Chapter 9
Reflections

James Scouras

Motivated by the importance of the perceived risk of nuclear deterrence failure in national security policy formulation, we began our study by asking whether more structured analytic approaches could improve on the highly intuitive manner by which the risk of deterrence failure has generally been assessed. For the likelihood dimension of risk, each of the approaches included in this book—case study, elicitation of expert judgment, probabilistic risk assessment, and application of complex systems theory—has something unique to offer. However, none of these approaches can do the job by itself. Rather we have reinforced the notion that multiple disciplines can each shed limited light on the question. We must extract from each of them whichever valuable insights they offer and do our best to synthesize these insights, using the art and science of knowledge integration, into a policy-relevant assessment. However daunting this task, discernible research paths hold significant promise. As for the physical consequences of nuclear use, it is clear that our knowledge base, derived primarily from concern about the military effectiveness of nuclear weapons, is inadequate to assess the potential consequences from the broader array of nuclear uses that now appear possible or from intangible consequences that could exceed even the physical consequences. This lack of knowledge is easier to address from an analytic perspective but requires an adequately funded research program. The dim prospects of such a program are yet another consequence of the complacency induced by our intuitive sense that nuclear weapon risks have largely abated.

It is remarkable how rapidly perspectives on nuclear risk have changed over the three post–Cold War decades. In the incredulity of our good fortune after the end of the Cold War and the demise of the Soviet Union, concern with the residual nuclear threat from Russia plummeted. Nuclear war was deemed highly improbable, and all things nuclear were relegated to much

lower priority in national security planning. A decade later, after the attacks of September 11, 2001, terrorism—including nuclear terrorism—ascended to the top of the threat priority list and remained there until about five years ago. More recently, the emerging nuclear threat from North Korea, the rebuilding of the Russian nuclear arsenal, and the gradual rise in the Chinese nuclear threat, all three of which have been accompanied by belligerent international behaviors, have emplaced interstate nuclear crises and war with these actors as the primary nuclear risks.

Beyond the question of the extent to which this shuffling of the deck accurately reflected international realities, several points are worth making. First, perspectives on nuclear risk have been almost always focused on likelihood. Consequences have been rarely considered as a coequal component of risk. This is most obvious in the elevation of nuclear terrorism after 9/11 to higher concern than global nuclear war with Russia. While the arsenals of the United States and Russia had dropped significantly by then from their Cold War levels, they still numbered in the thousands of weapons. There remained the possibility—perhaps remote, perhaps not—that deterrence might fail and these arsenals would be used. Thus, while the likelihood of nuclear terrorism appeared to be growing, a global conflagration that involved the US and Russian nuclear arsenals was—and remains—among the more horrific catastrophes we could imagine.

Second, long-term (a decade or more, for argument's sake) projections of nuclear risk are inherently suspect. The world has changed and can reasonably be expected to continue to change too rapidly to justify confidence in risk assessments beyond the short to intermediate term. Thus, those mathematical calculations that suggest any nonzero and nondecreasing annual risk of nuclear war compounded over many years will inevitably lead to catastrophe need to be reconsidered in light of the improbability of the assumptions about nondeclining future risk.

As this discussion suggests, this work has been motivated in large part by the concern that the conventional wisdom regarding possible nuclear weapon use—from global nuclear war involving the arsenals of the nuclear superpowers to terrorist nuclear use of a single weapon—has not been adequately challenged. Are there analytic approaches that will allow us to move beyond intuition and overly simplified analyses to a more rigorous basis for assessing the risk of deterrence failure?

What Have We Learned?

Clearly, each of the approaches included in this study that contribute to our understanding of the likelihood of nuclear use—historical case study, elicitation of expert judgment, probabilistic risk assessment, and the application of complex systems theory—has something unique to offer. Because none of these approaches is by itself a silver bullet, we must extract from each of them whatever valuable nuggets (or morsels) of insight they offer and do our best to synthesize these insights, using the art and science of knowledge integration, into a policy-relevant assessment. Then, utilizing a risk assessment framework, likelihood must be combined with an analysis of the prospective consequences of nuclear use, for which there is a very large body of accumulated knowledge but also large uncertainties and enormous gaps.

