actions that generate nonlinear positive effects on the system.
- **Innovation and experimentation:** Encourage constant creativity and new thinking to enhance adaptation to changing circumstances. Cultivate a degree of randomness and accident that can generate positive unplanned advances.
- **Loose control procedures:** Avoid top-down inflexibility. Govern through broad direction; accept and allow variation in implementation. Tightly-managed homogeneity is a weakness.
- **Emergent and responsive approach:** Assume that strategy will unfold through constant probing of environment. Prefer bottom-up to top-down solutions.
- **Versatile portfolio of options:** Build a range of options along a spectrum of risk (from no-brainers to general options to bold moves) to have available for multiple contingencies. Seek versatility—flexible responsiveness to many futures—in addition to robustness and resilience.
- **Patience for a series of gradual yet rapid adaptations:** Recognize that under uncertainty, responding to challenges too urgently can be reckless; patient incrementalism is the default approach—along with a decisive response once a challenge becomes clear.
- **Anticipatory shaping:** Shape the strategic environment to dampen organic aspects of uncertainty.
- **Resilience:** Pursue specific policies to increase ability to withstand shocks.

Taleb outlines the characteristics that make organizations vulnerable under uncertainty—the sources of weakness that must be mitigated by a strategy for managing uncertainty. “The large, optimized, overreliant on technology, overreliant on the so-called scientific method instead of age-tested heuristics”—such institutions are inherently vulnerable, he contends.²² Brittleness is a function of size, linear planning mindsets, high degrees of interconnections, and efforts to take uncertainty through objective models. Strength comes from multiplicity, flexibility, and a range of characteristics that make organizations supple and responsive to change. The goal is to become as robust as possible against the widest range of potential plausible futures. Each of these principles contributes in its own way to an overall concept of uncertainty management.

Diversity

The concept of generating a diverse portfolio of national security tools and strategies is that a wide range of capabilities will provide a hedge against many potential futures, threats, and missions. Because a single future cannot be forecast or engineered, actors must hedge against uncertainty with a range of capabilities. Efficiency can be the enemy of robustness. This suggests, for example, generating a balanced Joint force with a wide range of capabilities and different, and sometimes overlapping, means of achieving the same missions. In terms of equipment and procurement, this principle warns against relying on a single system or technology to achieve a mission or goal—a single communications platform, aircraft, or missile. Resilience comes from diversity; the more an organization relies on narrow technologies, the more vulnerable it becomes.²³

The principle of diversity also suggests not focusing myopically on any one threat. It recommends developing a broad set of non-military instruments of statecraft. In regard to specific challenges such as terrorism or climate change, this principle suggests building a number of distinct and overlapping approaches rather than placing all eggs into one basket. It also implies a mindset and continuous practice of experimentation, of constantly developing new potential tools and approaches to add to the mosaic of capabilities.

Redundancy

A second principle calls on organizations to reduce their obsession with efficiency, and accept the value of having partly redundant capabilities.²⁴ Unforeseeable events can threaten particular capabilities—anything from a weapons system to an economic tool to a specific technology like a radio. To avoid single-point vulnerabilities, institutions should carefully assess demands and requirements and build an overlapping portfolio of capabilities.

Modularity

In situations of uncertainty and many potential threats, national security strategy should aim to ensure that the loss of one node does not automatically cascade into systemic failure. The more each element of a network is secure, the more the system as a whole is robust against contagion. This principle holds that robustness in the system comes from the strength of individual units as much or more than system-wide guarantees. It suggests, for example, that the military must be very careful about investing in networks or systems that are so heavily interlinked that vulnerability in one point automatically creates vulnerability for the system as a whole. Information technology is an obvious example, but others might include weapons systems that rely on a single sensor or radar, or a battlefield management system without which specific units could not adequately fight. The “network of networks” concept is an invitation to vulnerability. This same principle applies to energy technology: The existing energy grid is both efficient and vulnerable because of its linkages; encouraging a much greater use of decentralized, independent generating systems, such as solar, localized wind, and fuel cell technology, could render the overall system much more robust.

These first three principles point to a specific area of needed reform: significantly reforming the way the Department of Defense builds requirements for, and acquires, major weapons systems. Managing uncertainty demands more of the reforms that have been proposed for the procurement system for decades: more rapid prototyping, more use of off-the-shelf technology, multiple lines of development, making incremental moves or true leaps ahead but not vastly expensive half-measures, and so forth.²⁵

They also produce a related insight about the socioeconomic foundations for responsiveness to an uncertain environment. Few of the policies advocated here will be effective over the long-term if they are forced to draw on an industrial base that is lagging behind world standards, or a society consumed with inequality and partisan tension. Just about every major strategic recalibration in modern US history—Eisenhower’s narrowing of the global ambitions of early Cold War US strategy, the Nixon administration’s effort to balance means and ends, the Obama administration’s recovery from global belligerence and overextension—has come with reminders that security demands attention to the home front. The economic and social foundations of US strength demand investment as much as any particular military capability; over the long-term, managing uncertainty will depend more on such general strengths than on any one technology or type of force.

Engineer stabilizing feedbacks

One of the primary sources of danger in an uncertain environment, as we have seen, is cascading or contagious risks. If a dangerous outcome of a

strategy is limited and constrained, it will not threaten disaster. It is when risks rush through whole systems that major crises emerge. An important strategy for managing uncertainty is therefore to find ways, if possible, to build inherent dampening effects into the system. This is the definition of a stable system—one that tends to return to a normal distribution.

In part this calls for a renewed agenda to bolster the international norms, rules, and institutions that help dampen and constrain the results of uncertainty. Stronger international norms enhance predictability and, as in the British case, allow the leader of the system to share the burden of stabilizing it. Conflict-management norms and institutions reduce the incentive for military solutions to problems. Today, as a number of commentators have recognized, the post-war liberal order is under significant strain. Investing in its health—through enforcement of norms, financial and political commitment to institutions, sharing decision-making authority with a wider array of countries to enhance the system's legitimacy, and more—should be a major focus of a strategy for managing uncertainty in the years ahead.

An example of such an initiative would be efforts to ensure that large-scale aggression is self-defeating and generates negative rather than positive feedback loops. In responding to Russian belligerence in Eastern Europe, for example, the United States should seek—as it has been doing—to reinforce the natural systemic reactions to aggressive overreaching. It is critical that the United States use such examples to set precedents about the feedbacks to be expected from aggression. It may be that the symbolic reaction—unanimity in the condemnations of the actions by governments in Europe and beyond—are as important as the concrete sanctions imposed. What is important is that the United States leads an effort to generate stabilizing feedback effects.

Discover leverage points

In an uncertain system exhibiting nonlinear complexity, inputs will not always be proportional to outputs. This can be true of risks—but it can also be true of US actions. Part of the goal of a strategy for managing uncertainty is to discover those initiatives that promise a disproportionate beneficial outcome. As a recent National Academy of Sciences study put it, “nonlinear behavior ... means that an effective remedy need not require a massive effort, just a well-targeted one.”²⁶ A conscious strategic focus in confronting uncertainty, especially with constrained resources, should be to discover the actions or initiatives with special leverage. An example might be a specific new military technology that would negate or sidestep whole classes of competitors' systems, or a diplomatic initiative that would change the playing field in a dramatic way.

In this context, while generally remaining patient and adapting gradually to an uncertain environment, the United States should also be constantly on the lookout for home-run opportunities that promise to tilt the strategic balance in its favor. These must be approached with intense care, and in particular,

outcome-oriented risk analysis must be comprehensive when considering them. All too often—as with the Iraq disaster—national leaders convince themselves that a wild swing of the bat will transform the strategic context, when in fact they have merely invited a calamity. And yet from time to time, bold steps—investments in particular technologies, new alliances, strategic recalibrations—do change the overall balance in ways that make US national security postures more sustainable and effective.

Innovation and experimentation

In order to generate a persistent stream of new capabilities to adapt to changing circumstances and suddenly rising threats, the national security architecture must be deeply innovative and constantly experimenting with a range of new approaches and systems. This principle is important at many levels. Specific military services, for example, must embody a highly creative mindset, constantly experimenting with new capabilities, concepts, and force structures. The principle recommends such policies as rapid prototyping and multiple lines of investment to provide the widest range of options for choice and creative space.

This principle also demands the new approaches to procurement and production suggested above, policies that emphasize faster responsiveness and greater experimentation. Richard Danzig emphasizes that a context of uncertainty creates a repeated requirement to “meet unanticipated needs.” He recommends increased attention to such concepts as open systems architecture, modular design, and spiral development. “Manufacturing adaptiveness” is a critical characteristic for future defense policy.²⁷

Another aspect of an innovative and experimental mindset is to cultivate a certain degree of randomness and variation in policies and capability development.²⁸ Organizations cannot adequately respond to uncertain and unpredictable contexts with planned and structured processes based on an anticipation of future events. In the management and innovation processes, in the corporate culture, even in the personality of some key leaders, organizations should cultivate a degree of unpredictability and willingness to try bold and untested new ideas. Through the somewhat accidental outcomes that result, the organization will generate potential capabilities, ideas, concepts, and more that could potentially be of use in responding to unpredictable future challenges.

Loose not tight control

In order to cultivate the agility and creativity required by an uncertain context, organizations should avoid attempts to impose rigid, top-down control of organizations.²⁹ Such control presumes that objectively “correct” answers can be determined and mandated onto organizations. That may be true in certain limited areas or on specific issues. For the most part, however, organizations responding to uncertainty will want to maintain an approach that

devolves significant initiative and decision-making authority through the ranks. Only this will provide the sort of grassroots responsiveness necessary for such environments.

An emergent and responsive approach

Managing uncertainty also demands an emergent, bottom-up approach to strategy. James G. March has drawn a distinction between rationalistic and adaptive risk management:³⁰ The degree of uncertainty involved in complex strategic judgments means that no process could address all the relevant information in a single phase and adequately assess the risks involved. A better alternative is to pursue what March and others have referred to as *adaptive strategies*: the gradual, repeated treatment of risk allowing an institution to adapt to an emerging situation rather than impose one meaning on it at one time. Adaptation, March and others have argued, is a more appropriate strategy for emergent uncertainty.

The lesson for national security strategists is to realize that the process of managing uncertainty is an unfolding task, one that must be built around continual discussions supported by rigorous information searches and analyses of the elements of strategy. When “managing risk” for a specific issue—such as the credibility of US promises in a situation of declining US preeminence—senior officials should conceive of the overall process as an ongoing adaptive process of managing uncertainty. Within that, a specific effort to assess the risks from the consequences of proposed strategies can play an important role.

Henry Mintzberg famously described an ecological approach to strategy. He contrasted “planning” strategy—a carefully laid out and deliberate process designed to generate rational control of the situation, based on systematic analysis—with “crafting” strategy. He described this as “not so much thinking and reason as involvement, a feeling of intimacy and harmony with the materials at hand, developed through long experience and commitment.”³¹ Effective strategies can emerge as well as be designed, arising in response to an unfolding situation. This implies a grass-roots approach of gradual discovery: “The strategist *finds* strategies no less than creates them, often in patterns that form inadvertently in its own behavior.”³²

Versatile portfolio of options

Versatility, Frank Hoffman has written, is “based on a breadth of competencies, instead of a collection of specialized organizations or players.”³³ An essential principle of effective strategy under uncertainty is to build a broad-based set of capabilities that can respond to a range of demands. At the same time the principle would recommend prioritizing those capabilities that are most adaptable, good for more than one contingency or issue.³⁴

Patience for a series of gradual yet rapid adaptations

Under uncertainty, an organization's first impulse of what to do in response to pressures and challenges will often be wrong. As Nassim Nicholas Taleb has stressed, what organizations do *not* do when confronting uncertainty is as important as what actions they rush to take.³⁵ Organizations should take a fundamentally incremental, patient approach under uncertainty, because they will not be able to know the future well enough, typically, to make grand bets.

A leading principle of managing uncertainty can be phrased simply as, "Avoid fool's errands." Given the wide range of potential demands, the United States should avoid unnecessarily throwing resources and credibility into hopeless causes that reflect secondary interests. There is some tension here, of course, because the United States often does those things to shore up a perceived credibility gap—credibility that is essential for the maintenance of the system as a whole. Two major lessons of the last half-century, though, are that national credibility is not a fragile flower, capable of being ruined with a single wrong step—and that, in any case, fool's errands rarely end up sustaining credibility.

Restraint when facing such temptations can emerge from the recognition that, as urgent as crises may look in the moment, the international system is often self-healing. Problems can sometimes solve themselves, and in a nonlinear environment actions cannot be guaranteed to have their intended effects in any case. The one paradoxical aspect of this requirement is that, once an immediate challenge becomes apparent for which a response *is* necessary, an organization must be able to act quickly and decisively. This suggests the need for versatile and responsive organizational processes even under the general approach of strategic patience.

Anticipatory shaping

Although the potential for accurate forecasting under uncertainty is limited, a strategy for managing uncertainty ought to seek to mold the strategic context to the degree possible. One way of mitigating uncertainty is to channel events in hopeful directions. Organizations should determine the future environment they want, not merely the one that seems to be emerging, and take actions to shape events in that direction.

Resilience

Finally, thinking in terms of managing uncertainty points us in the direction of building resilience rather than managing a quantified set of risks. The transition that must be made is from reliable foresight to resilience: Instead of identifying and anticipating specific risks, decision-makers should pursue more emphasis on understanding broad categories and then building general systemic resilience. It is a shift from risk management to uncertainty

tolerance, from predicting and avoiding dangers to hardening an institution against likely surprises. Joshua Cooper Ramo emphasizes the importance of systems “capable of managing the demands of a changing environment” through strong resilience.³⁶

A fundamental truth about managing uncertainty is that *resilience is a better strategy than anticipation*. Mary Douglas and Aaron Wildavsky discuss the importance of developing an adaptive organization that can constantly respond to shifts in the risk profile. Attempting linear control and anticipation will fail: “anticipation to secure stability is a bad bet for safety,” they write. “Without continuous experience in overcoming a variety of disturbances, organisms are likely to adapt to a steady state. When dramatic change does occur, these organisms are more likely to perish both because the search for stability has used up surplus resources and because they have suppressed their capacity to cope with the unexpected. ... Resilience is the capacity to use change to better cope with the unknown; it is learning to bounce back.”³⁷ Resilience is the fundamental response to uncertainty: “Since we do not know what risks we incur, our responsibility is to create resilience in our institutions.”³⁸

The objective of resilience can be served by many of the principles and concepts outlined above.³⁹ A key component of resilience, for example, is variety: Rather than betting on a single technology or approach, organizations should develop a range—for example, in energy policy. This is again the principle of diversity. Modularity is also integrally related to resilience: The more secure each node in the system can be made, the more resilient it will be.

The principles of a responsive institutional culture outlined in the previous chapter are essential for the success of this parallel set of effective strategy under uncertainty. Only if an organization can accept dissent and promote open-minded analysis of issues will any of these mechanisms for managing uncertainty work as intended. The question then is what institutional structures can be developed to create default processes for managing uncertainty.

Institutionalizing the management of uncertainty

In order for the management of uncertainty to become a habit in the national security enterprise, it will not be enough merely to outline promising steps and expect existing policy processes to respond. Different habits are too deeply ingrained in policy, planning, and risk management mechanisms today. Adaptive planning for uncertainty will have to be institutionalized and made a formal component of national security planning. In order to do this within the Department of Defense, a number of reforms are possible.

The Department could, for example, create a formal process to parallel the current requirements and capabilities-based approach with one focused on adaptive criteria for resilience and robustness. It could be guided by a

management council under the guidance perhaps of the Deputy Secretary, with a new senior position created to oversee specific analytical areas or programs—for example, an Assistant Secretary for Adaptive Strategies. The core staff underneath this office could be very lean, just a few dozen people; its major function would be coordination and sponsoring rather than overseeing large programs. (If existing offices with more programmatic responsibility were transferred to this effort, as suggested below, the number of staff would grow, but these positions would not reflect a net increase from current defense staffing.) Under the assistant secretary, various offices would be responsible for ensuring regular attention to different aspects of adaptive strategies. This new office could be responsible for several categories of policy and development.

1. *Open-minded and alternative analyses.* From an analytical standpoint, the foundation for any effort to manage uncertainty is a deep appreciation for alternative interpretations of the present situation, and a continuous effort to outline and build resilience and robustness against a range of future scenarios. Any institutional home for managing uncertainty should have a dedicated analytical function, one designed to bring together thinking from around government into ongoing and rigorous assessments of present and future realities. This function could reside under a unique office director or deputy assistant secretary and could include:
 - o A small, dedicated staff of analysts devoted to alternative analyses of present realities and future scenarios.
 - o A significant contracting vehicle to conduct scenario planning efforts. The office would be responsible for coordinating among the dozens of such exercises that take place across government—to share themes, create efficiencies, assess the analytical soundness of game outcomes and lessons, and generate common conclusions.
 - o A function to link offices and individuals throughout the intelligence community devoted to red-cell and alternative analyses. It could also use emerging approaches to crowdsourced forecasting and analysis to enhance such thinking.
2. *Sustainability analyses and policies.* The office could also promote resilience and robustness by ensuring that the US defense posture remained sustainable. A significant risk facing any institution is that trends in the competitive environment might render its advantages or fundamental strategies untenable. Several such risks arguably confront US defense planners today, both in terms of the strategies and initiatives of potential adversaries and unsustainable US policies. This function could address a number of aspects of the sustainability question:
 - o Sponsor research and analysis to understand the leading threats to sustainability.

- o Understand specific risks to technological advantages and areas especially vulnerable to being overtaken, and develop approaches for mitigating these risks and preserving a sustainable technological edge in key areas.
 - o Develop policies to preserve the sustainability of defense institutions relative to self-imposed risks (for example, changes to pay and benefits to avoid bow-wave costs).
3. *Innovation and experimentation.* This area of emphasis would aim to enhance the innovative character of the overall defense enterprise, and to sponsor the sort of broad experimentation necessary to keep an institution resilient against a wide range of threats. One question would be whether the full scope of defense research and development efforts would come under this adaptation office. Some of these efforts are closely tied to service capability requirements, and it might be difficult to break them out under the supervision of an independent agency. Perhaps a way could be found to preserve basic RDT&E on a set of core capabilities under the purview of existing service and DoD processes while segmenting out a handful of long-range, experimental, and prototyping activities under the new adaptation office. Specific actions under the purview of the organization could include:
- o Supporting specific, small-scale experimental technologies through prototyping and experimentation⁴⁰—as Fitzgerald and Cheney-Peters put it, “more ‘hackathons,’ ‘X-prize’ type challenges and crowdsourcing.”⁴¹
 - o Organizing a coherent long-term research program. Perhaps the Defense Advanced Research Project Agency could be placed under the responsibility of this office and be used as the centerpiece of such funding.
 - o Working with scholars and practitioners to develop the habits, practices, and structures of innovative cultures and helping to integrate them through DoD.
4. *Building a diverse array of force capabilities, structures, and technologies.* A major danger today is that the United States risks creating what Fitzgerald and Cheney-Peters call a “capability monoculture” from its massive investments in a few dominant systems.⁴² They point out that the F-35 is scheduled to consume a quarter of the entire DoD procurement budget. It is also designed to replace a suite of systems, creating some degree of the sort of single-point failure risk that resilient systems strive to avoid. Another function in this office could be responsible for developing criteria, plans, and programs to avoid such vulnerabilities and to ensure that the United States develops a broad spectrum of capabilities for future scenarios. This could entail a number of possible initiatives:
- o Develop a capability-to-task matrix to understand the ways in which various service capabilities serve specific requirements or missions—and uncovering areas with a dangerous lack of redundancy or single-point vulnerability. The office could conduct, for the first time, true cost-effectiveness comparisons among not just specific systems or force

- structures to provide needed capabilities, but also among sets or combinations of such capabilities designed to both address the need and provide diversity and redundancy.
- o Conduct assessments of vulnerabilities to determine areas of greatest risk from the standpoint of redundancy. This would include efforts to determine systemic risks that undermine redundancy, such as information technology threats that have the potential to affect a wide range of capabilities simultaneously.
 - o Tie these assessments together with technology experimentation to build resilience through research and development, targeting areas of current technology monoculture or future vulnerability.
5. *Innovate in concepts and doctrines.* Finally, the new function could serve as a catalyst for innovation in Joint approaches to operations. While individual services attend to service-level doctrine, this office could work to build innovative ways for the Joint force as a whole to employ its capabilities to achieve ends. Such thinking will be especially important at a time when a diverging ends-means gap calls for new ways—strategic and operational concepts—to promote US interests with fewer resources and less influence. Such a function could involve several initiatives, including:
- o Identifying issues and problems that demand new conceptual treatment, such as light-footprint irregular warfare and new approaches to regional deterrence under the shadow of area denial capabilities.
 - o Assessing areas of fit and conflict between and among service-specific concepts and doctrines.
 - o Building new Joint concepts and doctrines.