Historical Case Study

Knowledge of the history of actual use of nuclear weapons in World War II and close calls of potential use during the Cold War and the post–Cold War period provides the essential foundation for any assessment of future risk. This history helps to identify paths to close calls and use as well as to assess past risks and contemporaneous perceptions of those risks. Without an awareness of the history of close calls, it is difficult to appreciate the myriad and unexpected ways in which nuclear war could be triggered. The unanticipated and idiosyncratic nature of many of these close calls, such as the 1995 Norwegian meteorological rocket that Russia briefly considered to be a possible US nuclear attack, should also engender an appropriate humility in any prognostications about the future.[1]

Extracting historical lessons relevant to the future is not straightforward. Fundamentally, the future is not the past. History is easily misused, and it is difficult to generalize from diverse and infrequent close calls. Facts are limited and those known may be biased, and historians often differ in their interpretations. Most important, case study must be combined with expert input to assess future risks, and therein lies further opportunity for subjectivity and disagreement. Nevertheless, to know how nuclear weapons might be used, we must know how they could have been used. Thus, study of close calls provides a necessary, but incomplete and uncertain, guide to the future.

Our work suggests two primary directions for further historical research. First, past close calls should be revisited as new source materials become available. For example, the Center for Naval Analyses conducted a fresh study of the 1969 Sino-Soviet border conflict utilizing newly available primary and secondary sources.[2] Beyond understanding the roles of nuclear weapons and policies in the progression and outcome of past crises, the goals of such studies should include an evaluation of the potential applicability of lessons extracted from these experiences to current and future nuclear challenges. Second, because lessons from history are applied mainly through the use of analogy, a greater appreciation of the historical record, limitations, and legitimate usage of analogies when applied to nuclear close calls would be immensely helpful in avoiding the more common misuses of history by policy-makers and others.[3]

Elicited Expert Knowledge

For data-sparse, theory-poor problems such as assessing the risk of nuclear deterrence failure, heavy reliance is placed on knowledge from experts. Formal elicitation methods to extract the best-quality knowledge from experts have been developed and successfully applied to diverse problems over the past several decades. Moreover, expert elicitation is an active research field with techniques continuing to improve. Awareness of the biases peculiar to assessing the risk of deterrence failure with its rather unique aspects (e.g., the challenges of thinking about the unthinkable) should allow elicitation methods to also be successfully tailored to that problem.

Unfortunately, formal elicitation methods are time consuming and expensive to employ. They have not been applied in past analyses of the risk of nuclear deterrence failure, resulting in data with dubious quality and suspect conclusions. Major improvements in the utility of elicited information would result from such simple practices as not relying on self-elicitations and capturing experts' thinking and their uncertainties. In addition, elicitations that directly ask for experts' estimations of the risk of nuclear weapon use suffer from assuming that a single "expert" can make informed judgments across the entire scope of the problem, rather than parsing the problem into smaller pieces that can be addressed more authoritatively by different experts.

Notwithstanding the potential of utilizing elicited knowledge, a fundamental limitation of formal elicitation must be recognized. Formal elicitation at its best can only extract knowledge that exists or that elicitations can provoke experts to develop by thinking through a question. Many of the issues in assessing the risk of deterrence failure are beyond experts' knowledge and analytic capabilities. This reality is not an argument against using formal elicitation methods; rather it a caution that should be reflected in the uncertainties associated with experts' judgments.

Probabilistic Risk Assessment

As with formal expert elicitation, probabilistic risk assessment is being applied to an ever-broader range of problems. However, its major successes have been in assessing risks associated with engineered systems such as nuclear power plants and the space shuttle.[4] Probabilistic risk assessment is far less mature for problems involving complex human interactions, where the range of possible decisions and actions cannot be identified in advance and actors are adaptive. An elaborate attempt to apply it to such a problem, the Department of Homeland Security's Bioterrorism Risk Assessment, has been sharply criticized as fundamentally flawed in a National Academy of Sciences report.[5]

Two other challenges of utilizing probabilistic risk assessment—one real and one that is more accurately characterized as a misguided concern—derive from the quantification of likelihood using probability. The real challenge is that many subject-matter experts are not particularly adept at estimating probabilities and tend to underestimate associated uncertainties.[6] Somewhat surprisingly, this observation is valid for experts both trained and untrained in mathematically based disciplines, including even statisticians.[7] On-the-spot training and practice at the front end of an elicitation do not seem to help much.[8] The misguided concern is that a well-executed probabilistic risk assessment, with experts' uncertainties in their judgments accurately captured, would not be helpful because many experts' uncertainties could be very large and, consequently, the probability of deterrence failure would have significant nonzero values across a very broad range of possibilities. However, rather than being uninformative, such a result would support two critical considerations for policy: experts are highly uncertain as to the risk of deterrence failure and the probability could be high or low. This, of course, helpfully undermines unwarranted

confidence in the conventional wisdom that we can rest assured that the risk of global nuclear war is low. More precise conclusions may simply not be supported by our current state of knowledge.

While probabilistic risk assessment cannot be relied on to provide a definitive analysis when applied to the risk of nuclear deterrence failure, it does have important uses. In particular, it requires defining outcomes (e.g., global nuclear war between the United States and Russia, regional nuclear war in south Asia, or terrorist use of a single nuclear weapon in a European city) and identifying paths to these outcomes, each path starting with a triggering event and progressing step by step to an outcome of concern. These intermediate products are valuable even if all possible outcomes and paths are not identified and even if probabilities are not assigned to the various (or *all* the various) steps along each path. Developing these products in the context of a probabilistic risk assessment facilitates clarity in thinking and dialogue among experts to identify points of agreement and disagreement. To adapt the oft-quoted wisdom of George Box, essentially all probabilistic risk assessments are wrong, but some are useful.[9]

Finally, while not (yet) ready for prime time, probabilistic risk assessment has significant potential for improvement, as discussed in Martin Hellman's chapter. Focusing on historical examples of accidents and close calls should provide useful bases for both scenario development and expert elicitation of needed probabilities, however uncertain.

Complex Systems Theory

Our work has established that deterrence is an example of a complex system in operation. Thus, complex systems theory offers a vantage point for thinking about the potential failure of deterrence. However, like probabilistic risk assessment, while complex systems theory has had success in a variety of realms ranging from physical to biological and economic systems, relatively little research has focused on applications to international relations. The challenge lies in developing meaningful psychological and sociological behavior models for the human components of nuclear deterrence systems.

Nevertheless, general principles from complex systems theory can be applied to the risk of deterrence failure. Among them is the notion that system behaviors are not always presaged in component behaviors, which results in so-called emergent behaviors. Conventional systems assessment

approaches inevitably miss rare, high-consequence events that must occur in complex systems. Complex systems theory also reinforces the importance of defining system boundaries carefully lest important interactions be excluded, giving rise to unintended consequences of deterrence policies and actions.

Progressing beyond discussion of general principles from complex systems theory will require development of a model of the nuclear deterrence system. Given the criticality of the continued success of deterrence, it is somewhat surprising that such a model has apparently never been developed—or possibly even attempted—except at a high level of abstraction with a narrow focus on damage calculations in stylized scenarios of strike and counterstrike (so-called exchange analysis). If carefully constructed and continually improved, a complex systems model of deterrence would enable the analytic tools (e.g., simulation) of complex systems analysis to help identify otherwise hidden failure modes, anticipate the impact of alternative deterrence policies and strategies, and perhaps even estimate the overall risk of deterrence failure.

Knowledge Integration

We have seen that while each of the four approaches examined in this study has something of value to contribute to assessing the likelihood of deterrence failure, that value may be difficult to extract. Casting such a wide methodological net creates yet another challenge: how to summarize knowledge obtained from such diverse approaches. To illustrate, historical case studies may provide some evidence that, for example, the progression of close calls to their final resolutions has been unpredictable and may support the inference that the outcomes of future close calls are also likely to be so. Elicitation of experts on the risk of deterrence failure may provide a sampling of opinion, qualitative and quantitative, and supporting thought processes. Probabilistic risk assessment, if all goes well, may provide quantitative probability distributions of the likelihoods of various paths to nuclear use. And considering the risk of deterrence failure from the perspective of complex systems theory may provide insights into limitations of more traditional systems analyses. Combining such diverse forms of knowledge from multiple approaches, all with uncertainties and both confirming and conflicting data, interpretations, and conclusions, is the challenging objective of knowledge integration.