The Department of Defense has now acknowledged that the emerging strategic environment is characterized most of all by uncertainty. But it has yet to develop mechanisms to actually plan for uncertainty, and consciously employ national security activities and initiatives to build robustness and resilience for an uncertain future. Doing this requires a concept of how to plan for uncertainty as well as reforms to our processes and organizations to reflect the revised emphasis. Safety in the face of a rapidly changing, nonlinear future will not happen by itself—it will demand conscious effort. The time has come for a more serious commitment to the principles, processes, and institutional functions that will build adaptability into US national security institutions. Taken together, these components of a new defense planning process would formalize and institutionalize critical mechanisms of planning for uncertainty.

Taken together, the last two chapters have outlined a combined approach to the challenge of building a process for managing uncertainty. That approach begins by building the aspects of organizational culture necessary to create the right mindsets and habits, and then pursuing a number of principles of organizational strength and health under uncertainty. Combined with the aspects

of modified risk management suggested earlier, these concepts add up to a broad approach for managing risk and uncertainty in complex and ambiguous environments.

This chapter suggests three broad sets of recommendations. First is to set risk in its proper place in the strategy process, and address it in the most effective way. This study suggests taking many current approaches to risk and treating them as either threats or sufficiency analyses, and restricting the term and its associated processes to analyzing potential negative outcomes of proposed strategies. In the process, the study recommends an end to quantitative or threshold-based presentations of risk—"0.7 on a scale of 1 to 10" or "moderate" or "yellow"—and instead focus on a description of the types or categories of potential risk and the specific analytical conclusions that determine the judgment.

Second, the previous chapters then suggest setting the institutional culture to ensure an adequate consideration of risk. This means, in other words, producing a culture and series of habits and procedures that can take warnings seriously without being paralyzed by them. Chapter 13 outlined the basic principles necessary for such an outcome, including an analytically rigorous culture, the full engagement of senior leaders, and transparency in the identification and evaluation of risk.

Finally, this chapter has discussed a broader requirement than mere risk management—managing uncertainty for competitive advantage. It has laid out a number of specific principles for achieving that goal and a few recommendations for the character of US national security policy.

Taken together, these various recommendations point toward a somewhat revised mindset in terms of planning for national security. Rather than merely planning against defined contingencies and assessing risk in terms of threats and ends-means relationships, this new approach would take seriously the character of uncertainty and the tendency of organizations to miss outcome-oriented risk. It would place major emphasis on generating an agile, versatile, and adaptable national security posture, and on an unwavering commitment to deep understanding of the current situation and potential futures. Many senior leaders of the US defense establishment have argued that we confront an era of unprecedented uncertainty and potential for nonlinear surprises. This study agrees, and has argued that it is time to take that demand seriously from the standpoint of national security planning.

Notes

1 Risk, Judgment, and Uncertainty

1. Quoted in Scott Patterson, *The Quants: How a New Breed of Math Whizzes Conquered Wall Street and Nearly Destroyed It* (New York: Crown Business, 2010), 263–264.
2. Quoted in William D. Cohan, *House of Cards: A Tale of Hubris and Wretched Excess on Wall Street* (New York: Doubleday, 2009), 9–10.
3. David Wighton, “Imagination and Common Sense Brew a Safer Culture,” *FT.com*, November 28, 2007.
4. Anil K. Kashyap, “Lessons of the Financial Crisis for Risk Management,” Paper Prepared for the Financial Crisis Inquiry Commission February 27, 2010, 15; available at http://faculty.chicagobooth.edu/anil.kashyap/research/papers/lesson_for_fcic.pdf.
5. One “Compendium” which cites many of the DoD risk processes in place in the early 2000s can be found at <https://www.aiha.org/get-involved/VolunteerGroups/Documents/RISKVG-%28Attachment%20F%29%20Compendium%20of%20risk%20assessment-risk%20management%20resources.pdf>. See, in particular, Department of Homeland Security, “Risk Management Fundamentals,” Washington, DC: DHS, April 2011; Defense Systems Management College, *Risk Management Guide for DoD Acquisition* (Ft. Belvoir, VA: Defense Systems Management College Press, May 1999); United States Army Field Manual 100-14, “Risk Management,” April 23, 1998.
6. Australia has a particularly extensive approach, integrating business concepts of Integrated Risk Management into its defense planning process. See, for example, Svetoslav Gaidow and Seng Boey, “Australian Defence Risk Management Framework: A Comparative Study,” Canberra, Australian Government, Department of Defence, Systems Sciences Laboratory, 2005.
7. See, for example, Yasmin Tadjdeh, “TRADOC Leader: Uncertain Future Biggest Threat to Army,” *National Defense*, June 30, 2015; available at <http://www.nationaldefensemagazine.org/blog/Lists/Posts/Post.aspx?ID=1887>.
8. Quoted in Tara Copp, “Gen. Dempsey: This Is Not a Time to Be Withdrawing from the World,” *Washington Examiner*, May 6, 2015.
9. William J. Perry and John P. Abizaid, co-chairs, *Ensuring a Strong U.S. Defense for the Future: The National Defense Panel Review of the 2014 Quadrennial Defense Review* (Washington, DC: U.S. Institute of Peace, July 31, 2014), 58.
10. For general treatments, see Peter Bernstein, *Against the Gods: The Remarkable Story of Risk* (New York: Wiley, 1998); Thomas S. Coleman, *A Practical Guide to Measuring Risk* (New York: Research Foundation of CFA Institute, 2011); and Michael Crouhy, Dan Galai, and Robert Mark, *The Essentials of Risk Management* (New York: McGraw Hill, 2005).
11. J. F. Yates and E. R. Stone, “The Risk Construct,” in J. F. Yates, ed., *Risk-Taking Behavior* (New York: Wiley, 1992), 1.
12. John Adams, “Egg on Wall? It’s Rational,” *Times Higher Education Supplement*, October 6, 2000.
13. Jack Dowie, “Against Risk,” *Risk Decision and Policy*, Vol. 4, No. 1 (1999). Dowie’s analysis diverges from my perspective, however, in returning to an emphasis on quantified answers—risk should give way, he suggests, to a rigid effort to discover accurate measures of probability and consequence.

14. Admittedly, the aspects of strategic-level choices and uncertainty are two different concepts, and need not go together. My argument, however, is that they overlap very significantly: Complex strategic judgments *necessarily* imply a context of uncertainty, because of their nature.
15. James H. Barnes, Jr., "Cognitive Biases and Their Impact on Strategic Planning," *Strategic Management Journal*, Vol. 5, No. 2 (April–June 1984), 134.
16. See, for example, Simon Johnson and James Kwak, "Seduced by a Model," *New York Times*, October 1, 2009.
17. Phil Rosenzweig makes the argument for the importance of differentiating decision type—and sketches out a strategic-level category of decision close to what I have in mind here—in "What Makes Strategic Decisions Different," *Harvard Business Review*, November 2013, 89–90.
18. This same fact offers strong reason to distrust classic game theory, which is based on the presumption of an ability to anticipate causal relationships between actions and reactions.
19. Richard Danzig, "Driving in the Dark: Ten Propositions about Prediction and National Security," Washington, DC: Center for a New American Security, October 2011, 12–13.
20. These quotes are drawn from Bethany McLean and Peter Elkind, *The Smartest Guys in the Room: The Amazing Rise and Scandalous Fall of Enron* (New York: Penguin, 2003), 114–118.
21. Yaacov Vertzberger argues that this is a fundamental distinction between the economic and political realms—that decision-makers in financial markets "clearly define their acceptable risk levels" whereas foreign policy decision-makers "are not accustomed to defining precisely and systematically their level of acceptable risk prior to making a decision"; in business, he contends, "there exist shared consensual norms of what are reasonable business practices. In politics, on the other hand, a consensual normative framework that distinguishes the gambler from the astute responsible statesman is yet to emerge, let alone become an integral part of the political decision-making culture." See Vertzberger, "Rethinking and Reconceptualizing Risk in Foreign Policy Decision-Making: A Sociocognitive Approach," *Political Psychology*, Vol. 16, No. 2 (June 1995), 347–348.
22. Ben Bernanke, "Modern Risk Management and Banking Supervision," Speech at the Stonier Graduate School of Banking, June 12, 2006; available at <http://www.federalreserve.gov/newsevents/speech/bernanke20060612a.htm>.
23. Paul Davidson, "Is Risk Management a Science?" Unpublished conference paper, May 15, 2012, 1–2; available at <http://econ.bus.utk.edu/documents/davidsonpapers/primatalk15.pdf>.
24. Gretchen Morgenson and Joshua Rosner, *Reckless Endangerment: How Outsized Ambition, Greed and Corruption Led to Economic Armageddon* (New York: Times Books, 2011), 129.
25. See, for example, Benoit Mandelbrot and Richard L. Hudson, *The (Mis)Behavior of Markets: A Fractal View of Financial Turbulence* (New York: Basic Books, 2004), 5, 11.
26. A January 2012 survey by the consulting firm PwC, for example, concluded both that "the risk frameworks and processes that are currently in place in their organisations are no longer giving them the level of protection they need" and that they were seeing "rapid increases" in the speed of events and "the extent to which their impacts on the business are 'contagious.'" The implication was that "large organisations may now have blind spots from which high-impact risks could emerge to damage or potentially destroy their business." PriceWaterhouseCooper (PwC), "Black Swans Turn Grey: The Transformation of Risk," January 2012, 2–3.

27. Quoted in John Cassidy, *How Markets Fail: The Logic of Economic Calamities* (New York: Farrar, Straus and Giroux, 2009), 269.

2 Defining Risk

1. U.S. Department of Defense, *Quadrennial Defense Review 2014* (Washington, DC: DoD, 2014), 61–62.
2. U.S. DoD, *Quadrennial Defense Review 2014*, 61; see also 62–63.
3. On these aspects of strategic logic, see Harry R. Yarger, “Strategic Theory for the 21st Century: The Little Book on Big Strategy,” U.S. Army War College Strategic Studies Institute *Letort Papers*, February 2006.
4. Robert S. Kaplan and Anette Mikes, “Managing Risks: A New Framework,” *Harvard Business Review*, June 2012; available at <https://hbr.org/2012/06/managing-risks-a-new-framework>.
5. James Kitfield, “Dempsey Wants to ‘Rebalance’ Away from the Use of Military Force,” *Defense One*, May 12, 2014; available at <http://www.defenseone.com/ideas/2014/05/dempsey-wants-rebalance-away-use-military-force/84271/print/>.
6. John Adams, *Risk* (London: Routledge, 1995), 8.
7. Douglas W. Hubbard, *The Failure of Risk Management: Why It’s Broken and How to Fix It* (Hoboken, NJ: Wiley, 2009), 8.
8. Deloitte & Touche LLP, *Global Risk Management Survey: Sixth Edition: Risk Management in the Spotlight* (New York: Deloitte and Touche, Global Financial Services Industry Practice, June 2009).
9. United States Department of Homeland Security, *Risk Management Fundamentals* (Washington, DC: DHS, April 2011), 7.
10. Baruch Fischhoff and John Kadvany, *Risk: A Very Short Introduction* (London: Oxford University Press, 2011), 5, 22.
11. U.S. Department of Defense, *Risk Management Guide for DoD Acquisition*, 6th ed. (Washington, DC: DoD, August 2006), 1.
12. James G. March and Zur Shapira, “Managerial Perspectives on Risk and Risk Taking,” *Management Science*, Vol. 33, No. 11 (November 1987), 1405–1407.
13. The noted scholar Philip Tetlock has for several years conducted ground-breaking experiments into the ways in which groups of forecasters, using sophisticated analytical techniques and engaging in rigorous updating as new information becomes available, can significantly improve the success rate of probabilistic estimates; see Philip E. Tetlock and Dan Gardner, *Superforecasting: The Art and Science of Prediction* (New York: Crown, 2015). I will make reference to Tetlock’s work in a number of places, but even he admits the fundamental distinction between “clocklike” and “cloudlike” issues, and notes carefully that the more uncertain and nonlinear a question is, the less accurate the forecasters can be. In this study I am interested in precisely the sort of complex strategic judgments that are largely immune to accurate anticipation—and that characterize a majority of the biggest decisions made by senior officials in government and business.
14. See, for example, Mark Blyth, “Coping with the Black Swan: The Unsettling World of Nassim Taleb,” *Critical Review*, Vol. 21, No. 4 (2009), 448–449.
15. Adams, *Risk*, 26.
16. Ian Goldin and Mike Mariathan, in a thoughtful recent treatment of systemic risk, make the point that increasing global complexity and interaction is making uncertainty even more intense—and making it even more difficult to use probabilistic forms of risk management. Analyzing risk in terms of identifiable, quantifiable

- outcomes demands a clear view of causalities, and this is becoming more and more difficult. The “specification of contingencies,” they argue, becomes “progressively more difficult as transport, communication, and financial and other world systems become increasingly integrated. This is because we start to lose sight of the effects of individual actions, introducing uncertainty and hazard.” See Ian Goldin and Mike Mariathasan, *The Butterfly Defect: How Globalization Creates Systemic Risks, and What to Do about It* (Princeton, NJ: Princeton University Press, 2014), 26–27.
17. Yaacov Vertzberger, “Rethinking and Reconceptualizing Risk in Foreign Policy Decision-Making: A Sociocognitive Approach,” *Political Psychology*, Vol. 16, No. 2 (June 1995), 351.
 18. Of the many recent applications of complex models to financial markets in ways that reflect significant success, for example, is that of James H. Simons, the legendary founder of the Renaissance Technologies hedge fund. The fund seeks out non-random, predictable patterns in short-term prices, and crunches mind-boggling amounts of data to discover these patterns. The results are impressive, including average returns in its flagship fund of over 45% and returns of over 70% in one year. See, for example, “A Dozen Things I’ve Learned from Jim Simons,” *25iq*, n.d., available at <http://25iq.com/2014/07/09/a-dozen-things-ive-learned-from-jim-simons/>; Scott Patterson and Jenny Strausburg, “Pioneering Fund Stages Second Act,” *Wall Street Journal*, March 16, 2010. James Owen Weatherall places this fund in larger context in *The Physics of Wall Street: A Brief History of Predicting the Unpredictable* (New York: Mariner Books, 2013).
 19. Jonathan Baron defines and discusses various theories of probability in *Thinking and Deciding*, 4th ed. (New York: Cambridge University Press, 2008), 103–111.
 20. See, for example, Nabil I. Al-Najjar and Luciano de Castro, “Subjective Probability,” unpublished ms., Kellogg School of Business, Northwestern University, March 2010; available at http://www.kellogg.northwestern.edu/faculty/alnajjar/papers/Subjective_Probability.pdf.
 21. For a classic example specifically dealing with subjective probability, see Daniel Kahneman and Amos Tversky, “Subjective Probability: A Judgment of Representativeness,” *Cognitive Psychology*, Vol. 3 (1972). 430–454.
 22. See, for example, Tetlock and Gardner, *Superforecasting*, and Tara Isabella Burton, “Could You Be a Super-Forecaster?” *BBC Online*, January 20, 2015; available at <http://www.bbc.com/future/story/20150120-are-you-a-super-forecaster>. The Good Judgment Project itself can be found at <http://www.goodjudgmentproject.com/>.
 23. Prakash Loungani, “The Arcane Art of Predicting Recessions,” *Financial Times*, December 18, 2000; available at <http://www.imf.org/external/np/vc/2000/121800.htm>. On problems with more general economic forecasting see Mark Thoma, “Why Are Economic Forecasts Wrong So Often?” *CBS News*, September 29, 2014; available at <http://www.cbsnews.com/news/why-are-economic-forecasts-wrong-so-often/>.
 24. A simple but useful typology can be found at, “Risk Management Is More Than Just Risk Mitigation,” Oulixeus Thought Leadership Blog, March 5, 2010; available at <http://www.oulixeus.com/2010/03/risk-management-is-more-than-just-risk-mitigation/>.
 25. Jack Dowie, “Against Risk,” *Risk Decision and Policy*, Vol. 4 (1999). 57–73.
 26. Jack Dowie, “Communication for Better Decisions: Not about ‘Risk,’” *Health, Risk and Society*, Vol. 1, No. 1 (1999), 41.
 27. Jack Dowie, “A Risky Decision: Managing without Risk,” *Risk Management*, Vol. 2, No. 2 (2000), 55.
 28. Dowie, “A Risky Decision,” 57.
 29. Dowie, “Communication for Better Decisions,” 43.
 30. Dowie, “Communication for Better Decisions,” 49–50.

3 Approaches to Risk in National Security

1. For an interesting example of a unit-level ORM process, see the 2012 instruction from Training Air Wing Five at <http://www.cnatra.navy.mil/tw5/docs/instructions/3500.1D.pdf>.
2. For details see Frank Camm, Lauren Caston, Alexander C. Hou, Forrest E. Morgan, and Alan Vick, *Managing Risk in USAF Force Planning* (Santa Monica, CA: The RAND Corporation, 2009), 14–18.
3. The Air Force has since replaced that approach with a modified and updated risk management framework. The managers of this process understand well many of the potential dangers in the employment of risk management, and are trying to use the process to generate the sort of high-level strategic dialogues recommended in Part III.
4. See *Risk Management Guide for DoD Acquisition*, 6th ed., August 2006; available at <http://www.dau.mil/publications/publicationsDocs/RMG%206Ed%20Aug06.pdf>.
5. U.S. Department of Defense, *Quadrennial Defense Review Report* (Washington, DC: DoD, September 30, 2001), 57–65.
6. U.S. Department of Defense, *Quadrennial Defense Review Report* (Washington, DC: DoD, February 2010), 89–95. The 2010 QDR in fact contains perhaps the most elaborate discussion of the risk management framework in DoD. Under operational risk, it suggested that “key issues” posing risk included “providing sufficient enabling capabilities, building partnership capacity, and securing DoD systems in cyberspace.” Under force management risk, it listed “supporting operations in Afghanistan and Iraq, providing health care to DoD personnel, and ensuring the proper mix and roles of the Active Component and Reserve Component.” Leading risks on the minds of Defense decision-makers in the area of institutional risk “include reforming general acquisition processes, optimizing information technology acquisition processes, and maintaining the defense industrial base.” And finally, key future challenges risks were “managing uncertainty about the future environment and science and technology (S&T) trends.”
7. Steve Inskeep, “Afghanistan’s Way Forward: A Talk with Gen. John Campbell, Decoded,” NPR.org, November 11, 2014.
8. Air Force Instruction 10-281, “Force Development Concepts,” October 23, 2014, 14; available at http://static.e-publishing.af.mil/production/1/af_a3_5/publication/afi10-2801/afi10-2801.pdf.
9. Much of the detail below comes from the DHS Risk Management office web site, at <http://www.dhs.gov/office-risk-management-and-analysis-mission#wcm-survey-target-id>. See also Clark Murdock, Matt Squeri, Chris Jones, and Becca S. Smith., *Risk Management in Non-DoD U.S. Government Agencies and the International Community* (Washington, DC: Center for Strategic and International Studies, 2011), 43–47. For information on risk management fundamentals as viewed by DHS, see the “Risk Fundamentals Report” at <http://www.dhs.gov/risk-management-series#1>.
10. U.S. Department of Homeland Security, “Strategic National Risk Assessment,” December 2011; available at <http://www.dhs.gov/xlibrary/assets/rma-strategic-national-risk-assessment-ppd8.pdf>.
11. A nongovernmental report that reflects largely the same approach is the World Economic Forum’s “Global Risk Report.” It surveys a “global risk landscape” of things that could go wrong in the international system, from disease to financial crises to conflict. Again the focus is on defining and assessing contextual threats that pose risks only to general national interests. See, for example, World Economic Forum, *Global Risk Report*, 10th ed. (2015); available at <http://reports.weforum.org/global-risks-2015/>.