We cannot evade this challenge. Knowledge integration allows us to summarize and draw conclusions, facilitating policy-making and communication. Without integration and summary, we are free to cherry-pick only those facts and arguments that support our preconceived notions. The process is not dissimilar from what we ask trial jurors to do without the benefit of much guidance from judges' instructions. We can learn something by thinking through how a reasonable juror might approach the problem.

While the more esoteric mathematical approaches that characterize some of the advanced research in knowledge integration have their place, it might also be helpful to focus on fundamental principles and practical approaches that emphasize simplicity, completeness, traceability, and transparency. Structures for assessing the quality of evidence, arraying it in support of and opposed to a hypothesis, and presenting the chain of logic that connects the evidence to a hypothesis would be immensely helpful. Uncertainties, contradictory evidence, and alternative hypotheses must be included as well. Such a disciplined, structured process would facilitate constructive dialogues among experts, policy-makers, and the public. And then it is up to fallible human beings, relying on neither mathematical exotica nor intuition, to render final judgments.

Consequences of Nuclear Weapon Use

Finally, we have concentrated on the likelihood component of risk, but a complete risk assessment of deterrence failure must also address the consequences of nuclear weapon use. While assessing the likelihood of nuclear use may be daunting, assessing the consequences is not a trivial task either. Despite more than one thousand US nuclear tests between 1945 and 1992 and extensive analyses of the effects of the Hiroshima and Nagasaki atomic attacks, there remain significant uncertainties with regard to consequences of nuclear weapon use. These uncertainties derive from inadequate data and study of those nuclear weapons effects, such as fire, that are not the primary mechanisms for destroying targets of interest to military planners; from phenomena, such as electromagnetic pulse, that were discovered only late (and by accident!) in the US nuclear testing program; from insufficient understanding of environmental effects, such as climate change, and societal effects, such as the robustness/fragility of economies and political institutions; and from a limited ability to model global,

cascading, and/or long-term consequences of all types. In truth, we have only an inkling of the full scope of consequences of nuclear use, whether such use involves only one weapon or thousands. Military consequence assessments focus on damage to specific targets and miss the larger impacts of nuclear use on society and the environment. Ironically, it is these broader effects that will undoubtedly weigh most heavily on the minds of national leaders who must ultimately authorize nuclear weapon use.

It might be argued that it is enough to know that the consequences of even a single nuclear explosion would be horrific. While appealing in its simplicity, this perspective provides an inadequate basis for policy-making. Some nuclear wars will lead to 10,000 or fewer deaths; others to 1,000,000,000 or more. Beyond the 999,990,000 survivors of the smaller war who might take issue with the notion that all nuclear wars are essentially indistinguishable, there are other distinctions of importance. Some nuclear wars will bring about the demise of the United States as a political entity; others will not. Some nuclear wars will lead to a severe and long-term curtailment of civil liberties in the United States; others will not. Some nuclear wars will induce atmospheric changes that affect agricultural production across the planet, leading to mass starvation of hundreds of millions of human beings far beyond the borders of the belligerents; others will not. Some nuclear wars will encourage proliferation; others will not. Some nuclear wars will strengthen the nuclear taboo; others will not. While even the smallest nuclear war threatens consequences far worse than many calamities we can imagine, casting a blind eye to the varying degrees of horror across the spectrum of possible nuclear wars is simply irrational.

The perspective that the consequences of alternative potential nuclear wars are essentially indistinguishable in their horror is also not without policy dangers. In particular, it encourages the simplistic notion of independence between the likelihood of deterrence failure and the level of destruction threatened by nuclear retaliation. On the contrary, the relationship is complex. If the anticipated consequences to a would-be initiator of nuclear war are not perceived to be horrific enough—a criterion that depends on both the damage the initiator expects to inflict on the enemy and the damage the initiator expects to suffer in return—deterrence is undermined because such consequences may not be intolerable to the initiator. This is more than a purely theoretical point. Concern with the

inadequacy of damage threatened by retaliation drove the Cold War arms race and in the future may, for example, impede significant further reductions in nuclear arsenals. It is also possible that a more comprehensive understanding of consequences could enable deeper reductions in nuclear arsenals. If more of the effects of nuclear war are included in the calculus, fewer weapons might be required to inflict whatever is perceived to constitute intolerable damage. The point of this discussion is that successful nuclear policy development, employment strategy, and crisis management depend on an understanding of consequences more nuanced than that all nuclear wars would be bad.

Similarly, the notion that we need not delve into the details of the consequences of nuclear war encourages the dangerous belief that deterrence can be maintained with a small arsenal even in the face of an adversary with a much larger arsenal. There is a substantial historical basis for this belief. For example, during the Cuban missile crisis, President Kennedy was deterred from bombing nuclear ballistic missile sites because of his fear that not all of them would be destroyed and at least one American city would suffer a nuclear attack in response. And, he was deterred from attacking the Soviet Union for the same reason. As expressed by Robert McNamara:[10]

> During a recent visit to the Soviet Union I was asked by several Russian political and scientific leaders to define nuclear parity. I replied that parity exists when each side is deterred from initiating a strategic strike by the recognition that such an attack would be followed by a retaliatory strike that would inflict unacceptable damage on the attacker. I went on to say: "I will surprise you by stating that I believe parity existed in October 1962, at the time of the Cuban missile crisis. The United States then had approximately five thousand strategic warheads, compared to the Soviet's three hundred. Despite an advantage of seventeen to one in our favor, President Kennedy and I were deterred from even considering a nuclear attack on the USSR by the knowledge that, although such a strike would destroy the Soviet Union, tens of their weapons would survive to be launched against the United States. These would kill millions of Americans. No responsible political leader would expose his nation to such a catastrophe."

Notwithstanding recent increases in China's nuclear arsenal, another example is provided by China's relatively small nuclear deterrent, which is justified on the basis that threatening only a few American cities is enough to deter US nuclear attack and prevent US intimidation.

However, there are historical counterexamples, as well. As McNamara noted, during the Cuban missile crisis the United States enjoyed an overwhelming nuclear superiority over the Soviet Union. That superiority was important in motivating the Soviet Union to place missiles in Cuba in the first place and also important in inducing the Soviets to ultimately remove their missiles.[11] As another example, in the Soviet–China border dispute of 1969, Mao Zedong had little confidence in the ability of China's small nuclear arsenal to deter a preemptive Soviet nuclear attack.[12] Finally, recognition of the importance of maintaining a *balance* of terror, even at absurdly high levels, helps explain the otherwise inexplicable buildup of US and Soviet arsenals to staggering heights during the Cold War. Again we see that reality is more complex than simplistic notions about deterrence and the consequences of nuclear war might suggest.

One might also argue that beyond a certain point, the law of diminishing returns will apply. So, while the first ten million fatalities might ruin your day, the next ten million would make it only slightly worse. That is, beyond a certain level of destructiveness, deterrence will not be further enhanced by threatening even greater damage; we can stop concerning ourselves with a careful assessment of consequences beyond that level. Unfortunately, even this does not seem to be a valid conclusion. As nuclear war increases in destructiveness, and depending on the nature of the targeting, the possibility of an environmental catastrophe increases. Ground bursts of sufficient yield can loft smoke and soot created by fires ignited by nuclear thermal radiation into the stratosphere, where they can remain for years, circulate around the globe, and attenuate sunlight. The reduced level of sunlight penetrating to the surface can significantly reduce surface temperatures for long periods of time, thereby harming agriculture on a global scale. While not completely understood and still the subject of some debate among scientists, this phenomenon could cause casualties far in excess of those caused by the more direct nuclear effects. So, rather than diminishing in incremental destructiveness as nuclear wars get larger, they can actually increase as they traverse the domain where global climatic effects are produced. Of course, the deterrent effects and

policy implications of this depend on awareness of this phenomenon; it is not clear that official policies reflect such awareness.

In summary, a nuanced and comprehensive understanding of the consequence dimension of the risk of deterrence failure is essential but not easy. Target damage, the primary basis for evaluating nuclear war plans, provides only a very small piece of the knowledge we need. A comprehensive consequence assessment that includes the broader psychological and societal impacts of nuclear weapon use is beyond the reach of current analytic capabilities. Simplistic arguments that attempt to circumvent the need to assess these broader consequences can be dangerous. Finally, we cannot redress current deficiencies without a rejuvenated nuclear weapons effects enterprise.