12. For a summary of NASA's approach, see *NASA Risk Management Handbook* (Washington, DC: NASA Headquarters, November 2011); available at http://www.hq.nasa.gov/office/codeq/doctree/NHBK_2011_3422.pdf.
13. Many aspects of NASA's approach to risk can be found at the Agency's "Risk Management Page," available at <http://www.hq.nasa.gov/office/codeq/risk/>. See especially the articles on "A Paradigm Shift" and "Risk-Informed Decision Making."
14. Camm et al., *Managing Risk in USAF Force Planning*, 4.
15. United States Department of Defense, *Quadrennial Defense Review 2014* (Washington, DC: Department of Defense, March 2014), 53–57.
16. United States Army Training and Doctrine Command, *The U.S. Army Operating Concept: Win in a Complex World* (Joint Base Langley-Eustis, VA: TRADOC, October 31, 2014), 42–43.
17. Gina Harkins, "Amos: USMC of 175K Would Be 'Moderate Risk' Force," *Marine Times*, March 12, 2014.
18. David Vergun, "Odierno: Brigade Readiness Half of What It Should Be," March 11, 2015; available at http://www.army.mil/article/144302/Odierno_Brigade_readiness_half_what_it_should_be/.
19. Camm et al., *Managing Risk in USAF Force Planning*, 103–104.
20. For a depiction and description of this chart see Elizabeth Bumiller, "We Have Met the Enemy and He Is Powerpoint," *New York Times*, April 26, 2010; available at http://www.nytimes.com/2010/04/27/world/27powerpoint.html?_r=0.

4 Risk and Uncertainty

1. Scott Patterson, *The Quants: How a New Breed of Math Whizzes Conquered Wall Street and Nearly Destroyed It* (New York: Crown Business, 2010), 97.
2. Some analyses distinguish uncertainty from ambiguity and ignorance; uncertainty, in this account, is all about preparing for various possible future worlds, whereas ambiguity refers specifically to the lack of precision about specific variables and ignorance is lack of information. See, for example, Andy Stirling, "Keep It Complex," *Nature*, Vol. 468 (December 2010), 1029–1031., and Alan Bond, A. Morrison-Saunders, J. A. Gunn, J. Pope, and F. Retief, "Managing Uncertainty, Ambiguity and Ignorance in Impact Assessment by Embedding Evolutionary Resilience, Participatory Modelling and Adaptive Management," *Journal of Environmental Management*, Vol. 151 (2015), 98–99. My approach defines uncertainty as a combination of epistemological and ontological variants—that is, lack of information combined with nonlinear emergence.

International relations theory employs the term in a very specific way—as a representation of the doubt and ignorance that exists for states in an anarchic environment; see Brian C. Rathbun, "Uncertain about Uncertainty: Understanding the Multiple Meanings of a Crucial Concept in International Relations Theory," *International Studies Quarterly*, Vol. 51 (2007), 533–557. National decision-makers lack critical information about the intentions of other actors. In realism, this uncertainty is the engine of fear, and thus of spirals of misperception and conflict. Such ignorance can be one result of uncertainty as I mean it, but for the purposes of this analysis I am using the version of uncertainty more traditionally associated with risk.

3. An interesting discussion of the definition of uncertainty is in Raanan Lipshitz and Orna Strauss, "Coping with Uncertainty: A Naturalistic Decision-Making Analysis," *Organizational Behavior and Human Decision Processes*, Vol. 69, No. 2 (February 1997), 149–150. They stress the importance of subjective perception of uncertainty—when

- people feel a situation to be uncertain, then it is. On the same distinction see Fernando J. Cardim de Carvalho, "Keynes on Probability, Uncertainty, and Decision Making," *Journal of Post-Keynesian Economics*, Vol. XI, No. 1 (Fall 1988), 66–67. This wonderful piece is perhaps the best single description of the character of classic uncertainty in social contexts. Philip Tetlock discusses epistemic versus aleatory uncertainty in Philip Tetlock and Dan Gardner, *Superforecasting: The Art and Science of Prediction* (New York: Crown, 2015), 143–144.
4. Daniel Ellsberg, "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics*, Vol. 75, No. 4 (November 1961), 650–656, 666.
 5. Nate Silver, *The Signal and the Noise: Why So Many Predictions Fail—But Some Don't* (New York: Penguin Press, 2012), 390.
 6. Peter Earl, *The Economic Imagination: Towards a Behavioural Analysis of Choice* (Armonk, NY: M. E. Sharpe, 1983), 124.
 7. Silver, *The Signal and the Noise*, 118–119.
 8. Barry Schacter, "The New Risk Management Framework after the 208 Financial Crisis," Presentation to Global Association of Risk Professionals, July 13, 2010, slide 14; available at <http://www.slideshare.net/BarrySchachter/the-new-risk-management-framework-after-the-2008-financial-crisis>. These factors are ably discussed in Carvalho, "Keynes on Probability," 71, 76–78. See also Ian Goldin and Mike Mariathasan, *The Butterfly Defect: How Globalization Creates Systemic Risks, and What to Do about It* (Princeton, NJ: Princeton University Press, 2014), 22.
 9. Tetlock uses these terms in *Superforecasting*, 13.
 10. Paul K. Davis, "Uncertainty-Sensitive Planning," in Stuart E. Johnson, Martin Libicki, and Gregory F. Treverton, eds., *New Challenges, New Tools for Defense Decisionmaking* (Santa Monica, CA: The RAND Corporation, 2003), 133–134.
 11. Federal Open Market Committee, Transcript of Meeting, January 28–29, 2003, 21–24 and 37–38; available at <http://www.federalreserve.gov/monetarypolicy/files/FOMC20030129meeting.pdf>.
 12. Christopher Brown and Cheng Hao, "Treating Uncertainty as Risk: The Credit Default Swap and the Paradox of Derivatives," *Journal of Economic Issues*, Vol. XLVI, No. 2 (June 2012), 307.
 13. G. L. S. Shackle, *Uncertainty in Economics* (Cambridge, UK: Cambridge University Press, 1968), 6. This form of uncertainty is related to George Soros's notion of "reflexivity," the idea that classical economic models don't account for the interaction of actors and the environment. See George Soros, "A General Theory of Reflexivity," *The Financial Times*, October 26, 2009; available at <http://www.ft.com/cms/s/2/0ca06172-bfe9-11de-aed2-00144feab49a.html#axzz31nmz0Ct4>. See also the discussion of interdependence in Mark Blyth, "Ideas, Uncertainty and Evolution," in Robert Cox and Daniel Beland, eds., *Ideas and Politics in Social Science Research* (New York: Oxford University Press, 2010), 92–93.
 14. Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long Term Capital Management* (New York: Random House, 2000), 65; see also 8, 44, 61. For a detailed analysis of the quantitative mindset see also 65–77.
 15. Lowenstein, *When Genius Failed*, 123, 127, 146.
 16. Lowenstein, *When Genius Failed*, 34–35, 59.
 17. See, for example, Martin Wolf, *The Shifts and the Shocks: What We've Learned—and Have Still to Learn—from the Financial Crisis* (New York: Penguin Press, 2014), 132–133.
 18. Benoit Mandelbrot and Richard L. Hudson, *The (Mis)Behavior of Markets: A Fractal View of Financial Turbulence* (New York: Basic Books, 2004), 4.
 19. Andrew Ross Sorkin, *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves* (New York: Penguin, 2010), 4.

20. Nassim Nicholas Taleb, *The Black Swan: The Impact of the Highly Improbable* (New York: Random House, 2007), 282–283. Scott Patterson has a fascinating account of the origins of Black-Scholes and the “random walk” theory of market distributions in *The Quants*, 27–46.
21. Diane Vaughan, *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA* (Chicago: University of Chicago Press, 1996), 421.
22. Joshua Cooper Ramo, *The Age of the Unthinkable* (New York: Little Brown, 2009), 33–34, 49. See also Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011), 215.
23. PwC (PriceWaterHouse Cooper), “Black Swans Turn Grey: The Transformation of Risk,” January 2012, 4. James G. March briefly surveys the history of rationalistic risk analysis in “Rationality, Foolishness, and Adaptive Intelligence,” *Strategic Management Journal*, Vol. 27 (2006), 202–203.
24. Frank H. Knight, *Risk, Uncertainty and Profit*, Reprint of 3rd ed. (New York: Augustus Kelley, 1964), 4–5.
25. Stephen C. Nelson, “Governing Risky and Uncertain Financial Markets,” *Stiftung Wissenschaft und Politik: SWP Comments*, No. 45, December 2012, 2.
26. Doreen M. Tulloch, “Sartrean Existentialism,” *The Philosophical Quarterly*, Vol. 2, No. 6 (January 1952), 33.
27. This is the essential story told in Paul L. Bernstein, *Against the Gods: The Remarkable Story of Risk* (New York: Wiley, 1998).
28. Quoted in Leonard Mlodinow, *The Drunkard’s Walk: How Randomness Rules Our Lives* (New York: Pantheon Books, 2008), 192. For another discussion of the issue see Nate Silver, *The Signal and the Noise: Why So Many Predictions Fail—But Some Don’t* (New York: Penguin Press, 2012), 112–114.
29. Mlodinow, *The Drunkard’s Walk*, 192.
30. “The existence of a complete set of markets for future and contingent claims is a requirement for Pareto-efficiency in Arrow-Debreu economies,” two scholars have argued. “Viewed in this context, financial engineers were seen as doing the important work of completing markets, or of creating financial instruments that permit gambles on all possible future states of the world. Thus, the proliferation of derivatives served the desideratum of pushing real life to closer approximation of the Arrow-Debreu economy, conceived of as a kind of Platonic form.” See Brown and Hao, “Treating Uncertainty as Risk,” 305.
31. Douglas W. Hubbard, *The Failure of Risk Management: Why It’s Broken and How to Fix It* (Hoboken, NJ: Wiley, 2009), 61, 71–72. His critique is that risk management is too subjective and qualitative and doesn’t use “sophisticated” models employed by actuaries; but in this he’s rushing past distinction between risk and uncertainty (cf. 145–164). Actuarial models don’t work for radical uncertainty.
32. See, for example, J. M. Hampton, P. G. Moore, and H. Thomas, “Subjective Probability and Its Measurement,” *Journal of the Royal Statistical Society, Series A (General)*, Vol. 136, No. 1 (1973), 21–42.
33. Taleb, *The Black Swan*, 35, 146. An interesting discussion of the same basic discussion can be found in Charles R. Schwenk, “Cognitive Simplification Processes in Strategic Decision Making,” *Strategic Management Journal*, Vol. 5, No. 2 (April–June 1984), 112. The challenges with probabilistic assumptions in a nonlinear context are also described in Patterson, *The Quants*, 192–195, and Roger Lowenstein, *The End of Wall Street* (New York: Penguin Press, 2010), 44–46.
34. Brown and Hao, “Treating Uncertainty as Risk,” 307. For a similar distinction see Robin M. Hogarth, “Subways, Coconuts, and Foggy Minefields,” in Erwann Michale-Kerjan and Paul Slovic, eds., *The Irrational Economist: Making Decisions in a Dangerous World* (New York: PublicAffairs Books, 2010), 22–24.

35. Brown and Hao, "Treating Uncertainty as Risk," 307. See also Paul Davidson, "Risk and Uncertainty in Economics," unpublished conference paper, February 6, 2009, 5–6; available at <http://econ.bus.utk.edu/documents/davidsonpapers/RISKANDUNCERTAINTYINECONOMICS.pdf>.
36. Bernstein, *Against the Gods*, 119–121, 330–331; see also Mlodinow, *The Drunkard's Walk*, 194.
37. Knight, *Risk, Uncertainty and Profit*, 251.
38. David H. Krantz, "Constructed Preference and the Quest for Rationality," in Erwann Micel-Kerjan and Paul Slovic, eds., *The Irrational Economist: Making Decisions in a Dangerous World* (New York: PublicAffairs, 2010), 65–66.
39. Mary Douglas and Aaron Wildavsky, *Risk and Culture: An Essay on the Selection of Technological and Environmental Dangers* (Berkeley: University of California Press, 1982), 192.
40. The problem of conceiving judgment under uncertainty in this sense is similar to the concept of "future-choice" decisions, whose basic elements match a number of these criteria. Hogarth, "Subways, Coconuts and Foggy Minefields," 21. In one very useful work on judgment under uncertainty, Young Bak Choi defines uncertainty in part through a number of specific characteristics. These include the "relative complexity of calculation," the unpredictability of the future, the interdependence of human actions, and the limits of our information. Young Bak Choi, *Paradigms and Conventions: Uncertainty, Decision Making and Entrepreneurship* (Ann Arbor: University of Michigan Press, 1993), 11–16.
41. Taleb, *The Black Swan*, 36.
42. Knight, *Risk, Uncertainty and Profit*, 216–218 and 226. He also suggests (16) that "We have no way of discussing a force or change except to describe its effects or results under given conditions."
43. Shackle, *Uncertainty in Economics*, 4–5, 63. Elsewhere (26) he refers to "non-divisible, non-serial" experiments in which frequency-distribution analysis simply doesn't apply.
44. Bernstein, *Against the Gods*, 14, 44, 144–145.
45. Silver, *The Signal and the Noise*, 237–245.
46. Lowenstein, *The End of Wall Street*, 44.
47. See, for example, Paul Davidson, "Is Risk Management a Science?" Unpublished conference paper, May 15, 2012, 3–4; available at <http://econ.bus.utk.edu/documents/davidsonpapers/primatalk15.pdf>. See also Mandelbrot and Hudson, *The (Mis) Behavior of Markets*, 8. Mark Blyth discusses the assumption of linear causality in "Ideas, Uncertainty and Evolution," in Robert Cox and Daniel Beland, eds., *Ideas and Politics in Social Science Research* (New York: Oxford University Press, 2010), 84–86. His analysis of hidden generators of events is really an argument about hidden causality; see Mark Blyth, "Coping with the Black Swan: The Unsettling World of Nassim Taleb," *Critical Review*, Vol. 21, No. 4 (2009), 450.
48. Schacter, "The New Risk Management Framework," slide 6.
49. Nelson, "Governing Risky and Uncertain Financial Markets," 6.
50. PwC, "Black Swans Turn Grey," 5.
51. See, for example, Ranaan Lipschitz and Orna Strauss, "Coping with Uncertainty: A Naturalistic Decision Making Analysis," *Organizational Behavior and Human Decision Processes*, Vol. 69, No. 2 (February 1997), 149–163, and James H. Barnes, Jr., "Cognitive Biases and Their Impact on Strategic Planning," *Strategic Management Journal*, Vol. 5, No. 2 (April–June 1984), 134.
52. Schacter, "The New Risk Management Framework," slide 24.
53. One powerful example was the "Gaussian copula function," an equation designed to comprehend correlated risks which proved unable to grasp the nonlinear complexities of actual markets. See Felix Salmon, "Recipe for Disaster: The Formula

- That Killed Wall Street," *Wired*, February 23, 2009, available at <http://www.wired.com/2009/02/wp-quant/>.
54. Cited in Patterson, *The Quants*, 294.
 55. Lowenstein, *The End of Wall Street*, 45.
 56. Silver, *The Signal and the Noise*, 80, 270–271, 426.
 57. Stephen C. Nelson and Peter J. Katzenstein, "Uncertainty, Risk and the Financial Crisis of 2008," *International Organization*, Vol. 68, No. 2 (March 2014), 378. See also John Cassidy, *How Markets Fail: The Logic of Economic Calamities* (New York: Farrar, Straus and Giroux, 2009), 274–279.
 58. Joe Nocera, "Risk Management," *New York Times*, January 2, 2009.
 59. Simon Johnson and James Kwak, "Seduced by a Model," *New York Times*, October 1, 2009.
 60. Cassidy, *How Markets Fail*, 269.
 61. Brown and Hao, "Treating Uncertainty as Risk," 305.
 62. Barry Schacter, "The New Risk Management Framework," slide 4.

5 Risk Is What We Make of It

1. Brian Burrough, "Bringing Down Bear Sterns," in Graydon Carter, ed., *The Great Hangover* (New York: HarperPerennial, 2010), 4.
2. Burrough, "Bringing Down Bear Sterns," 16.
3. Lehman Brothers' leader Dick Fuld constantly blamed short sellers and the perceptions they created for destroying his company. The problem wasn't their liquidity, it was perceptions of it. Andrew Ross Sorkin, *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves* (New York: Penguin, 2010), 14–16, 108, 434. The perceptual character of the risk involving Bear is also discussed in William D. Cohan, *House of Cards: A Tale of Hubris and Wretched Excess on Wall Street* (New York: Doubleday, 2009), 5–7.
4. See Paul Blustein, *The Chastening: Inside the Crisis That Rocked the Global Financial System and Humbled the IMF* (New York: PublicAffairs, 2001), 2, 7–8.
5. Diane Vaughn, "The Dark Side of Organizations: Mistake, Misconduct, and Disaster," *Annual Review of Sociology*, Vol. 25 (1999), 295. See also Mark Simon, Susan M. Houghton, and Karl Aquino, "Cognitive Biases, Risk Perception, and Venture Formation: How Individuals Decide to Start Companies," *Journal of Business Venturing*, Vol. 15 (1999), 114–115.
6. Quoted in Peter Earl, *The Economic Imagination: Towards a Behavioural Analysis of Choice* (Armonk, NY: M. E. Sharpe, 1983), 117.
7. Paul Slovic, ed., *The Perception of Risk* (London: Earthscan Publications, 2000).
8. Diane Vaughan, *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA* (Chicago: University of Chicago Press, 1996), 62.
9. Baruch Fischhoff and John Kadvany, *Risk: A Very Short Introduction* (London: Oxford University Press, 2011), 33–34.
10. See Scott Plous, *The Psychology of Judgment and Decision Making* (New York: McGraw-Hill, 1993), 18–20.
11. Fischhoff and Kadvany, *Risk*, 73. See also Leonard Mlodinow, *The Drunkard's Walk: How Randomness Rules Our Lives* (New York: Pantheon Books, 2008), 173.
12. Mlodinow, *The Drunkard's Walk*, 189.
13. Mark Blyth, "Coping with the Black Swan: The Unsettling World of Nassim Taleb," *Critical Review*, Vol. 21, No. 4 (2009), 449.
14. Nassim Nicholas Taleb, *The Black Swan: The Impact of the Highly Improbable* (New York: Random House, 2007), 63–64.