Conclusions and Recommendations

The introduction to this study concluded by posing two questions: Is a risk assessment of deterrence failure worth pursuing? And, if so, what is the most promising path forward? This final section provides responses to these questions.[13]

Is an assessment of the risk of deterrence failure worth pursuing? Note that this question is not the same as asking whether the risk of deterrence failure is worth knowing. We should not presume that an attempted assessment will be successful. The main argument for trying boils down to this: Nuclear deterrence is a high-stakes strategy, gambling with hundreds of millions of lives and perhaps even with the survival of civilization. Prudence dictates doing all we can to reduce the risks associated with that strategy. Assessing those risks is the first step toward this end.

Risk assessment is a prerequisite of risk management. Without assessing the risks of deterrence failure, we are flying blind with respect to the need to reduce them. If a risk assessment finds risks are "unacceptably" high, we can more vigorously pursue risk reduction policies and programs. If significant risk abatement is not feasible, we might consider giving greater consideration to developing long-term alternatives to deterrence as our central strategy for preventing nuclear war and brinkmanship for managing nuclear crises.

If the assessment fails to come to any conclusion at all on the risk of deterrence failure, that too could be a useful result. While decisions

and policies would undoubtedly continue to be made based on intuitive assessments of risk, it would perhaps help recalibrate the confidence we have in our intuition and motivate further research. Finally, in terms of further research, a well-executed risk assessment could set an appropriately high bar for the quality of such work, a welcome departure from many analyses conducted to date.

The arguments against pursuing an assessment of the risk of deterrence failure are perhaps more straightforward, but not as persuasive. First, *a risk assessment will cost money;* $10 million is not an unreasonable estimate for a comprehensive assessment using and improving state-of-the-art methods. Even in this era of high deficits and high national debt, this is clearly affordable and could result in significant savings if we develop a clearer understanding of the requirements for deterrence.

A more compelling counterargument is that it just cannot be done. *We are simply not able to develop a credible quantitative estimate of the risk of deterrence failure.* This presumption may or may not be valid, but it misses the broader picture. Related alternative goals that might be more achievable could result in useful policy-relevant insights. Such goals include:

- **Assessing the utility of determining the risk of deterrence failure.** How might the results inform policy-making? What form of answer would be useful for what decisions?
- **Assessing whether the question is researchable.** Can we come up with an analytic approach to the question? What is a practical research path forward? Our work has only scratched the surface on this.
- **Evaluating past analyses,** which could restore some balance to the current uncritical citation of such analyses in the literature.
- **Understanding risk perceptions,** including the question of how to judge risk acceptability. How much risk is too much and how should we decide?
- **Evaluating risk management policies** without an actual assessment of the risk of deterrence failure. Are we doing all we should to reduce risk?
- **Assessing future risk relative to past risk.** Is the risk of nuclear use increasing or decreasing? Is it larger or smaller than in past times?

- **Providing a lower bound on the risk,** which solves the objection to probabilistic risk assessment that it cannot possibly identify all paths to nuclear use.
- **Making qualitative judgments** without trying to quantify the risk.

While quantification of risk may represent the holy grail of risk assessment, accomplishing—or even just making progress toward—these objectives would be extremely useful and would be important steps toward this ultimate goal.

A third argument against performing a risk assessment is that we do not want to know the risk of deterrence failure. A determination of higher than "acceptable" risk could undermine deterrence with no practical alternative available. This position is not irrational but does not consider the more constructive reaction of pursuing a less-riskier variant of deterrence.

If a risk assessment of deterrence failure is worth pursuing, why hasn't it been done? The primary reason lies in the nature of the federal bureaucracy. First, there has been no leadership push for analysis on this question. In fact, independent analysis can be a dangerous thing to political leaders; one can never be sure how it will turn out. Analysis is too often just another weapon in the arsenal to promote an agenda, useful only to the extent it promotes preconceived views. Moreover, political leaders, in general, see nothing wrong with intuition, or at least *their* intuition. And the propensity to surround themselves with like-thinking advisers only reinforces their intuition.

Second, no single government agency—including the Departments of Defense and State—has purview over this question. It is not explicitly in the charters of the US Strategic Command, the Office of Net Assessment, the Defense Threat Reduction Agency, or any other component of the federal government to address the risk of nuclear deterrence failure. So, why would any of these organizations choose to expend scarce resources on it? While $10 million may be a rounding error in the entire federal budget, it is not a minor portion of the discretionary research budget of many government agencies.