15. Paul L. Bernstein, *Against the Gods: The Remarkable Story of Risk* (New York: Wiley, 1998), 202–203, 206–207.
16. Frank H. Knight, *Risk, Uncertainty and Profit*, Reprint of 3rd ed. (New York: Augustus Kelley, 1964), 213–214.
17. John D. Steinbruner, *The Cybernetic Theory of Decision: New Dimensions of Political Analysis* (Princeton, NJ: Princeton University Press, 1974), 15. See also Fernando J. Cardim de Carvalho, “Keynes on Probability, Uncertainty, and Decision Making,” *Journal of Post-Keynesian Economics*, Vol. XI, No. 1 (Fall 1988), 74–75. In uncertain contexts, Carvalho explains, “the decision maker has to fill the voids, has to ‘create’ the additional premises which may be needed.” The result is judgment under “imagined premises.”
18. Mlodinow, *The Drunkard’s Walk*, 171.
19. Described in Earl, *The Economic Imagination*, 136–137.
20. Earl, *The Economic Imagination*, 57.
21. Young Bak Choi, *Paradigms and Conventions: Uncertainty, Decision Making and Entrepreneurship* (Ann Arbor: University of Michigan Press, 1993), 21.
22. Stephen C. Nelson and Peter J. Katzenstein, “Uncertainty, Risk and the Financial Crisis of 2008,” *International Organization*, Vol. 68, No. 2 (March 2014), 362.
23. Knight, *Risk, Uncertainty and Profit*, 199, 211.
24. Knight, *Risk, Uncertainty and Profit*, 223, 226, 230.
25. Earl, *The Economic Imagination*, 65–66, 138–139. Alan Greenspan has suggested a similar mental process. He aligns himself with Kahneman’s intuitive System 1 for most everyday thinking, especially in areas where we have acquired some experience and knowledge. “We pour information into our mind,” Greenspan writes, “and, with a delay, out pop epiphanies”; Alan Greenspan, *The Map and the Territory: Risk, Human Nature, and the Future of Forecasting* (New York: Penguin, 2013), 35.
26. Bruce Mangan, “Sensation’s Ghost: The Non-Sensory ‘Fringe’ of Consciousness,” *Psyche*, Vol. 7, No. 18 (October 2001). An article that offers some critiques of his views, and contrasts some aspects with Damasio’s thinking, is Erik Woody and Henry Szechtman, “The Sensation of Making Sense,” *Psyche*, Vol. 8, No. 20 (October 2002).
27. Robert A. Burton, *On Being Certain: Believing You Are Right Even When You’re Not* (New York: St. Martin’s Griffin, 2008).
28. Vaughan, *The Challenger Launch Decision*, 37.
29. Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long Term Capital Management* (New York: Random House, 2000), 220, 228–229.
30. See, for example, Joseph P. Simmons and Leif D. Nelson, “Intuitive Confidence and the Prominence Effect”; and Simmons and Nelson, “Intuitive Confidence: Choosing Between Intuitive and Nonintuitive Alternatives,” *Journal of Experimental Psychology: General*, Vol. 135, No. 3 (August 2006), 409–428.
31. See, for example, the analysis of decision style on risk perception in John C. Henderson and Paul C. Nutt, “The Influence of Decision Style on Decision Making Behavior,” *Management Science*, Vol. 26, No. 4 (April 1980), 371–386. “Disputes about risk in which the participants hurl charges of stupidity and irrationality at each other,” John Adams writes, “are usually seen upon dispassionate inspection to be arguments in which the participants are arguing from different premises, different paradigms, different world views—different myths of nature, both physical and human.” John Adams, *Risk* (London: Routledge, 1995), 37, 50.
32. René M. Stulz, “Risk Management Failures: What Are They and When Do They Happen?” Cornerstone Research Organization report, 2009.
33. Nassim Taleb uses that term in *Fooled by Randomness: The Hidden Role of Chance in Life and Markets* (New York: Random House, 2004), xliii.

34. George Loewenstein and Erik Angner, "Predicting and Indulging Changing Preferences," in Loewenstein, Daniel Read, and Roy F. Baumeister, eds., *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice* (New York: Sage Foundation, 2003), 352.
35. Steinbruner, *The Cybernetic Theory of Decision*, 113.
36. Shackle discusses the gap between our imagination of the future and the reality and says that the way we fill it is with belief; G. Shackle, *Uncertainty in Economics* (Cambridge, UK: Cambridge University Press, 1968), 21.
37. Burton, *On Being Certain*, 177–178.

6 Indifferent to Consequences

1. Johnson quoted in Michael Beschloss, ed., *Reaching for Glory: Lyndon Johnson's Secret White House Tapes, 1964–1965* (New York: Simon and Schuster, 2001), 214–215.
2. Johnson quoted in Beschloss, ed., *Reaching for Glory*, 211–213, 345, 365.
3. Fernando J. Cardim de Carvalho, "Keynes on Probability, Uncertainty, and Decision Making," *Journal of Post-Keynesian Economics*, Vol. XI, No. 1 (Fall 1988), 68–69. Eldar Shafir and Amos Tversky write that "Most conceptions of decision-making under uncertainty ... are *consequentialist* in the sense that decisions are determined by an assessment of the potential consequences and their perceived likelihood." See Eldar Shafir and Amos Tversky, "Thinking through Uncertainty: Nonconsequential Reasoning and Choice," *Cognitive Psychology*, Vol. 24 (1992), 450.
4. Raanan Lipshitz, "The Road to Desert Storm," *Organization Studies*, Vol. 16, No. 2 (1995), 243–244. James G. March writes that "Rational choice involves two kinds of guesses: guesses about future consequences of current actions and guesses about future preferences for those consequences"; see James G. March, "Bounded Rationality, Ambiguity, and the Engineering of Choice," *The Bell Journal of Economics*, Vol. 9, No. 2 (Autumn 1978), 589.
5. He gave essentially the same speech to the journalist Drew Pearson on March 23: "I don't believe I can walk out," he said. "I'd be another Chamberlain and ... we'd have another Munich. The aggressors feed on blood. ... *They* may get another president, but I'm not going to pull out"; Johnson quoted in Beschloss, ed., *Reaching for Glory*, 181–182, 238.
6. See, for example, Robert Dallek, *Flawed Giant: Lyndon Johnson and His Times, 1961–1973* (New York: Oxford University Press, 1988), 340, 343, 350, and 379.
7. Gordon M. Goldstein, *Lessons in Disaster: McGeorge Bundy and the Path to War in Vietnam* (New York: Times Books, 2008), 160.
8. Shafir and Tversky survey a number of mechanisms that interfere with consequentialist decision-making in "Thinking through Uncertainty," 451–465.
9. Such a preference has been termed procedural utility. See Bruno S. Frey and Alois Stutzer, "Beyond Outcomes: Measuring Procedural Utility," *Oxford Economic Papers*, Vol. 57 (2005), 90–91.
10. See, for example, Paul Slovic, "Choice Between Equally Valued Alternatives," *Journal of Experimental Psychology: Human Perception and Performance*, Vol. 1, No. 3 (1975), 280–287; Paul Slovic, "The Construction of Preference," *American Psychologist*, Vol. 50, No. 5 (1995), 364–371; Amos Tversky, Shmuel Sattath, and Paul Slovic, "Contingent Weighting in Judgment and Choice," *Psychological Review*, Vol. 95, No. 3 (July 1988), 371–384; Joseph P. Simmons and Leif D. Nelson, "Intuitive

- Confidence and the Prominence Effect: When Consumer Choices Are Sensitive to Matching Prices," unpublished paper, n.d., available at <http://pages.stern.nyu.edu/~lnelson0/Simmons%20%20Nelson%20JMR%20Submission.pdf>; and Gregory W. Fischer, Ziv Carmon, Dan Ariely, and Gal Zauberman, "Goal-based Construction of Preferences: Task Goals and the Prominence Effect," *Management Science*, Vol. 45, No. 8 (August 1999), 1057–1075.
11. This factor is not identical to the imperative effect. The prominence effect examines cases in which people's main goal is justifying their action, and so they focus on the most prominent defensible objective and downgrade other considerations. In imperative-driven thinking, leaders' vision is similarly narrowed but in a slightly different way. The drive for justification may well be present, but the perceived urgency of the imperative is the primary engine of the judgment.
 12. Paul Slovic, Melissa L. Finucane, Ellen Peters, and Donald G. MacGregor, "Risk as Analysis and Risk as Feelings: Some Thoughts about Affect, Reason, Risk, and Rationality," *Risk Analysis*, Vol. 24, No. 2 (2004), 311–314. See also George F. Loewenstein, Elke U. Webber, Christopher K. Hsee, and Ned Welch, "Risk as Feelings," *Psychological Bulletin*, Vol. 127, No. 2 (2001), 267–286.
 13. See, for example, George Loewenstein, "Out of Control: Visceral Influences on Behavior," *Organizational Behavior and Human Decision Processes*, Vol. 65, No. 3 (March 1996), 272–292.
 14. Shane Frederick, George Loewenstein, and Ted O'Donoghue, "Time Discounting and Time Preference: A Critical Review," in George Loewenstein, Daniel Read, and Roy F. Baumeister, eds., *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice* (New York: Sage Foundation, 2003), 372.
 15. Slovic et al., "Risk as Analysis and Risk as Feelings," 318–319. The fact that present-day decision-makers cannot adequately account for the outcomes of their actions in their judgments, these scholars conclude, counts as "a strong repudiation of the model of informed rational choice."
 16. Loewenstein, "Out of Control," 273. See also p. 289 on "arational" decision-making. See also Frederick et al., "Time Discounting and Time Preference," 40.
 17. Shane Frederick, George Loewenstein, and Ted O'Donoghue, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, Vol. 40, No. 2 (June 2002), 353–354. See also the chapters in Loewenstein, Read, and Baumeister, *Time and Decision*.
 18. Alex Kacelnik, "The Evolution of Patience," in Loewenstein, Read, and Baumeister, *Time and Decision*, 117.
 19. Daniel Kahneman and Amos Tversky, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, Vol. 47, No. 2 (March 1979), 265.
 20. A very good analysis is Jonathan Baron, "Nonconsequentialist Decisions," *Behavioral and Brain Sciences*, Vol. 17, No. 1 (1994), 1–10.
 21. This is similar to the concept of "procedural utility," in which decision-makers pay more attention to aspects of the process by which the decision emerges (for example, how democratic it is) than to outcomes; see, for example, Bruno S. Frey and Alois Stutzer, "Beyond Outcomes: Measuring Procedural Utility," *Oxford Economic Papers*, Vol. 57, No. 1 (January 2005), 90–111. Interestingly, one could argue that the US political system favors procedural over outcome utility: The system is judged not by the quality of its choices but by its ability to rationalize various social interests in the process of making them. It can also be seen as a form of "rule utilitarianism," in which the adherence to a strict rule is viewed as providing the greatest good.

22. Jonathan Baron and Mark Spranca, "Protected Values," *Organizational Behavior and Human Decision Processes*, Vol. 70, No. 1 (April 1997), 3. Baron takes care to distinguish protected values from heuristics, which are general simplifying rules whose application can be somewhat flexible.
23. Philip Tetlock, "Thinking the Unthinkable: Sacred Values and Taboo Cognitions," *Trends in Cognitive Sciences*, Vol. 7, No. 7 (July 2003), 320.
24. Tetlock, "Thinking the Unthinkable," 323–324.
25. Christopher Clark, *The Sleepwalkers: How Europe Went to War in 1914* (New York: HarperCollins, 2012), 429, 449–450, 519.
26. Shafir and Tversky, "Thinking through Uncertainty," 465.
27. In some cases, there can be more than one imperative, and they can collide: In the cases of Vietnam and, more recently, the Obama administration's choices on Afghanistan, leaders faced colliding imperatives to win and to withdraw. I would argue that the default example of this effect is one in which a single, unidirectional imperative dictates behavior, but situations can be much more complicated than that. From a definitional perspective, the key question is whether the factors influencing choice reflect the characteristics of an imperative.
28. Imperative-driven thinking can be a function of naturalistic approaches. In naturalistic models of decision-making, decision-makers quickly choose "a 'promising' alternative" rather than the optimal one, and then often engage in the "restructuring of preferences and beliefs to accentuate its superiority over other alternatives." Such a process, which involves a relatively quick, and often highly intuitive, endorsement of the first reasonable alternative, imperatives can play a dominant role, shaping the choice of what that alternative should be. See Raanan Lipshitz, "The Road to Desert Storm," *Organization Studies*, Vol. 16, No. 2 (1995), 243–244, 247.
29. This example was famously used by Gerd Gigerenzer to illustrate the advantages of intuition. See Gerd Gigerenzer, *Gut Feelings: The Intelligence of the Unconscious* (New York: Penguin Books, 2008).
30. In this sense, imperatives reflect aspects of both satisficing and naturalistic decision models, both of which emphasize the importance of the first reasonable alternative that appears to decision-makers.
31. Michael Lewis, "The Man Who Crashed the World," in Graydon Carter, ed., *The Great Hangover* (New York: Harper Perennial, 2010), 106, 116.
32. Bethany McLean, "Fannie Mae's Last Stand," in Carter, ed., *The Great Hangover*, 174.
33. McLean, "Fannie Mae's Last Stand," 161.
34. Andrew Ross Sorkin, *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves* (New York: Penguin, 2010), 14, 88, 125–126, 146–147.
35. Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long Term Capital Management* (New York: Random House, 2000), 84.
36. Lewis, "The Man Who Crashed the World," 115, 122.
37. Lowenstein, *When Genius Failed*, 130.

7 The Swans to Worry About Are Gray

1. Nassim Nicholas Taleb, *The Black Swan: The Impact of the Highly Improbable* (New York: Random House, 2007), xvi–xviii.
2. Taleb, *The Black Swan*, xviii–xix.
3. The PwC analysis "Black Swans Turn Gray: The Transformation of Risk" makes this distinction, although its argument is that black swans are simply becoming

- more frequent. The briefing is available at http://www.pwccn.com/webmedia/doc/635116518906857384_ia_risk_transform_aug2013.pdf.
4. These issues are discussed in Joe Nocera, "Risk Management," *New York Times*, January 9, 2009. To calculate probabilistic risk they require consistent trends likely to continue into the future. Strategists and risk managers, then, without necessarily being explicit about it, typically focus on the majority of risks that can be identified and, often, quantified.
 5. Taleb, *The Black Swan*, 272–273.
 6. Robert Walker, "Beware the 'Gray Swan,'" *The World Post*, May 18, 2011; available at http://www.huffingtonpost.com/robert-walker/beware-the-gray-swan_b_863237.html.
 7. *Investopedia* definition of gray swan; available at <http://www.investopedia.com/terms/g/gray-swan.asp>.
 8. Taleb, *The Black Swan*, 36–37, 309. His use of the term "Mandelbrotian" is to refer to a class of events that follows some rules of order even in a generally nonlinear context.
 9. Taleb, *The Black Swan*, 272.
 10. Adrienne LaFrance, "Disasters That Were Foretold," *The Atlantic*, September 4, 2015.
 11. Taleb, *The Black Swan*, xix.
 12. Richard Clarke, *Against All Enemies: Inside America's War on Terror* (New York: Free Press, 2004), 227.
 13. A declassified version of the memo is available at <http://www2.gwu.edu/~nsarchiv/NSAEBB/NSAEBB147/clarke%20memo.pdf>.
 14. The quote is from the executive summary; see http://www.9-11commission.gov/report/911Report_Exec.htm.
 15. Nate Silver, *The Signal and the Noise: Why So Many Predictions Fail—But Some Don't* (New York: Penguin Press, 2012), 22–26. See also the analysis in Carolyn Kousky, John Pratt, and Richard Zeckhauser, "Virgin Versus Experienced Risks," in Erwann Micel-Kerjan and Paul Slovic, eds., *The Irrational Economist: Making Decisions in a Dangerous World* (New York: PublicAffairs, 2010), 99–103.
 16. See, for example, Charles Mackay, *Extraordinary Popular Delusions and the Madness of Crowds* (Renaissance Classics, 2012).
 17. John Cassidy, *How Markets Fail: The Logic of Economic Calamities* (New York: Farrar, Straus and Giroux, 2009), 18–19, 282.
 18. Andrew Ross Sorkin, *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves* (New York: Penguin, 2010), 5, 49, 60, 146–147.
 19. Michael Lewis, "The Man Who Crashed the World," in Graydon Carter, ed., *The Great Hangover* (New York: HarperPerennial, 2010), 120.
 20. Among other things, two prominent economists published an essay in the *Journal of Finance* that predicted the sort of market runs that could ruin a leveraged arbitrage firm like LTCM. See Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long Term Capital Management* (New York: Random House, 2000), 111.
 21. Cassidy, *How Markets Fail*, 17.
 22. Gretchen Morgenson and Joshua Rosner, *Reckless Endangerment: How Outsized Ambition, Greed and Corruption Led to Economic Armageddon* (New York: Times Books, 2011), 193.
 23. See, for example, Robin Dillon-Merrill, Catherine H. Tinsley, and Matthew A. Cronin, "How Near-Miss Events Amplify or Attenuate Risky Decision Making," *Published Articles & Papers*, Paper 93 (2012); available at http://research.create.usc.edu/published_papers/93.

24. See, for example, Roy F. Baumeister and Kathleen D. Vohs, "Willpower, Choice, and Self-Control," in George Loewenstein, Daniel Read, and Baumeister, eds., *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice* (New York: Russell Sage Foundation, 2003), 206–209.
25. Dion Scott-Kakures, "Motivated Believing: Wishful and Unwelcome," *Noûs*, Vol. 34, No. 3 (2000), 349.
26. Ziva Kunda, "The Case for Motivated Reasoning," *Psychological Bulletin*, Vol. 108, No. 3 (1990), 484.
27. See, for example, Monica Prasad, Andrew J. Perrin, Kieran Bezila, Steve G. Hoffman, Kate Kindleberger, Kim Manturuk, and Ashleigh Smith Powers, "'There Must Be a Reason': Osama, Saddam, and Inferred Justification," *Sociological Inquiry*, Vol. 79, No. 2 (2009), 143–144. See also Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011).
28. These examples come from Scott-Kakures, "Motivated Believing," 351. The quote at the end of the list is from p. 352.
29. Alexander L. George, *Presidential Decisionmaking in Foreign Policy: The Effective Use of Information and Advice* (Boulder, CO: Westview Press, 1980), 38–39.
30. Kunda, "The Case for Motivated Reasoning," 487.
31. Kunda, "The Case for Motivated Reasoning," 495.
32. David P. Redlawsk, "Hot Cognition or Cool Consideration? Testing the Effects of Motivated Reasoning on Political Decision Making," *The Journal of Politics*, Vol. 64, No. 4 (November 2002), 1040; Serena Chen, "Getting at the Truth or Getting Along: Accuracy- Versus Impression-Motivated Heuristic and Systemic Processing," *Journal of Personality and Social Psychology*, Vol. 71, No. 2 (1996), 262–275; Kunda, "The Case for Motivated Reasoning," 481.
33. Philip E. Tetlock, "Accountability and the Perseverance of First Impressions," *Social Psychology Quarterly*, Vol. 46, No. 4 (1983), 285–292. Yet other research in "debi-asing" has found mixed results—even some counterproductive results, as when "motivating people to think harder will often backfire, amplifying biases rather than attenuating them." See Philip E. Tetlock and Barbara A. Mellers, "The Great Rationality Debate," *Psychological Science*, Vol. 13, No. 1 (January 2002), 97.
34. Scott-Kakures, "Motivated Believing," 366–367.
35. G. L. S. Shackle, *Uncertainty in Economics* (Cambridge, UK: Cambridge University Press, 1968), 8–9.
36. For one account see Walter Pincus and Karen DeYoung, "Analysts Warning of Iraq Chaos Detailed," *Washington Post*, May 26, 2007.

8 Risk Becomes Personalized

1. Bethany McLean and Peter Elkind, *The Smartest Guys in the Room: The Amazing Rise and Scandalous Fall of Enron* (New York: Penguin, 2003), 17.
2. McLean and Elkind, *The Smartest Guys in the Room*, 19–22.
3. Andrew Ross Sorkin, *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves* (New York: Penguin, 2010), 543.
4. There is a substantial empirical literature on the role of personality traits and personal style in shaping risk perception on the part of senior decision-makers. See, for example, John C. Henderson and Paul C. Nutt, "The Influence of Decision Style on Decision Making Behavior," *Management Science*, Vol. 26, No. 4 (April 1980), 387–400; and Paul C. Nutt, "Decision Style and Strategic Decisions of Top Executives," *Technological Forecasting and Social Change*, Vol. 30, No. 1 (August 1986), 39–62.

5. Orlando Figes, *The Crimean War: A History* (New York: Metropolitan Books, 2010), 156.
6. Sorkin, *Too Big to Fail*, 119.
7. Roger Lowenstein, *The End of Wall Street* (New York: Penguin Press, 2010), 32–33.
8. Scott Patterson, *The Quants: How a New Breed of Math Whizzes Conquered Wall Street and Nearly Destroyed It* (New York: Crown Business, 2010), 199–202.
9. William D. Cohan, *House of Cards: A Tale of Hubris and Wretched Excess on Wall Street* (New York: Doubleday, 2009), 209–211. For a discussion of Greenberg's complex views on risk, see pp. 196–197.
10. James G. March and Zur Shapira, "Managerial Perspectives on Risk and Risk Taking," *Management Science*, Vol. 33, No. 11 (November 1987), 1408.
11. McLean and Elkind, *The Smartest Guys in the Room*, 27–28, 40, 77–78, 101.
12. Brian Burrough, "Bringing Down Bear Sterns," in Graydon Carter, ed., *The Great Hangover* (New York: HarperPerennial, 2010), 8.
13. Carol Dweck, "Warning Signs of an Unhealthy Appetite for Risk," *Harvard Business Review*, June 14, 2012; available at <http://blogs.hbr.org/2012/06/how-a-fixed-mindset-feeds-an-u/>.
14. John Cassidy, *How Markets Fail: The Logic of Economic Calamities* (New York: Farrar, Straus and Giroux, 2009), 199.
15. Michael Lewis, "The Man Who Crashed the World," in Carter, ed., *The Great Hangover*, 115–116, 124–125.
16. Michael Lewis, "Wall Street on the Tundra: The Implosion of Iceland's Economy," in Carter, ed., *The Great Hangover*, 213–214, 216–217.
17. Lewis, "Wall Street on the Tundra," 220.
18. James H. Barnes, Jr., "Cognitive Biases and Their Impact on Strategic Planning," *Strategic Management Journal*, Vol. 5, No. 2 (April–June 1984), 133–134.
19. Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011), 87, 255, 257. See also Robert J. Shiller, *Irrational Exuberance*, 2nd ed. (Princeton, NJ: Princeton University Press, 2005), 152–155.
20. McLean and Elkind, *The Smartest Guys in the Room*, 241.
21. Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long Term Capital Management* (New York: Random House, 2000), 13, 166, 197, 211.
22. Kahneman, *Thinking, Fast and Slow*, 263, 212.
23. Cassidy, *How Markets Fail*, 28.
24. Alan Greenspan, *The Map and the Territory: Risk, Human Nature, and the Future of Forecasting* (New York: Penguin, 2013), 5–6.
25. Greg Farrell, *Crash of the Titans: Greed, Hubris, the Fall of Merrill Lynch, and the Near-Collapse of Bank of America* (New York: Crown Business, 2010), 22–29, 69–70, 480.
26. Lowenstein, *When Genius Failed*, xix, 83, 87, 110.
27. McLean and Elkind, *The Smartest Guys in the Room*, 28, 349.
28. McLean and Elkind, *The Smartest Guys in the Room*, 320.
29. Bethany McLean, "Fannie Mae's Last Stand," in Carter, ed., *The Great Hangover*, 155.
30. Gretchen Morgenson and Joshua Rosner, *Reckless Endangerment: How Outsized Ambition, Greed and Corruption Led to Economic Armageddon* (New York: Times Books, 2011), 185–188.
31. Daniel L. Byman and Kenneth M. Pollack, "Let Us Now Praise Great Men: Bringing the Statesman Back In," *International Security*, Vol. 25, No. 4 (Spring 2001), 107–146.
32. Nate Silver, *The Signal and the Noise: Why So Many Predictions Fail—But Some Don't* (New York: Penguin Press, 2012), 57.

9 What You Don't Know Can Destroy You: Ignorance and Correlated Risk

1. Greg Farrell, *Crash of the Titans: Greed, Hubris, the Fall of Merrill Lynch, and the Near-Collapse of Bank of America* (New York: Crown Business, 2010), 16–20.
2. Roger Lowenstein, *The End of Wall Street* (New York: Penguin, 2010), 49, 67, 94.
3. Brian Burrough, “Bringing Down Bear Sterns,” *Vanity Fair*, August 2008, available at http://www.vanityfair.com/news/2008/08/bear_stearns200808.
4. William D. Cohan, *House of Cards: A Tale of Hubris and Wretched Excess on Wall Street* (New York: Doubleday, 2009), 280–282, 291, 307, 311–313.
5. Bethany McLean and Peter Elkind, *The Smartest Guys in the Room: The Amazing Rise and Scandalous Fall of Enron* (New York: Penguin, 2003), 3, 15, 85–86.
6. Michael Lewis, “The Man Who Crashed the World,” in Graydon Carter, ed., *The Great Hangover* (New York: HarperPerennial, 2010), 107.
7. Lewis, “The Man Who Crashed the World,” 109, 117.
8. Andrew Ross Sorkin, *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves* (New York: Penguin, 2010), 90, 124.
9. Lewis, “The Man Who Crashed the World,” 121.
10. Michael Lewis, “Wall Street on the Tundra: The Implosion of Iceland’s Economy,” in Carter, ed., *The Great Hangover*, 218.
11. Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long Term Capital Management* (New York: Random House, 2000), 104.
12. Paul Blustein, *The Chastening: Inside the Crisis That Rocked the Global Financial System and Humbled the IMF* (New York: PublicAffairs Press, 2003), 13, 15.
13. Lewis, “The Man Who Crashed the World,” 119.
14. Of course firms realized the dangers of scale. A key risk management practice was limiting the size of positions that any one trader could amass. But these rules were often enough ignored, modified or bypassed, especially under the influence of the incentives and personalities examined in other chapters.
15. Lowenstein, *When Genius Failed*, 162, 188.
16. Sorkin, *Too Big to Fail*, 5, 59.
17. Blustein, *The Chastening*, 124.
18. Lewis, “The Man Who Crashed the World,” 113, 118, 121.
19. Cited in Gretchen Morgenson and Joshua Rosner, *Reckless Endangerment: How Outsized Ambition, Greed and Corruption Led to Economic Armageddon* (New York: Times Books, 2011), 125.
20. Ian Goldin and Mike Mariathan, *The Butterfly Defect: How Globalization Creates Systemic Risks, and What to Do about It* (Princeton, NJ: Princeton University Press, 2014), 2–3.
21. Goldin and Mariathan, *The Butterfly Defect*, 28.
22. Joshua Cooper Ramo, *The Age of the Unthinkable* (New York: Little, Brown, 2009), 150.
23. Stephen C. Nelson, “Governing Risky and Uncertain Financial Markets,” *Stiftung Wissenschaft und Politik, SWP Comments*, No. 45, December 2012, 4.
24. Goldin and Mariathan, *The Butterfly Defect*, 38, 59.

10 Risk, Incentives, and Culture

1. Mary Douglas and Aaron Wildavsky, *Risk and Culture: An Essay on the Selection of Technological and Environmental Dangers* (Berkeley: University of California Press, 1982), 9, 186, 80, 88.

2. Andrew Ross Sorkin, *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves* (New York: Penguin, 2010), 14, 88, 125–126, 146–147.
3. Greg Farrell, *Crash of the Titans: Greed, Hubris, the Fall of Merrill Lynch, and the Near-Collapse of Bank of America* (New York: Crown Business, 2010), 478.
4. Scott Patterson, *The Quants: How a New Breed of Math Whizzes Conquered Wall Street and Nearly Destroyed It* (New York: Crown Business, 2010), 200–202.
5. Roger Lowenstein, *The End of Wall Street* (New York: Penguin Press, 2010), 33.
6. Bethany McLean, “Fannie Mae’s Last Stand,” in Graydon Carter, ed., *The Great Hangover* (New York: HarperPerennial, 2010), 174. Roger Lowenstein also catalogues the competitive pressures on Fannie and Freddie in *The End of Wall Street*, 11–13.
7. McLean, “Fannie Mae’s Last Stand,” 161.
8. Gretchen Morgenson and Joshua Rosner, *Reckless Endangerment: How Outsized Ambition, Greed and Corruption Led to Economic Armageddon* (New York: Times Books, 2011), 130.
9. Martin Wolf, *The Shifts and the Shocks: What We’ve Learned—and Have Still to Learn—from the Financial Crisis* (New York: Penguin Press, 2014), 134.
10. Lowenstein, *The End of Wall Street*, 74.
11. Morgenson and Rosner, *Reckless Endangerment*, 23.
12. Lewis, “The Man Who Crashed the World,” 106, 116.
13. Lowenstein, *The End of Wall Street*, 3, 51.
14. This argument is made in Morgenson and Rosner, *Reckless Endangerment*, 205.
15. Paul Blustein, *The Chastening: Inside the Crisis That Rocked the Global Financial System and Humbled the IMF* (New York: PublicAffairs, 2001), 48.
16. Andrew Ross Sorkin, *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves* (New York: Penguin, 2010), 124, 130, 146–147.
17. Bethany McLean and Peter Elkind, *The Smartest Guys in the Room: The Amazing Rise and Scandalous Fall of Enron* (New York: Penguin, 2003), 118.
18. Michael Lewis, “The Man Who Crashed the World,” in Carter, ed., *The Great Hangover*, 117.
19. Lewis, “The Man Who Crashed the World,” 115, 122.
20. Dwight M. Jaffee, “Catastrophe Insurance and Regulatory Reform after the Subprime Mortgage Crisis,” in Erwann Micel-Kerjan and Paul Slovic, eds., *The Irrational Economist: Making Decisions in a Dangerous World* (New York: PublicAffairs, 2010), 162.
21. Nate Silver, *The Signal and the Noise: Why So Many Predictions Fail—But Some Don’t* (New York: Penguin Press, 2012), 355–356.
22. This issue is discussed at fascinating length in Donald C. Langevoort, “Chasing the Greased Pig Down Wall Street: A Gatekeeper’s Guide to the Psychology, Culture, and Ethics of Financial Risk Taking,” *Cornell Law Review*, Vol. 96, No. 5 (July 2011), 1224–1234.
23. Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long Term Capital Management* (New York: Random House, 2000), 130.
24. John Cassidy, *How Markets Fail: The Logic of Economic Calamities* (New York: Farrar, Straus and Giroux, 2009), 11–13, 177.
25. Joe Nocera, “Risk Management,” *New York Times*, January 2, 2009.
26. Stephen C. Nelson and Peter J. Katzenstein, “Uncertainty, Risk and the Financial Crisis of 2008,” *International Organization*, Vol. 68, No. 2 (March 2014), 373–374.
27. Lowenstein, *When Genius Failed*, 17, 21, 89.
28. Philippe Jorion, “Risk Management Lessons from LTCM,” *European Financial Management* Vol. 6 (September 2000), 277–300.

29. Langevoort, "Chasing the Greased Pig down Wall Street," 1219–1220.
30. Farrell, *Crash of the Titans*, 23.
31. For examples from the Merrill case, see Farrell, *Crash of the Titans*, 27, 122.
32. William D. Cohan, *House of Cards: A Tale of Hubris and Wretched Excess on Wall Street* (New York: Doubleday, 2009), 266–267.
33. Lowenstein, *When Genius Failed*, 58, 129.
34. Farrell, *Crash of the Titans*, 73, 88–89, 478.
35. Sorkin, *Too Big to Fail*, 186, 160; Morgenson and Rosner, *Reckless Endangerment*, 10, 18, 22–30, 194.
36. McLean and Elkind, *The Smartest Guys in the Room*, 121, 152, 213, 216–217.
37. Farrell, *Crash of the Titans*, 154.
38. Morgenson and Rosner, *Reckless Endangerment*, 265, 282–283.
39. Diane Vaughan, *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA* (Chicago: University of Chicago Press, 1996), 409.
40. Vaughan, *The Challenger Launch Decision*, 37, 400.
41. McLean and Elkind, *The Smartest Guys in the Room*, 132.
42. Morgenson and Rosner, *Reckless Endangerment*, 93.
43. Nassim Nicholas Taleb, *The Black Swan: The Impact of the Highly Improbable* (New York: Random House, 2007), 275.
44. Clark Murdock, Becca S. Smith, Matt Squeri, and Chris Jones, *Risk Management in Non-DoD U.S. Government Agencies and the International Community* (Washington, DC: Center for Strategic and International Studies, 2011), 12.

11 The Role of Risk in Strategy

1. Frank Camm, Lauren Caston, Alexander C. Hou, Forrest E. Morgan, and Alan Vick, *Managing Risk in USAF Force Planning* (Santa Monica, CA: The RAND Corporation, 2009), 4.
2. As suggested in the framework above, this concept of risk is very similar to some aspects of what has become known as "strategic risk." One Deloitte survey defines this as "risks that affect or are created by an organization's business strategy and strategic objectives." See Deloitte Consulting, "Exploring Strategic Risk: A Global Survey," 2013, 4; available at http://deloitte.wsj.com/riskandcompliance/files/2013/10/strategic_risk_survey.pdf.
3. 10 U.S.C. §118; available at <http://www.gpo.gov/fdsys/pkg/USCODE-2011-title10/pdf/USCODE-2011-title10-subtitleA-partI-chap2-sec118.pdf>.
4. This insight comes out in Camm et al., *Managing Risk in USAF Force Planning*, 80–100.
5. See, for example, Douglas W. Hubbard, *The Failure of Risk Management: Why It's Broken and How to Fix It* (Hoboken, NJ: Wiley, 2009), 117–143.
6. Hubbard, *The Failure of Risk Management*, 122.

12 Outcome Assessment of the Emerging US National Security Strategy

1. One of the best such treatments is Barry R. Posen and Andrew S. Ross, "Competing Visions for U.S. Grand Strategy," *International Security*, Vol. 21, No. 3 (Winter 1996–1997), 5–53.
2. See Thomas Wright, "The Rise and Fall of the Unipolar Concert," *The Washington Quarterly*, Vol. 37, No. 4 (Winter 2015), 7–24.

3. See, for example, Robert Art, "Geopolitics Updated: The Strategy of Selective Engagement," *International Security*, Vol. 23, No. 3 (Winter 1998–1999), 5–42.
4. Michael J. Mazarr and the NDU Study Group, "Discriminate Power: A Strategy for a Sustainable National Security Posture," Foreign Policy Research Institute, *Philadelphia Papers*, No. 2, May 2013, 10.
5. Walter Russell Mead, remarks at Council on Foreign Relations conference, "Hegemony and the Decline of Empires: Is History Repeating Itself?" April 28, 2003; available at <http://www.cfr.org/history-and-theory-of-international-relations/hegemony-decline-empires-history-repeating-itself/p5940>.
6. Paul Kennedy, *The Rise and Fall of the Great Powers* (New York: Vintage Books, 1987), 151, 228.
7. David Lake, "British and American Hegemony Compared: Lessons for the Current Era of Decline," in Michael Fry, ed., *History, the White House and the Kremlin: Statesmen as Historians* (London: Pinter Publishers, 1991), 108, 111.
8. Kennedy, *The Rise and Fall of the Great Powers*, 226–227.
9. Kennedy, *The Rise and Fall of the Great Powers*, 230, 315.
10. Aaron Friedberg, *The Weary Titan* (Princeton, NJ: Princeton University Press, 1988), 103.
11. Andrew Gamble, "Hegemony and Decline: Britain and the United States," in Patrick Karl O'Brien and Armand Clesse, eds., *Two Hegemonies: Britain 1846–1914 and the United States 1941–2001* (Aldershot: Ashgate Publishing, 2002), 127–140.
12. Jean-Marie Guehenno, *The End of the Nation-State* (Minneapolis: University of Minnesota Press, 1985).
13. Kennedy, *The Rise and Fall of the Great Powers*, 232, 317.
14. Charles A. Kupchan, "Grand Strategy and Power Transitions: What We Can Learn from Great Britain," New America Foundation essay, July 2011, 3; available at <http://newamerica.net/sites/newamerica.net/files/policydocs/Kupchan.%20Grand%20Strat%20and%20Power%20Formatted%20PDF.pdf>.
15. Kennedy, *The Rise and Fall of the Great Powers*, 319.
16. Kupchan, "Grand Strategy and Power Transitions," 2.
17. For an interesting argument about identity-formation and the construction of a strong strategic relationship between the former enemies of London and Washington, see Feng Yongping, "The Peaceful Transition of Power from the UK to the U.S.," *Chinese Journal of International Politics*, Vol. 1, No. 1 (2006); available at <http://cjp.oxfordjournals.org/content/1/1/83.full>.
18. Eugene Gholz, "The Nixon Doctrine in the 21st Century," *World Politics Review*, July 22, 2009, available at <http://www.worldpoliticsreview.com/articles/4106/the-nixon-doctrine-in-the-21st-century>.
19. Chad Pillai, "The Return of Great Power Politics: Re-Examining the Nixon Doctrine," *War on the Rocks*, March 27, 2014; available at <http://warontherocks.com/2014/03/the-return-of-great-power-politics-re-examining-the-nixon-doctrine/>.
20. Henry A. Kissinger, *Diplomacy* (New York: Simon and Schuster, 1994), 707.
21. Kissinger, *Diplomacy*, 679–685, surveys various options considered in Vietnam.
22. Kissinger, *Diplomacy*, 707–708.
23. Richard Nixon, "Address to the Nation on the War in Vietnam," November 3, 1969; available at http://www.nixonlibrary.gov/forkids/speechesforkids/silentmajority/silentmajority_transcript.pdf.
24. Richard Nixon, *U.S. Foreign Policy for the 1970's, a New Strategy for Peace* (Washington, DC: U.S. Government Printing Office, February 18, 1970).
25. Kissinger, *Diplomacy*, 708–709; and Henry A. Kissinger, *White House Years* (Boston: Little, Brown, 1979), 224.

26. Richard M. Nixon, "Asia After Viet Nam," *Foreign Affairs*, October, 1967. See also Mark Safranski, "A New Nixon Doctrine: Strategy for a Polycentric World," *War on the Rocks*, July 8, 2014; available at <http://warontherocks.com/2014/07/a-new-nixon-doctrine-strategy-for-a-polycentric-world/>.
27. Earl C. Ravenal, "The Nixon Doctrine and Our Asian Commitments," *Foreign Affairs*, January 1971, available at <https://www.foreignaffairs.com/articles/asia/1971-01-01/nixon-doctrine-and-our-asian-commitments>.
28. Ravenal points out that the supposed "two-and-a-half war" planning construct, offering forces capable of a several-month defensive action in Europe and an Asian war and a smaller contingency, was never fulfilled. In practice the United States deployed forces capable of one-and-a-half wars. This approach was "explicitly founded on the improbability of two simultaneous major contingencies. Thus demands on the planned general purpose forces are to be considered alternative rather than additive." Ravenal, "The Nixon Doctrine and Our Asian Commitments."
29. For a number of recent treatments of the credibility problem, see Christopher Fettweis, "Credibility and the War on Terror," *Political Science Quarterly*, Vol. 122, No. 4 (Winter 2007), 607–633; Stephen M. Walt, "The Credibility Addiction," *Foreign Policy*, January 6, 2015, available at <http://foreignpolicy.com/2015/01/06/the-credibility-addiction-us-iraq-afghanistan-unwinnable-war/>; Jonathan Mercer, *Reputation and International Politics* (Ithaca: Cornell University Press, 2010); and Daryl Press, *Calculating Credibility: How Leaders Assess Military Threats* (Ithaca: Cornell University Press, 2007).
30. Ravenal, "The Nixon Doctrine and Our Asian Commitments."
31. Bill Sweetman, "Nuclear Bow Wave Builds with Program Starts," *Aviation Week*, February 2, 2015, available at <http://aviationweek.com/defense/nuclear-bow-wave-builds-program-starts>.
32. Lawrence J. Korb, Alex Rothman, and Max Hoffman, "Reforming Military Compensation," Center for American Progress, May 2012, available at <https://www.americanprogress.org/issues/security/report/2012/05/07/11573/reforming-military-compensation/>; and Defense Business Board, "Modernizing the Military Retirement System," 2011; available at http://dbb.defense.gov/Portals/35/Documents/Reports/2011/FY11-5_Modernizing_The_Military_Retirement_System_2011-7.pdf.
33. See the discussion in Mark Gunzinger, *Shaping America's Future Military: Toward a New Force Planning Construct* (Washington, DC: Center for Strategic and Budgetary Assessments, 2013).