Finally, it requires some nontraditional thinking to even pose the question. Most government agencies are too busy fighting daily fires to think long term or out of the box. An agency runs the risk of criticism both for stepping outside its lane of responsibility and for wasting money, as well as potential ridicule for tackling such a seemingly esoteric problem. Failing to take the initiative in this area will not ruin anyone's government career.

If an assessment of the risk of deterrence failure is worth pursuing, what is the most promising path forward? While the National Academies of Science, Engineering, and Mathematics (NASEM) is conducting a study of analytic methods for assessing the risks of nuclear war and nuclear terrorism, it is not actually performing a risk assessment, nor is it well suited to do so. Follow-on studies that actually do try to assess nuclear risks conducted by other competent organizations (e.g., federally funded research and development centers or university affiliated research centers) would be a logical follow-on.

The important concept is to undertake *several* independent studies. Absent collusion, it seems unlikely they would come up with the same answer, and whatever disagreement exists will inspire further thinking and evaluation. Several independent studies will also reduce the risk of mindless recitation of the results of a single study, as has been the case with the Lugar survey.

As an alternative to a mega-study or mega-studies, what could be even more worthwhile over the long run is to make this subject a more respectable topic for academic and government-sponsored research. Not many scholars have pursued it, and not many government or private funding sources have recognized its importance. Smaller studies tackling portions of the problem could result in a significant improvement in understanding and could pave the way, eventually, for a mega-study. Smaller studies also enable foundation support and even individual academics to pursue research without any external support at all.

One mechanism to kick-start interest in such studies is a community workshop that involves academics, think tanks, government agencies, and foundations. Ideally, multiple disciplines would be included, including physical and social scientists; political scientists, international relations experts, and nuclear strategy analysts; weapon designers, weapons effects experts and weapons operations experts; and experts in the methodologies addressed in this study and others. Beyond providing a forum for exchanging ideas, one objective could be to craft the outlines of a research agenda.

Such a research agenda might be usefully divided into studies that could be undertaken using current methods and advancement in methodologies that would enable future studies. The former category includes:

- Selected historical case studies and studies looking across the broad spectrum of close calls designed to draw lessons for assessing past and future risks

- A state-of-the-art expert elicitation to assess contemporary perspectives on current and future nuclear risks
- Development of a taxonomy of alternative paths to nuclear use
- A comprehensive assessment of the physical and social consequences of various scenarios of nuclear use

Given the consequences of even the smallest nuclear war and given the challenges of assessing nuclear risks, taking prudent steps to reduce the risk of nuclear war should not await a definitive risk assessment. Thus, many research topics will naturally blend risk assessment and risk management. While there are myriad potential topics for such studies, one example of particular importance relates to the three-quarters-of-a-century-long tradition of nonuse of nuclear weapons. We do not fully understand the dynamics of this unexpected phenomenon, nor the extent to which it has been the result of wise policy or good fortune, nor its cultural dependencies, nor how robust or fragile it is, nor how we can nurture its continuance. As Thomas Schelling argues in his Nobel Prize lecture in 2005:[15]

> This attitude, or convention, or tradition, that took root and grew over these past five decades, is an asset to be treasured. It is not guaranteed to survive; and some possessors or potential possessors of nuclear weapons may not share the convention. How to preserve this inhibition, what kinds of policies or activities may threaten it, how the inhibition may be broken or dissolved, and what institutional arrangements may support or weaken it, deserves serious attention. How the inhibition arose, whether it was inevitable, whether it was the result of careful design, whether luck was involved, and whether we should assess it as robust or vulnerable in the coming decades, is worth examining.

Research to advance methodology development has significant potential over the longer term. In this area, three advances would be particularly important:

1. An improved ability to model human behaviors and estimate probabilities in probabilistic risk assessments
2. A complex model of the nuclear deterrence system
3. Improved methods to assess social, economic, and political effects of nuclear weapon use

In addition, preliminary examination of myriad approaches to assessing the risk of deterrence failure not examined in this study should be undertaken.

A Final Thought

While we have tried to make the case for structured multidisciplinary analysis as an improvement over intuition in assessing the risk of nuclear deterrence failure, we have also come to appreciate that as much depends on the quality of the analysis as on the methodology employed. Too much of what little analysis on this topic has been produced is arguably deficient. In particular, many analyses are plagued by lack of transparency in assumptions or reasoning, inadequate treatment of uncertainties, failure to apply state-of-the-art elicitation techniques, or all these. Until such problems are addressed, it should not be surprising that analyses of the risk of nuclear deterrence failure continue to change few minds.