13 Principles of Effective Risk Management

1. See <https://www.deloitte.com/us/riskprism>.
2. Robert S. Kaplan and Anette Mikes, "Managing Risks: A New Framework," *Harvard Business Review*, June 2012; available at <https://hbr.org/2012/06/managing-risks-a-new-framework>.
3. Jonathan Baron, *Thinking and Deciding*, 4th ed. (New York: Cambridge University Press, 2008), 199–228.
4. Philip Tetlock and Dan Gardner, *Superforecasting: The Art and Science of Prediction* (New York: Crown, 2015).
5. Kaplan and Mikes, "Managing Risks: A New Framework."
6. See, for example, Gretchen Morgenson and Joshua Rosner, *Reckless Endangerment: How Outsized Ambition, Greed and Corruption Led to Economic Armageddon* (New York: Times Books, 2011), 127.

7. Robert S. Kaplan and Anette Mikes, "Managing Risks: A New Framework," *Harvard Business Review*, June 2012; available at <https://hbr.org/2012/06/managing-risks-a-new-framework>.

14 Managing Uncertainty

1. Sanyin Siang, "Gen. Martin Dempsey, Chairman of the Joint Chiefs, on the Networked World," *Fortune*, May 22, 2015.
2. Gen. Gordon R. Sullivan (Ret.), "Unprecedented Uncertainty Is Army's Biggest Challenge," *The Hill*, April 10, 2015.
3. U.S. Department of Defense, *Quadrennial Defense Review 2014* (Washington, DC: DoD, March 2014), 3.
4. Quoted in Michel Syrett and Marion Devine, *Managing Uncertainty: Strategies for Surviving and Thriving in Turbulent Times* (London: Economist Books, 2012), 3.
5. A leading example is Syrett and Devine, *Managing Uncertainty*. Nassim Nicholas Taleb approaches the challenge from a very different and useful perspective in *Antifragile: Things That Gain from Disorder* (New York: Random House, 2012). See also Yvan Allaire and Mibaela E. Firsirotu, "Coping with Strategic Uncertainty," *Sloan Management Review*, Vol. 7 (Spring 1989), 7–16.
6. This concept is mentioned in Ian Goldin and Mike Mariathan, *The Butterfly Defect: How Globalization Creates Systemic Risks, and What to Do About It* (Princeton, NJ: Princeton University Press, 2014), 54–55.
7. Richard Danzig, *Driving in the Dark: Ten Propositions about Prediction and National Security* (Washington, DC: Center for a New American Security, October 2011), 14–15. For a discussion of this distinction in relation to uncertainty-based policies, see Taleb, *Antifragile*, 4, 7, 9, 132–133.
8. See Hugh Courtney, Jane Kirkland, and Patrick Viguier, "Strategy Under Uncertainty," *Harvard Business Review*, November–December 1997, available at <https://hbr.org/1997/11/strategy-under-uncertainty>.
9. Stephen Ward and Chris Chapman, "Transforming Project Risk Management into Project Uncertainty Management," *International Journal of Project Management*, Vol. 21 (2003), 98–99. See also Silvio Funtowicz and Jerome Ravetz, "Uncertainty, Complexity and Post-Normal Science," *Environmental Toxicology and Chemistry*, Vol. 13, No. 12 (1994), 1881–1885, and T. J. Ross, J. M. Booker, and A. C. Montoya, "New Developments in Uncertainty Assessment and Uncertainty Management," *Expert Systems with Applications*, Vol. 40 (2013), 841–974.
10. Ward and Chapman, "Transforming Project Risk Management," 101.
11. Some sources indicate a more quantitative approach to "managing uncertainty," built around assigning numerical values to as many uncertain variables as possible. See the discussion in J. M. Booker and T. J. Ross, "An Evolution of Uncertainty Assessment and Quantification," *Scientia Iranica*, Vol. 18, No. 3 (2011), 669–676.
12. An example of such thinking is Paul K. Davis, "Uncertainty-Sensitive Planning," in Stuart E. Johnson, Martin Libicki, and Gregory F. Treverton, eds., *New Challenges, New Tools for Defense Decisionmaking* (Santa Monica, CA: The RAND Corporation, 2003), esp. 134–138.
13. Michael Fitzsimmons makes an argument for the continuing value of probabilistic situational analysis in national security planning in "The Problem of Uncertainty in Strategic Planning," *Survival*, Vol. 48, No. 4 (2006), 131–146.
14. This mindset is closely related to what Jonathan Baron has described as "actively open-minded thinking," which is a disposition that remains open to alternative

- explanations and hypotheses that differ from one's favor explanation. See Baron, *Thinking and Deciding*, 4th ed. (New York: Cambridge University Press, 2008), 199–228.
15. See Central Intelligence Agency, "A Tradecraft Primer: Structured Analytic Techniques for Improving Intelligence Analysis," March 2009, 14–15; available at <https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/Tradecraft%20Primer-apr09.pdf>.
 16. I am grateful to Professor Martin Krieger of USC for bringing to my attention this concept of "probing" the environment as a critical analytical habit under uncertainty.
 17. As much as I find Fitzgerald and Cheney-Peters's analysis wonderfully persuasive, I do not prefer the term "hedging" as a shorthand for a US defense strategy. Its synonyms are words like prevarication, equivocation, and fudging, and that is not the spirit behind their approach—or mine. What they have in mind is an active search for competitive advantage through bold and creative investment and experimentation in a wide range of capabilities. I think the term "adaptive strategy" better hints at this concept.
 18. Ben Fitzgerald and Scott Cheney-Peters, "How the Military Can Keep Its Edge: Don't Offset—Hedge," *War on the Rocks*, April 29, 2015.
 19. Taleb, *Antifragile*, 3–4, 10–11.
 20. Andrew G. Haldane, "The Dog and the Frisbee," Speech at Federal Reserve Bank of Kansas City's 366th Economic Policy Symposium, August 31, 2012.
 21. An excellent source is John Kambhu, Scott Weidman, and Neel Krishnan, Rapporteurs, *New Directions for Understanding Systemic Risk: A Report on a Conference Cosponsored by the Federal Reserve Bank of New York and the National Academy of Sciences* (Washington, DC: National Academies Press, 2007), Chapter 3, "Systemic Risk in Ecology and Engineering." Syrett and Devine argue for such principles as anticipation, agility, and resilience; see *Managing Uncertainty*, Chapters 2, 4 and 5.
 22. Taleb, *Antifragile*, 332.
 23. On this principle see Danzig, "Driving in the Dark," 26–28.
 24. Taleb, *Antifragile*, 44–45.
 25. Danzig, "Driving in the Dark," 19–21.
 26. Kambhu et al., *New Directions for Understanding Systemic Risk*, 31.
 27. Danzig, "Driving in the Dark," 21–23.
 28. Taleb stresses this in *Antifragile*, 102.
 29. Taleb, *Antifragile*, 100, 119.
 30. James G. March, "Rationality, Foolishness, and Adaptive Intelligence," *Strategic Management Journal*, Vol. 27 (2006), 201–214.
 31. Henry Mintzberg, "Crafting Strategy," *Harvard Business Review*, July 1987, 66.
 32. Mintzberg, "Crafting Strategy," 73.
 33. Frank Hoffman, "Black Swans and Pink Flamingos: Five Principles for Force Design," *War on the Rocks*, September 3, 2015.
 34. Danzig, "Driving in the Dark," 23–25.
 35. Taleb, *Antifragile*, 122, 302, 338.
 36. Joshua Cooper Ramo, *The Age of the Unthinkable* (New York: Little, Brown, 2009), 61.
 37. Mary Douglas and Aaron Wildavsky, *Risk and Culture: An Essay on the Selection of Technological and Environmental Dangers* (Berkeley: University of California Press, 1982), 196.
 38. Douglas and Wildavsky, *Risk and Culture*, 198.
 39. Myron S. Scholes and Tom Kimner, "Evolving from Quantitative Risk Management to a High-Performance Risk Management Analytic Framework: Insights on a New Direction for Risk Management," SAS Institute White Paper, n.d., 2.

40. For a recent analysis of the requirements for such a process see Edie Williams and Alan R. Shaffer, "The Defense Innovation Initiative: The Importance of Capability Prototyping," *Joint Force Quarterly*, No. 77 (Second Quarter 2015), available at <http://ndupress.ndu.edu/Media/News/NewsArticleView/tabid/7849/Article/581867/jfq-77-the-defense-innovation-initiative-the-importance-of-capability-prototypi.aspx>.
41. Fitzgerald and Cheney-Peters, "How the Military Can Keep Its Edge."
42. Fitzgerald and Cheney-Peters, "How the Military Can Keep Its Edge."

Bibliography

- Abolafia, Mitchel Y. 2012. "Central Banking and the Triumph of Technical Rationality." In Katrin Knorr Cetina and Alex Preda (eds.), *Handbook of the Sociology of Finance*. New York: Oxford University Press, pp. 94–114.
- . "Narrative Construction as Sensemaking: How a Central Bank Thinks." *Organization Studies*, Vol. 31, No. 3 (2010): 349–367.
- . "Framing Moves: Interpretive Politics at the Federal Reserve." *Journal of Public Administration Research and Theory*, Vol. 14, No. 3 (2004): 349–370.
- Adams, John. *Risk*. London: Routledge, 1995.
- . "Egg on Wall? It's Rational." *Times Higher Education Supplement*, October 6, 2000.
- Akerloff, George A., and Robert J. Shiller. *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*. Princeton, NJ: Princeton University Press, 2009.
- Allaire, Yvan, and Mibaela E. Firsirotu. "Coping with Strategic Uncertainty." *Sloan Management Review*, Vol. 7 (Spring 1989).
- Al-Najjar, Nabil I., and Luciano de Castro. "Subjective Probability." Unpublished m.s., Kellogg School of Business, Northwestern University, March 2010.
- Art, Robert. "Geopolitics Updated: The Strategy of Selective Engagement." *International Security*, Vol. 23, No. 3 (Winter 1998–1999).
- Assmuth, Timo, and Adam M. Finkel. "Choices and Rationalities Under Radical Uncertainty: Ideals and Principles Behind Responses to Risks and Risk Information." Unpublished paper, n.d.
- Barnes, James H. "Cognitive Biases and Their Impact on Strategic Planning." *Strategic Management Journal*, Vol. 5, No. 2 (April–June 1984).
- Baron, Jonathan Baron. "Nonconsequentialist Decisions." *Behavioral and Brain Sciences*, Vol. 17, No. 1 (1994): 1–10.
- Baron, Jonathan Baron, and Mark Spranca. "Protected Values." *Organizational Behavior and Human Decision Processes*, Vol. 70, No. 1 (April 1997).
- . *Thinking and Deciding*, 4th Ed. New York: Cambridge University Press, 2008.
- Basili, Marcello, and Carlo Zappia, "Shackle and Modern Decision Theory." Unpublished paper, April 2006.
- Ben-Haim, Yakov. *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*. London: Elsevier, 2001.
- Bernanke, Ben. "Modern Risk Management and Banking Supervision." Speech at the Stonier Graduate School of Banking, June 12, 2006.
- Bernstein, Paul L. *Against the Gods: The Remarkable Story of Risk*. New York: John Wiley and Sons, 1998.
- Beschloss, Michael, ed. *Reaching for Glory: Lyndon Johnson's Secret White House Tapes, 1964–1965*. New York: Simon and Schuster, 2001.
- Bixter, Michael T., and Christian C. Luhmann. "Evidence for Implicit Risk: Delay Facilitates the Processing of Uncertainty." *Journal of Behavioral Decision Making*, Vol. 28 (2015).
- Blustein, Paul. *The Chastening: Inside the Crisis that Rocked the Global Financial System and Humbled the IMF*. New York: PublicAffairs, 2001.
- Blyth, Mark. "Great Punctuations: Prediction, Randomness, and the Evolution of Comparative Political Science." *American Political Science Review*, Vol. 100, No. 4 (2006): 493–498.

- . “Coping with the Black Swan: The Unsettling World of Nassim Taleb.” *Critical Review*, Vol. 21, No. 4 (2009).
- . “Ideas, Uncertainty and Evolution.” In Robert Cox and Daniel Beland (eds.), *Ideas and Politics in Social Science Research*. New York: Oxford University Press, 2010, pp. 83–101.
- Bond, Alan, A. Morrison-Saunders, J. A. Gunn, J. Pope, and F. Retief. “Managing Uncertainty, Ambiguity and Ignorance in Impact Assessment by Embedding Evolutionary Resilience, Participatory Modelling and Adaptive Management.” *Journal of Environmental Management*, Vol. 151 (2015).
- Booker, J. M., and T. J. Ross. “An Evolution of Uncertainty Assessment and Quantification.” *Scientia Iranica*, Vol. 18, No. 3 (2011), 669–676.
- Brown, Christopher, and Cheng Hao. “Treating Uncertainty as Risk: The Credit Default Swap and the Paradox of Derivatives.” *Journal of Economic Issues*, Vol. XLVI, No. 2 (June 2012).
- Burton, Robert A. *On Being Certain: Believing You Are Right Even When You’re Not*. New York: St. Martin’s Griffin, 2008.
- Byman, Daniel L., and Kenneth M. Pollack. “Let Us Now Praise Great Men: Bringing the Statesman Back In.” *International Security*, Vol. 25, No. 4 (Spring 2001), 107–146.
- Camm, Frank, Lauren Caston, Alexander C. Hou, Forrest E. Morgan, and Alan Vick. *Managing Risk in USAF Force Planning*. Santa Monica, CA: The RAND Corporation, 2009.
- Carter, Graydon, ed. *The Great Hangover*. New York: HarperPerennial, 2010.
- Carvalho, Fernando J. Cardim. “Keynes on Probability, Uncertainty, and Decision Making.” *Journal of Post-Keynesian Economics*, Vol. XI, No. 1 (Fall 1988).
- Cassidy, John. *How Markets Fail: The Logic of Economic Calamities*. New York: Farrar, Straus and Giroux, 2009.
- Chen, Serena. “Getting at the Truth or Getting Along: Accuracy- Versus Impression-Motivated Heuristic and Systemic Processing.” *Journal of Personality and Social Psychology*, Vol. 71, No. 2 (1996).
- Choi, Young Bak. *Paradigms and Conventions: Uncertainty, Decision Making and Entrepreneurship*. Ann Arbor: University of Michigan Press, 1993.
- Clarke, Richard. *Against All Enemies: Inside America’s War on Terror*. New York: Free Press, 2004.
- Cohan, William D. *House of Cards: A Tale of Hubris and Wretched Excess on Wall Street*. New York: Doubleday, 2009.
- Coleman, Thomas S. *A Practical Guide to Measuring Risk*. New York: Research Foundation of CFA Institute, 2011.
- Committee of Sponsoring Organizations of the Treadway Commission. “Enterprise Risk Management—Integrated Framework.” September 2004.
- Crouhy, Michael, Dan Galai, and Robert Mark. *The Essentials of Risk Management*. New York: McGraw Hill, 2005.
- Courtney, Hugh, Jane Kirkland, and Patrick Viguerie. “Strategy Under Uncertainty.” *Harvard Business Review*, November–December 1997.
- Dallek, Robert. *Flawed Giant: Lyndon Johnson and His Times, 1961–1973*. New York: Oxford University Press, 1988.
- Danielsson, Jon, and Hyung Son Shin. “Endogenous Risk.” Unpublished paper, September 21, 2002.
- Danzig, Richard. *Driving in the Dark: Ten Propositions About Prediction and National Security*. Washington, DC: Center for a New American Security, October 2011.
- Davidson, Paul. “Risk and Uncertainty in Economics.” Unpublished conference paper, February 6, 2009.

- . “Is Risk Management a Science?” Unpublished conference paper, May 15, 2012.
- Davis, Paul K. “Uncertainty-Sensitive Planning.” In Stuart E. Johnson, Martin Libicki, and Gregory F. Treverton (eds.), *New Challenges, New Tools for Defense Decisionmaking*. Santa Monica, CA: The RAND Corporation, 2003.
- Defense Systems Management College. *Risk Management Guide for DoD Acquisition*. Ft. Belvoir, VA: Defense Systems Management College Press, May 1999.
- Deloitte & Touche LLP. “Global Risk Management Survey: Sixth Edition: Risk Management in the Spotlight.” June 2009.
- . “Nine Principles of a Risk Intelligent Framework for Risk Management—A Briefing for Federal Executives.” 2012.
- . “Exploring Strategic Risk: A Global Survey.” 2013.
- Dillon-Merrill, Robin, Catherine H. Tinsley, and Matthew A. Cronin. “How Near-Miss Events Amplify or Attenuate Risky Decision Making.” *Published Articles & Papers*, Paper 93 (2012).
- Douglas, Mary, and Aaron Wildavsky. *Risk and Culture: An Essay on the Selection of Technological and Environmental Dangers*. Berkeley: University of California Press, 1982.
- Dowie, Jack. “Against Risk.” *Risk Decision and Policy* Vol. 4, No. 1 (1999), 57–73.
- . “Communication for Better Decisions: Not About ‘Risk’.” *Health, Risk and Society*, Vol. 1, No. 1 (1999).
- . “A Risky Decision: Managing Without Risk.” *Risk Management*, Vol. 2, No. 2 (2000).
- Earl, Peter. *The Economic Imagination: Towards a Behavioural Analysis of Choice*. Armonk, NY: M. E. Sharpe, 1983.
- Ellsberg, Daniel. “Risk, Ambiguity, and the Savage Axioms.” *The Quarterly Journal of Economics*, Vol. 75, No. 4 (November 1961).
- Farrell, Greg. *Crash of the Titans: Greed, Hubris, the Fall of Merrill Lynch, and the Near-Collapse of Bank of America*. New York: Crown Business, 2010.
- Fettweis, Christopher. “Credibility and the War on Terror.” *Political Science Quarterly*, Vol. 122, No. 4 (Winter 2007).
- Figes, Orlando. *The Crimean War: A History*. New York: Metropolitan Books, 2010.
- Fischer, Gregory W., Ziv Carmon, Dan Ariely, and Gal Zauberman. “Goal-based Construction of Preferences: Task Goals and the Prominence Effect.” *Management Science*, Vol. 45, No. 8 (August 1999), 1057–1075.
- Fischhoff, Baruch, and John Kadvaný. *Risk: A Very Short Introduction*. London: Oxford University Press, 2011.
- Fitzgerald, Ben, and Scott Cheney-Peters. “How the Military Can Keep Its Edge: Don’t Offset—Hedge.” *War on the Rocks*, April 29, 2015.
- Fitzsimmons, Michael. “The Problem of Uncertainty in Strategic Planning.” *Survival*, Vol. 48, No. 4 (2006), 131–146.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue. “Time Discounting and Time Preference: A Critical Review.” *Journal of Economic Literature*, Vol. 40, No. 2 (June 2002).
- Frey, Bruno S., and Alois Stutzer. “Beyond Outcomes: Measuring Procedural Utility.” *Oxford Economic Papers*, no. 57 (2005), 90–111.
- Friedberg, Aaron. *The Weary Titan*. Princeton, NJ: Princeton University Press, 1988.
- Funtowicz, Silvio, and Jerome Ravetz. “Uncertainty, Complexity and Post-Normal Science.” *Environmental Toxicology and Chemistry*, Vol. 13, No. 12 (1994), 1881–1885.
- Gaidow, Svetoslav, and Seng Boey. *Australian Defence Risk Management Framework: A Comparative Study*. Canberra: Australian Government, Department of Defence, Systems Sciences Laboratory, 2005.

- Gamble, Andrew. "Hegemony and Decline: Britain and the United States." In Patrick Karl O'Brien and Armand Clesse (eds.), *Two Hegemonies: Britain 1846–1914 and the United States 1941–2001*. Aldershot: Ashgate Publishing, Ltd., 2002.
- George, Alexander L. *Presidential Decisionmaking in Foreign Policy: The Effective Use of Information and Advice*. Boulder, CO: Westview Press, 1980.
- Gholz, Eugene. "The Nixon Doctrine in the 21st Century." *World Politics Review*, July 22, 2009.
- Gigerenzer, Gerd. *Gut Feelings: The Intelligence of the Unconscious*. New York: Penguin Books, 2008.
- Goldin, Ian, and Mike Mariathan. *The Butterfly Defect: How Globalization Creates Systemic Risks, and What to Do About It*. Princeton, NJ: Princeton University Press, 2014.
- Goldstein, Gordon M. *Lessons in Disaster: McGeorge Bundy and the Path to War in Vietnam*. New York: Times Books, 2008.
- Greenspan, Alan. *The Map and the Territory: Risk, Human Nature, and the Future of Forecasting*. New York: Penguin, 2013.
- Guehenno, Jean-Marie. *The End of the Nation-State*. Minneapolis: University of Minnesota Press, 1985.
- Güney, Şule, and Ben R. Newell. "Overcoming Ambiguity Aversion Through Experience." *Journal of Behavioral Decision Making*, Vol. 28 (2015).
- Gunzinger, Mark. *Shaping America's Future Military: Toward a New Force Planning Construct*. Washington, DC: Center for Strategic and Budgetary Assessments, 2013.
- Haldane, Andrew G. "The Dog and the Frisbee." Speech at Federal Reserve Bank of Kansas City's 366th Economic Policy Symposium, August 31, 2012.
- Hammond, Kenneth R., Thomas R. Stewart, Berndt Brehmer, and Derick O. Steinman. "Social Judgment Theory." In Martin Kaplan and Steven Schwartz (eds.), *Human Judgment and Decision Processes*. New York: Academic Press, 1986.
- Henderson, John C., and Paul C. Nutt. "The Influence of Decision Style on Decision Making Behavior." *Management Science*, Vol. 26, No. 4 (April 1980), 387–400.
- Hoffman, Frank. "Black Swans and Pink Flamingos: Five Principles for Force Design." *War on the Rocks*, September 3, 2015.
- Hubbard, Douglas W. *The Failure of Risk Management: Why It's Broken and How to Fix It*. Hoboken, NJ: John Wiley and Sons, 2009.
- Jorion, Philippe. "Risk Management Lessons from Long Term Capital Management." *European Financial Management*, Vol. 6 (September 2000).
- Kahneman, Daniel. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.
- Kahneman, Daniel, and Amos Tversky. "Subjective Probability: A Judgment of Representativeness." *Cognitive Psychology*, Vol. 3 (1972).
- . "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica*, Vol. 47, No. 2 (March 1979).
- Kambhu, John, Scott Weidman, and Neel Krishnan, Rapporteurs. *New Directions for Understanding Systemic Risk: A Report on a Conference Cosponsored by the Federal Reserve Bank of New York and the National Academy of Sciences*. Washington, DC: National Academies Press, 2007.
- Kaplan, Robert, and Anette Meeks. "JP Morgan's Loss: Bigger than 'Risk Management'." *HBR Blog post*, May 23, 2012.
- . "Managing Risks: A New Framework." *Harvard Business Review*, June 2012.
- Kashyap, Anil K. "Lessons of the Financial Crisis for Risk Management." Paper Prepared for the Financial Crisis Inquiry Commission, February 27, 2010.
- Kennedy, Paul. *The Rise and Fall of the Great Powers*. New York: Vintage Books, 1987.
- Kissinger, Henry. *White House Years*. Boston: Little, Brown, 1979.
- . *Diplomacy*. New York: Simon and Schuster, 1994.

- Knight, Frank H. *Risk, Uncertainty and Profit*, Reprint of 3rd Ed. New York: Augustus Kelley, 1964.
- Kunda, Ziva. "The Case for Motivated Reasoning." *Psychological Bulletin*, Vol. 108, No. 3 (1990).
- Kunreuther, Howard, et al. "High-Stakes Decision Making: Normative, Descriptive and Prescriptive Considerations." Unpublished conference summary paper, October 8, 2001.
- Kupchan, Charles A. "Grand Strategy and Power Transitions: What We Can Learn from Great Britain." New America Foundation essay, July 2011.
- Lafrance, Adrienne. "Disasters that Were Foretold." *The Atlantic*, September 4, 2015.
- Lake, David. "British and American Hegemony Compared: Lessons for the Current Era of Decline." In Michael Fry (ed.), *History, the White House and the Kremlin: Statesmen as Historians*. London: Pinter Publishers, 1991.
- Lipshitz, Raanan. "The Road to Desert Storm." *Organization Studies* Vol. 16, No. 2 (1995).
- Lipshitz, Raanan, and Orna Strauss. "Coping with Uncertainty: A Naturalistic Decision-Making Analysis." *Organizational Behavior and Human Decision Processes*, Vol. 69, No. 2 (February 1997), 149–163.
- Loewenstein, George. "Out of Control: Visceral Influences on Behavior." *Organizational Behavior and Human Decision Processes*, Vol. 65, No. 3 (March 1996), 272–292.
- Loewenstein, George, Elke U. Webber, Christopher K. Hsee, and Ned Welch. "Risk as Feelings." *Psychological Bulletin*, Vol. 127, No. 2 (2001), 267–286.
- Loewenstein, George, Daniel Read, and Roy F. Baumeister, eds. *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*. New York: Sage Foundation, 2003.
- Lowenstein, Roger. *When Genius Failed: The Rise and Fall of Long Term Capital Management*. New York: Random House, 2000.
- Mackay, Charles. *Extraordinary Popular Delusions and the Madness of Crowds*. Renaissance Classics, 2012.
- Mandelbrot, Benoit, and Richard L. Hudson. *The (Mis)Behavior of Markets: A Fractal View of Financial Turbulence*. New York: Basic Books, 2004.
- March, James G. "Bounded Rationality, Ambiguity, and the Engineering of Choice." *The Bell Journal of Economics*, Vol. 9, No. 2 (Autumn 1978).
- _____. "Rationality, Foolishness, and Adaptive Intelligence." *Strategic Management Journal*, Vol. 27 (2006), 201–214.
- March, James G., and Zur Shapira. "Managerial Perspectives on Risk and Risk Taking." *Management Science*, Vol. 33, No. 11 (November 1987).
- Mazarr, Michael J., and the NDU Study Group. "Discriminate Power: A Strategy for a Sustainable National Security Posture." Foreign Policy Research Institute *Philadelphia Papers* No. 2, May 2013.
- McLean, Bethany, and Peter Elkind. *The Smartest Guys in the Room: The Amazing Rise and Scandalous Fall of Enron*. New York: Penguin, 2003.
- Mead, Walter Russell. "Hegemony and the Decline of Empires: Is History Repeating Itself?" Remarks at Council on Foreign Relations conference, April 28, 2003.
- Michale-Kerjan, Erwann, and Paul Slovic, eds. *The Irrational Economist: Making Decisions in a Dangerous World*. New York: PublicAffairs Books, 2010.
- Mercer, Jonathan. *Reputation and International Politics*. Ithaca: Cornell University Press, 2010.
- Mlodinow, Leonard. *The Drunkard's Walk: How Randomness Rules Our Lives*. New York: Pantheon Books, 2008.
- Morgenson, Gretchen, and Joshua Rosner. *Reckless Endangerment: How Outsized Ambition, Greed and Corruption Led to Economic Armageddon*. New York: Times Books, 2011.

- Murdock, Clark, Matt Squeri, Chris Jones, and Becca S. Smith. *Risk Management in Non-DoD U.S. Government Agencies and the International Community*. Washington, DC: Center for Strategic and International Studies, 2011.
- National Aeronautics and Space Administration. *NASA Risk Management Handbook*. Washington, DC: NASA Headquarters, November 2011.
- Nelson, Stephen C. "Governing Risky and Uncertain Financial Markets." *Stiftung Wissenschaft und Politik: SWP Comments* No. 45, December 2012.
- Nelson, Stephen C., and Peter J. Katzenstein. "Uncertainty, Risk and the Financial Crisis of 2008." *International Organization*, Vol. 68, No. 2 (March 2014).
- Nocera, Joe. "Risk Management." *New York Times*, January 2, 2009.
- Nutt, Paul C. "Decision Style and Strategic Decisions of Top Executives." *Technological Forecasting and Social Change*, Vol. 30, No. 1 (August 1986), 39–62.
- Patterson, Scott. *The Quants: How a New Breed of Math Whizzes Conquered Wall Street and Nearly Destroyed It*. New York: Crown Business, 2010.
- Perry, William, and John P. Abizaid, co-chairs. *Ensuring a Strong U.S. Defense for the Future: The National Defense Panel Review of the 2014 Quadrennial Defense Review*. Washington, DC: U.S. Institute of Peace, July 31, 2014.
- Pillai, Chad. "The Return of Great Power Politics: Re-Examining the Nixon Doctrine." *War on the Rocks*, March 27, 2014.
- Plous, Scott. *The Psychology of Judgment and Decision Making*. New York: McGraw-Hill, 1993.
- Posen, Barry, and Andrew S. Ross. "Competing Visions for U.S. Grand Strategy." *International Security*, Vol. 21, No. 3 (Winter 1996–1997).
- Prasad, Monica, Andrew J. Perrin, Kieran Bezila, Steve G. Hoffman, Kate Kindleberger, Kim Maturuk, and Ashleigh Smith Powers, "'There Must Be a Reason': Osama, Saddam, and Inferred Justification." *Sociological Inquiry*, Vol. 79, No. 2 (2009).
- Press, Daryl. *Calculating Credibility: How Leaders Assess Military Threats*. Ithaca: Cornell University Press, 2007.
- Pretz, Jean E. et al. "Development and Validation of a New Measure of Intuition: The Types of Intuition Scale." *Journal of Behavioral Decision Making*, Vol. 27 (2014).
- Protiviti Inc. "Ten Common Risk Management Failures and How to Avoid Them." *The Bulletin*, Vol. 3, No. 6 (2008).
- PwC (PriceWaterHouse Cooper). "Black Swans Turn Grey: The Transformation of Risk." January 2012.
- Ramo, Joshua Cooper. *The Age of the Unthinkable*. New York: Little, Brown, 2009.
- Rathbun, Brian C. "Uncertain about Uncertainty: Understanding the Multiple Meanings of a Crucial Concept in International Relations Theory." *International Studies Quarterly*, Vol. 51 (2007), 533–557.
- Ravenal, Earl C. "The Nixon Doctrine and Our Asian Commitments." *Foreign Affairs*, January 1971.
- Redlawsk, David P. "Hot Cognition or Cool Consideration? Testing the Effects of Motivated Reasoning on Political Decision Making." *The Journal of Politics*, Vol. 64, No. 4 (November 2002).
- Rosenzweig, Phil. "What Makes Strategic Decisions Different." *Harvard Business Review*, November 2013.
- Ross, T. J., J. M. Booker, and A. C. Montoya. "New Developments in Uncertainty Assessment and Uncertainty Management." *Expert Systems with Applications*, Vol. 40 (2013), 841–974.
- Rossi, Cliff. *A Risk Professional's Survival Guide: Applied Best Practices in Risk Management*. New York: Wiley Finance, 2014.
- Salmon, Felix. "Recipe for Disaster: The Formula That Killed Wall Street." *Wired*, February 23, 2009.

- Schacter, Barry. "The New Risk Management Framework After the 2008 Financial Crisis." Presentation to Global Association of Risk Professionals. July 13, 2010. Available at <http://www.slideshare.net/BarrySchachter/the-new-risk-management-framework-after-the-2008-financial-crisis>.
- Scholes, Myron S., and Tom Kimner. "Evolving from Quantitative Risk Management to a High-Performance Risk Management Analytic Framework: Insights on a New Direction for Risk Management." SAS Institute White Paper, n.d.
- Schwenk, Charles R. "Cognitive Simplification Processes in Strategic Decision Making." *Strategic Management Journal*, Vol. 5, No. 2 (April–June 1984).
- Scott-Kakures, Dion. "Motivated Believing: Wishful and Unwelcome." *Noûs*, Vol. 34, No. 3 (2000).
- Shafir, Eldar, and Amos Tversky. "Thinking Through Uncertainty: Nonconsequential Reasoning and Choice." *Cognitive Psychology* Vol. 24 (1992).
- Shiller, Robert. *Irrational Exuberance*, 2nd Ed. Princeton, NJ: Princeton University Press, 2005.
- Sikka, Prem. "Financial Crisis and the Silence of the Auditors." *Accounting, Organizations and Society* (2009).
- Shackle, G. L. S. *Uncertainty in Economics*. Cambridge, UK: Cambridge University Press, 1968.
- . *Epistemics and Economics*. New Brunswick, NJ: Transaction Publishers, 2009.
- Silver, Nate. *The Signal and the Noise: Why So Many Predictions Fail—But Some Don't*. New York: Penguin Press, 2012.
- Simmons, Joseph P., and Leif D. Nelson. "Intuitive Confidence and the Prominence Effect"; and Simmons and Nelson, "Intuitive Confidence: Choosing Between Intuitive and Nonintuitive Alternatives." *Journal of Experimental Psychology: General*, Vol. 135, No. 3 (August 2006), 409–428.
- Simon, Mark, Susan M. Houghton, and Karl Aquino. "Cognitive Biases, Risk Perception, and Venture Formation: How Individuals Decide to Start Companies." *Journal of Business Venturing*, Vol. 15 (1999).
- Slovic, Paul. "Choice Between Equally Valued Alternatives." *Journal of Experimental Psychology: Human Perception and Performance*, Vol. 1, No. 3 (1975), 280–287.
- . "The Construction of Preference." *American Psychologist*, Vol. 50, No. 5 (1995).
- , ed. *The Perception of Risk*. London: Earthscan Publications, 2000.
- Slovic, Paul, Melissa L. Finucane, Ellen Peters, and Donald G. MacGregor. "Risk as Analysis and Risk as Feelings: Some Thoughts About Affect, Feelings, Risk, and Rationality." *Risk Analysis*, Vol. 24, No. 2 (2004).
- Soros, George. "A General Theory of Reflexivity." *The Financial Times*, October 26, 2009.
- Steinbruner, John D. *The Cybernetic Theory of Decision: New Dimensions of Political Analysis*. Princeton, NJ: Princeton University Press, 1974.
- Stirling, Andy. "Keep It Complex." *Nature*, Vol. 468 (December 2010), 1029–1031.
- Stulz, René M. "Risk Management Failures: What Are They and When Do They Happen?" Cornerstone Research Organization report, 2009.
- Syrett, Michel, and Marion Devine. *Managing Uncertainty: Strategies for Surviving and Thriving in Turbulent Times*. London: Economist Books, 2012.
- Tapscott, Bob, and Don Tapscott. "Risk Management 2.0: Overcoming the Current Financial Crisis and Restoring Stability and Prosperity with a New Perspective on Risk." nGenera Insight *Enterprise 2.0*, October 2008.
- Taleb, Nassim Nicholas. *Foiled by Randomness: The Hidden Role of Chance in Life and Markets* (New York: Random House, 2004).
- . *The Black Swan: The Impact of the Highly Improbable*. New York: Random House, 2007.

- . *Antifragile: Things That Gain from Disorder*. New York: Random House, 2012.
- Taleb, Nassim Nicholas, Daniel G. Goldstein, and Mark W. Spitznagel. "Six Mistakes Executives Make in Risk Management." *Harvard Business Review*, October 2009.
- Tetlock, Philip, and Dan Gardner. *Superforecasting: The Art and Science of Prediction*. New York: Crown, 2015.
- . "Thinking the Unthinkable: Sacred Values and Taboo Cognitions." *Trends in Cognitive Sciences*, Vol. 7, No. 7 (July 2003).
- Tetlock, Philip, and Barbara A. Mellers. "The Great Rationality Debate." *Psychological Science*, Vol. 13, No. 1 (January 2002).
- . "Accountability and the Perseverance of First Impressions." *Social Psychology Quarterly*, Vol. 46, No. 4 (1983), 285–292.
- Tversky, Amos. "Subjective Probability: A Judgment of Representativeness." *Cognitive Psychology*, Vol. 3 (1972), 430–454.
- Tversky, Amos, Shmuel Sattath, and Paul Slovic. "Contingent Weighting in Judgment and Choice." *Psychological Review*, Vol. 95, No. 3 (July 1988), 371–384.
- United States Army. "Risk Management." Field Manual 100–14, April 23, 1998.
- United States Army Training and Doctrine Command. *The U.S. Army Operating Concept: Win in a Complex World*. Joint Base Langley-Eustis, Virginia: TRADOC, October 31, 2014.
- United States Central Intelligence Agency. "A Tradecraft Primer: Structured Analytic Techniques for Improving Intelligence Analysis." March 2009.
- United States Department of Defense. *Quadrennial Defense Review Report*. Washington, DC: DoD, multiple years.
- . "Risk Management Guide for DoD Acquisition." 6th Ed. August 2006.
- United States Department of Homeland Security. *Risk Management Fundamentals*. Washington, DC: DHS, April 2011.
- Vaughan, Diane. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. Chicago: University of Chicago Press, 1996.
- . "The Dark Side of Organizations: Mistake, Misconduct, and Disaster." *Annual Review of Sociology*, Vol. 25 (1999).
- Vertzberger, Yaacov. "Rethinking and Reconceptualizing Risk in Foreign Policy Decision-Making: A Sociocognitive Approach." *Political Psychology*, Vol. 16, No. 2 (June 1995).
- Walker, Robert. "Beware the 'Gray Swan'." *The World Post*, May 18, 2011.
- Walt, Stephen. "Rigor or Rigor Mortis: Rational Choice and Security Studies." *International Security*, Vol. 23, No. 4 (Spring 1999).
- . "The Credibility Addiction." *Foreign Policy*, January 6, 2015.
- Ward, Stephen, and Chris Chapman. "Transforming Project Risk Management into Project Uncertainty Management." *International Journal of Project Management*, Vol. 21 (2003).
- Weatherall, James Owen. *The Physics of Wall Street: A Brief History of Predicting the Unpredictable*. New York: Mariner Books, 2013.
- Wighton, David. "Imagination and Common Sense Brew a Safer Culture." *FT.com*, November 28, 2007.
- Wolf, Martin. *The Shifts and the Shocks: What We've Learned—And Have Still to Learn—From the Financial Crisis*. New York: Penguin Press, 2014.
- World Economic Forum. *Global Risk Report*, 10th Ed. Geneva: WEF, 2015.
- Wright, Thomas. "The Rise and Fall of the Unipolar Concert." *The Washington Quarterly*, Vol. 37, No. 4 (Winter 2015).
- Yates, J. F., ed. *Risk-Taking Behavior*. New York: Wiley, 1992.

Index

- Adams, John, 7, 28, 217n31
Afghanistan, 18, 37, 44–6, 113, 151–2,
161, 211n6, 220n27
al Qaeda, 91–2
ambiguity
and causalities, 150
and personality, 103, 153
and strategic choices, 11, 12, 161, 169,
170
and uncertainty, 52–3, 61, 97, 146, 150,
189, 191
American International Group (AIG), 8,
85–6, 93, 95, 106–7, 115–19, 129,
131, 135
Asian financial crisis, 51, 65–6, 117, 119,
130

Bank of America, 17
Baron, Jonathan, 80, 88, 176–7, 220n22,
229–30n14
Bay of Pigs invasion, 113, 146, 184
Bear Sterns, 8, 65, 105–6, 114, 134–5
Bernanke, Ben, 15, 182
biases. *See* cognitive biases
big bets, 4, 13, 36, 38–9
Bissell, Richard, 106
Brigade Combat Teams (BCTs), 41–2
Black-Scholes model, 51, 56, 214n20
black swans, 89–92, 98, 132, 139. *See also*
gray swans; Taleb, Nassim Nicholas
Blustein, Paul, 66
Blyth, Mark, 68
Borget, Louis, 101–2
Bundy, McGeorge, 78
Burrough, Brian, 65, 106, 114
Burton, Robert, 72, 75
Bush, George W., 91, 96, 122, 158
Byman, Daniel, 111

Campbell, John, 37
Capability Review and Risk Assessment
(CRRA), 35
capability-to-task matrix, 204–5
Cassano, Joe, 86, 93, 95, 106, 119, 129,
131

Cassidy, John, 93, 108, 132
causality, use of the term, 12
Cayne, Jimmy, 105, 115
Central Command (CENTCOM), 41–2
certainty effect, 80
Challenger disaster, 56, 67, 136–7
Chanos, Jim, 110
Cheney-Peters, Scott, 193, 204, 230n17
China, 5, 41, 152, 166–9, 172
Cioffi, Ralph, 114
Clark, Christopher, 81
Clarke, Richard, 91–2
cognitive biases, 7, 9, 29–30, 33, 53, 68,
74, 125, 127, 146, 182
confirmation bias, 69
illusion of control, 61
invented patterns, 69
optimistic bias, 107
overconfidence, 4–5, 58, 61, 74, 92,
104, 106–8, 113, 133–5
and personality, 103–4, 106–7, 110, 112
wishful thinking, 9, 14, 30, 58, 61, 75,
94–5, 97, 106–7, 113, 171
Cohan, William, 105, 114–15, 134–5
complex strategic judgments
and adaptive strategies, 200
and ambiguity, 93–4, 170
and anticipation, 209n13
and causalities, 12, 150, 170
challenges of outcome analysis, 152–5
clarifying theories of success, 151–2
and consequences, 14–15
criteria for judging outcome-oriented
risk, 152
and data-based analysis, 11, 45, 59–60, 67
defined, 4, 10–12, 143–4
emergent and responsive approach, 200
and feeling-of-knowing 32–3
importance of alternative views, 150–1
and multiple values/objectives, 26
objective of driving necessary
conversations, 8, 144–6
and national security, 17
and perception, 32–3, 67, 69, 155
and risk, 21–2, 26–8, 46, 67, 89

- complex strategic judgments – *continued*
 risk as outcome analysis, 146–8
 risk as relative to strategic concept,
 148–50
 and uncertainty, 11–13, 27–8, 36, 52,
 58–62, 200, 208n14
- consequences, 77–8
 and the financial crisis, 85–8
 and imperative-driven thinking, 81–5
 and national security, 77–81
 and nonconsequentialist decision-
 making, 78–81
- Countrywide, 93, 110, 178
- Crimean War, 103–4
- Cuba. *See* Bay of Pigs invasion
- culture
 culture of dissent, 177
 monoculture, 204–5
 organizational, 4–5, 30–1, 69, 94, 98,
 101, 104–8, 125–6, 133–41, 144, 146,
 175–80, 199, 202, 204–6
 political, 97, 160, 162–4
- Danzig, Richard, 13, 190, 199
- decision-making
 and affect, 79
 and decision fatigue, 94–5
 and future-oriented activities, 79–80
 and hedgehogs versus foxes, 111–12
 and imperative-driven, 81–8
 and incentives, 129
 and institutional culture, 126
 and loose control, 199–200
 and methodology, 8
 nonconsequentialist, 78–81
 procedures and frameworks, 179–80
 and prominence effect, 79, 219n11
 and protected values, 80
 and sacred values, 80–1
 and sharing authority, 198
 and understandings of risk, 23–4, 32,
 35–8
 and warning overload, 94
- Dempsey, Martin, 6, 19, 21, 189
- Department of Defense (DoD), 77, 144–5
 and fragmentation and big bets, 38–9
 framework for risk, 36–8
 and institutionalizing management of
 uncertainty, 202–5
 and modularity, 197
 and oversight, 122
 and quantification, 44
 Quadrennial Defense Reviews (QDRs), 5,
 6, 19, 36–7, 39, 42, 147, 189, 211n6
 regional commands, 41–2
 and stoplight charts, 46
- Department of Homeland Security (DHS),
 5, 22–3, 37–8
- deterministic and non-deterministic
 environments, 10–14, 16, 51–8, 64,
 66, 96
- deterministic risk analysis, 24, 27–9, 33,
 192–3
- Dillon-Merrill, Robin, 93
- Dirksen, Everett, 78
- Douglas, Mary, 126, 202
- Dowie, Jack, 7, 31–3, 207n13
- Earl, Peter, 72
- efficient market hypothesis, 55–6, 58, 129
- Eisenhower, Dwight, 103, 149, 162, 197
- Elkind, Peter, 14, 101, 105, 107, 109–10,
 135
- Ellsberg, Daniel, 53
- Ellsberg Paradox, 53
- Enron Corporation, 8, 14, 101–2, 105,
 107–10, 114–15, 130–1, 135–6, 138,
 146, 176, 178
- European Command (EUCOM), 41
- false precision, 43–5, 183
- Fannie Mae, 86, 110, 128–9, 138
- Farrell, Greg, 109, 136
- Federal Open Market Committee, 54
- feeling-of-knowing, 33, 66, 69–75, 82, 83,
 85, 150, 155
- Figes, Orlando, 103–4
- financial crisis of 2007–2008
 and consequences, 85–8
 and gray swans, 97–8
 and human factors, 101–2, 104–9
 and ignorance, 114–21
 and imperatives, 85–8
 and incentives, 13, 63, 65, 73, 86, 95,
 110, 117, 118, 126–36, 182
 risk in, 15–17, 65–75
 and threat, 51, 55, 126, 129, 182
 and uncertainty, 51–62, 92–3
*See also individual organizations and
 individuals; Department of Defense
 (DoD)*
- Fitzgerald, Ben, 193, 204, 230n17

- force management risk, 35, 36, 211n6
 Freddie Mac, 135
 Friedman, Thomas, 65
 Fuld, Richard, 86, 104, 105, 114, 127,
 129–30, 216n3
 future challenges risk, 36, 211n6
- Geithner, Timothy, 3
 Goldin, Ian, 120–1, 209–10n16
 Goldman Sachs, 8, 98, 104–5, 110, 134–5,
 136, 177, 181
 gray swans, 32, 62, 75, 89–93, 95–9, 111,
 176, 187, 220–1n3. *See also* Taleb,
 Nassim Nicholas
 Great Britain, 157, 159–65, 168–70, 173,
 198
 Greenberg, Alan “Ace,” 105
 Greenberg, Hank, 116
 Greenspan, Alan, 3, 15, 54, 56, 108, 116,
 129, 182, 217n25
 groupthink, 4, 7, 14, 85, 107, 113, 130,
 139
- Haldane, Andrew, 194
 hedgehogs and foxes, 111–12
 hedging strategy, 193–4, 230n17
 herding, 4, 9, 55, 92, 108, 121, 130
 Hoffman, Frank, 200
 human factors, 4, 5, 8–10, 14–17, 21,
 29–33, 143–5, 175, 184, 187
 herding, 4, 9, 55, 92, 108, 121, 130
 institutional culture, 4–5, 30–1, 69, 94,
 98, 101, 104–8, 125–6, 133–41, 144,
 146, 175–80, 199, 202–6
 See also cognitive biases; incentives;
 personality
 Hurricane Katrina, 91
 Hussein, Saddam, 96, 122
- Iceland, 107, 116
 ignorance, 114–23, 146–7, 180, 212n2
 incentives
 and the financial crisis, 13, 63, 65, 73,
 86, 95, 110, 117, 118, 126–36, 182
 and group behavior, 130–1
 in group settings, 131–2
 malign/skewed, 4, 7, 14
 and national security, 138–9, 168, 198
 and organizational culture, 132, 176,
 178
 and Value at Risk, 63
- incomparability of the specific and the
 general, 28–9
 institutional risk, 36, 211n6
 International Monetary Fund (IMF), 65–6,
 117, 130
 Iraq, 10, 31, 44, 92, 94, 96–7, 113, 122,
 146, 180, 199, 211n6
- Johnson, James, 86, 110, 128–9, 135, 138
 Johnson, Lyndon, 77–8, 166
 Johnson, Simon, 63
 judgment, 4–5, 7–8, 10–18
 emergent, 69–71, 74–5, 85, 155
 feeling-of-knowing, 71–4
 use of the term, 12
 See also complex strategic judgments
- Kahneman, Daniel, 80, 107, 184, 217n25
 Kaplan, Robert, 20, 176, 178, 181
 Katzenstein, Peter J., 132
 Kennedy, John F., 103, 184
 Kennedy, Paul, 160–1
 Keynes, John Maynard, 132
 Killinger, Kerry, 104
 Kissinger, Henry, 13, 165, 167–8
 Knight, Frank, 57, 58–9, 72
 Korea, 10, 13, 40, 166, 186
 North Korea, 18, 27, 40, 67, 148
 South Korea, 18, 65, 67, 148
 Kronthal, Jeffrey, 93
 Kupchan, Charles, 163
 Kwak, James, 63
- Lafrance, Adrienne, 91
 Langevoort, Donald C., 133–4
 Laplace, Pierre-Simon de, 57–8
 Lay, Ken, 115–16
 Lewis, Michael, 87, 106–7, 116, 199, 131
 Liddy, Ed, 116
 Lindblom, Charles, 138
 Lindsey, Lawrence, 132
 Lipshitz, Raanan, 77
 Long-Term Capital Management (LTCM),
 8, 16, 51–2, 55–6, 73, 86–7, 93, 107–9,
 117–21, 132–3, 135
 Lowenstein, Roger, 55, 61–2, 87, 104, 114,
 117, 119, 128–9, 132–3, 135
- Mack, 104–5, 127–8
 Mandelbrot, Benoit, 56
 Mandelbrotian Gray Swans, 90

- Mangan, Bruce, 72
 March, James, 23–4, 200
 Mariathasan, Mike, 120–1, 209–10n16
 Mark, Rebecca, 105
 McLean, Bethany, 14, 86, 101, 105, 107, 109–10, 128, 135
 McNamara, Robert, 13, 77
 Mead, Walter Russell, 159
 Meriwether, John, 51, 109, 119
 Merrill Lynch, 8, 16–17, 62, 87, 93, 105, 109, 114–15, 127, 130–5
 Merton, Robert, 51, 133. *See also* Black-Scholes model
 methodology, 8–10
 Microsoft, 133
 Mikes, Anette, 20, 176, 178, 181
 Mintzberg, Henry, 200
 Mlodinow, Leonard, 69
 Morgan Stanley, 104–5, 127–8
 Morgenson, Gretchen, 16, 93, 129, 138
 motivated reasoning, 75, 92, 94–5.
See also wishful thinking
 Mozilo, Angelo, 110
 Muckleroy, Mike, 101–2
- National Aeronautics and Space Administration's (NASA), 4–5, 38, 56, 73, 136–7
 national security
 approaches to risk in, 35–47
 and consequences, 77–81
 Great Britain case study, 159–65
 and human factors, 74–5, 103–4, 106–7, 111–12
 and ignorance, 121–3
 and incentives, 138–9, 168, 198
 Nixon-era case study, 165–9
 and outcome-based risk assessment, 157–73
 and risk, 35–47
 and rise of “risk,” 5–8
 selective engagement, 157–8, 165, 169–73
 and threats, 6, 35–41, 43, 45, 67, 92, 148–9, 152, 157, 159–60, 164–72
 and uncertainty, 62–4, 91–2
See also individual conflicts, leaders, and nations
 Nelson, Stephen C., 57, 132
 Nicholas I of Russia, 103–4
 Nixon, Richard, 103, 165–72, 197
 Nixon Doctrine, 157, 166–9
 Nocera, Joe, 63, 132
 nonergodic systems, 29, 59
 normalization of deviance, 136–8
 North Korea, 18, 27, 40, 67, 148. *See also* Korea
 Odierno, Raymond, 42
 O’Neal, Stan, 16–17, 93, 109, 114, 127, 130, 135
 operational risk, 35, 36–7, 41, 44, 211n6
 optimistic bias, 107
 outcome, use of the term, 12
 outcome-oriented risk analysis, 83, 87, 180, 184–7
 and alternative views, 150–1
 categories of outcome risks, 154
 challenges with, 148, 152
 and cognitive biases, 146
 criteria for judging, 152–5
 defined, 77–9
 Great Britain case study, 159–65
 and national security, 77–8, 157–73
 Nixon-era case study, 165–9
 outcome-oriented risks of selective engagement, 169–73
 and theories of success, 151–2
 and uncertainty, 198–9
 overconfidence, 4–5, 58, 61, 74, 92, 104, 106–8, 113, 133–5
 overprecision. *See* false precision
- Patterson, Scott, 51, 62, 104, 214n20
 Paulson, Henry, 93, 110
 Perrow, Charles, 120
 personality, 101–2
 and crisis, 104–8
 as engine of interpretation, 103–4
 and the financial crisis, 101–2, 104–9
 and national security, 103–4, 111
 personalization of risk, 102–3
 and relationships, 108–10
See also human factors
 Pollack, Kenneth, 111–12
 Prince, Charles, 132
- quantification, 9, 11, 15, 17, 44–6, 64, 144, 146, 207n13, 209–10n16, 221n4
 Raines, Franklin, 86, 128
 Ramo, Joshua Cooper, 56–7, 202
 RAND Corporation, 38, 54, 92, 145

- rationality, 29, 56–8, 62–3, 72, 77, 80–1, 84, 96, 103, 199, 200
- Ravenal, Earl, 168–9, 228n28
- resilience, 64, 89, 98–9, 120, 129, 175, 187, 191–6, 201–5
- Rice, Condoleezza, 122
- risk and risk management, 13–15
 - accumulating risk, 117–18
 - cascading risk, 121–2
 - and chance, 23
 - and consequences, 77–88
 - correlated (systemic) risk, 118–21
 - dangers of risk analysis, 31–3
 - defining risk, 19–24
 - Department of Defense framework for, 36–8
 - deterministic risk analysis, 24, 27–9, 33, 192–3
 - and ends-means gaps, 41–4
 - and feeling-of-knowing and emergent judgment, 69–74
 - force management risk, 35, 36, 211n6
 - and fragmentation and big bets, 38–9
 - future challenges risk, 36, 211n6
 - and the general and particular, 28–9
 - and ignorance, 114–23
 - and incentives, 125–39
 - interpretative risk, 68–9
 - mitigation strategies, 31
 - and organizational culture, 133–41, 144, 146, 175–80
 - and personality, 101–12
 - principles of effective risk management, 175–87
 - and quantification, 44–7
 - revised approach to categorizing, 186
 - and reward, 24–7
 - and strategy, 17–18, 31, 143–55
 - as threat, 39–41
 - and uncertainty, 27–31, 33–4, 51–64*See also* outcome-oriented risk analysis
 - Risk Assessment and Control (RAC), 14
- Rosner, Joshua, 16, 93, 129, 138
- Roubini, Nouriel, 92
- Rubin, Howie, 105
- Rubin, Robert, 114
- Rumsfeld, Donald, 106, 122
- Russell, Richard, 77
- Russia, 5, 12–13, 25, 27, 40, 45, 51, 73, 103–4, 118–19, 148–9, 152, 186, 193, 198
- Salomon Brothers, 51
- Savage, Tom, 106
- sandpile effect, 57, 120
- Schacter, Barry, 62
- Schein, Edgar, 133
- Scholes, Myron, 51, 56, 133
- Semerci, Osman, 109, 114–15
- September 11, 2001, 89, 91–3, 96
- Shackle, G. L. S., 54, 60, 96, 215n43, 218n36
- Shapira, Zur, 23–4
- Shafir, Eldar, 81, 218n3
- Shiller, Robert, 92
- Silver, Nate, 53, 60, 62, 92, 131
- Skilling, Jeffrey, 14, 105, 107, 109, 116
- Sorkin, Andrew Ross, 86, 93, 103, 116, 119, 127, 130
- South Korea, 18, 65, 67, 148. *See also* Korea
- Spector, Warren, 115
- Stalin, Josef, 103
- Steinbruner, John, 69, 71
- stoplight charts, 46–7
- Sullivan, Gordon, 189
- superforecasters, 30, 177, 213n9
- swans. *See* black swans; gray swans
- Syria, 27, 29, 36
- Taleb, Nassim Nicholas
 - and 9/11, 91
 - and “antifragile” systems, 194
 - and Black-Scholes equality, 56
 - and black swans, 89–92, 98, 132, 139
 - and Mandelbrotian gray swans, 90, 221n8
 - and Mediocristan and Extremistan, 58–9
 - and narrative fallacy, 68
 - and uncertainty, 4, 58–9, 96–7, 194, 196, 201*See also* gray swans
 - Tetlock, Philip, 30, 80–1, 177, 184, 209n13, 213n9
- threat, 18–26, 31–4, 121
 - and causalities, 12
 - defined, 40
 - and the financial crisis, 51, 55, 126, 129, 182
 - and imperatives, 84
 - and managing uncertainty, 189–91, 193–4, 196–9, 202–5
 - and mitigation strategies, 31

threat – *continued*

- and national security, 6, 35–41, 43, 45, 67, 92, 148–9, 152, 157, 159–60, 164–72
- and outcome risk, 152–3, 157
- and resilience, 129
- risk as, 39–41, 147–8, 152
- See also* black swans; gray swans
- transmutability, 11, 16, 54, 61, 68–71, 170
- transparency, 133, 182–3, 185, 206
- Tversky, Amos, 80–1, 218n3

uncertainty

- and alternative perspectives, 193–4
- and ambiguity, 52–3, 61, 97, 146, 150, 189, 191
- and anticipatory shaping, 201
- characteristics of uncertain challenges, 60–2
- diagnosis of the situation, 192–3
- and diversity, 196
- and emergent and responsive approach, 200
- and the financial crisis, 92–3
- and innovation and experimentation, 199
- institutionalizing management of, 202–5
- and leverage points, 198–9
- and loose control, 199–200

- managing, 189–206
- and modularity, 197
- and national security, 62–4, 91–2
- ontological (aleatory) uncertainty, 59–61, 70–1
- and patience, 201
- and probability, 55–60
- and redundancy, 196
- and resilience, 187, 191–6, 201–5
- and risk, 27–31, 33–4, 51–64
- and robustness, 192–7, 202–3, 205
- and stabilizing feedbacks, 197–8
- and swans, 89–99
- understanding, 52
- and versatility, 200

- Value at Risk (VaR), 15, 30–1, 63, 118, 132
- Vaughan, Diane, 56, 67, 73, 136–8
- Vertzberger, Yaacov, 28, 208n21
- Vietnam War, 25, 53, 56, 77–8, 82, 165–7, 169, 220n27

- Washington Mutual (WaMu), 86, 128
- Welch, Jack, 111
- Wildavsky, Aaron, 126, 202
- wishful thinking, 9, 14, 30, 58, 61, 75, 94–5, 97, 106–7, 113, 171
- Wolf, Martin, 129
- World War I, 81, 159–61, 163–4
- World War II, 149, 162, 163