Notes

1. Peter Vincent Pry, *War Scare: Russia and America on the Nuclear Brink* (Westport, CT: Praeger Publishers, 1999).
2. Michael S. Gerson, *The Sino-Soviet Border Conflict: Deterrence, Escalation, and the Threat of Nuclear War in 1969* (Alexandria, VA: Center for Naval Analyses, 2010), https://www.cna.org/cna_files/pdf/d0022974.a2.pdf. The author of this chapter was involved in defining the objectives for Gerson's research, which was sponsored by the Defense Threat Reduction Agency's Advanced Systems and Concepts Office.
3. Richard Neustadt and Ernest May, *Thinking in Time: The Uses of History for Decision-Makers* (New York: Free Press, 1986), 58–62. See also Yuen Foong Khong, *Analogies at War: Korea, Munich, Dien Bien Phu, and the Vietnam Decisions of 1965* (Princeton, NJ: Princeton University Press, 1992).
4. George E. Apostolakis, "How Useful Is Quantitative Risk Assessment?" *Risk Analysis* 24, no. 3 (2004): 515–520, https://doi.org/10.1111/j.0272-4332.2004.00455.x. See also Elisabeth Pate-Cornell and Robin Dillon, "Probabilistic Risk Analysis for the NASA Space Shuttle: A Brief History and Current Work," *Reliability Engineering and System Safety* 74, no. 3 (2001): 345–352, https://doi.org/10.1016/S0951-8320(01)00081-3; and J. S. Wu and G. E. Apostolakis, "Experience with Probabilistic Risk Assessment in the Nuclear Power Industry," *Journal of Hazardous Materials* 29, no. 3 (1992): 313–345, https://doi.org/10.1016/0304-3894(92)85040-8.

5. Committee on Methodological Improvements to the Department of Homeland Security's Biological Agent Risk Analysis; Board on Mathematical Sciences and Their Applications; Division on Engineering and Physical Sciences; Board on Life Sciences; Division on Earth and Life Studies; National Research Council, *Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change* (Washington, DC: National Academy Press, 2008), https://doi.org/10.17226/12206.
6. Mary A. Meyer and Jane M. Booker, *Eliciting and Analyzing Expert Judgment: A Practical Guide* (Philadelphia: Society for Industrial and Applied Mathematics, 2001), 44.
7. Meyer and Booker, *Eliciting and Analyzing Expert Judgment*, 44.
8. Meyer and Booker, *Eliciting and Analyzing Expert Judgment*, 44.
9. George Box, Professor Emeritus of Statistics at the University of Wisconsin, said this of models. See George E. P. Box and Norman R. Draper, *Empirical Model-Building and Response Surfaces* (New York: John Wiley & Sons, 1987), 424.
10. Robert McNamara, *Blundering into Disaster: Surviving the First Century of the Nuclear Age* (New York: Pantheon Books, 1986), 44–45.
11. Not all historians agree on the latter point, but one need only imagine a counterfactual replay of the crisis in which the Soviet Union, rather than the United States, enjoyed an overwhelming advantage in nuclear weapons. Could Kennedy have credibly delivered his threat that "it shall be the policy of this Nation to regard any nuclear missile launched from Cuba against any nation in the Western Hemisphere as an attack by the Soviet Union on the United States, requiring a full retaliatory response on the Soviet Union"? Even more to the point, the emplacement of missiles in Cuba in the first place might not have been necessary from the Soviet perspective had the nuclear balance favored them at the time.
12. Gerson, *The Sino-Soviet Border Conflict*.
13. Several of the thoughts presented in response to these two questions were articulated earlier by Hellman and were refined through further private conversations between Hellman and the author. See Martin E. Hellman, "Risk Analysis of Nuclear Deterrence," *The Bent of Tau Beta Pi* 99, no. 2 (2008): 14–22, https://ee.stanford.edu/~hellman/publications/74.pdf.
14. Hellman, "Risk Analysis of Nuclear Deterrence," 18.
15. Thomas C. Schelling, "An Astonishing Sixty Years: The Legacy of Hiroshima," Nobel Prize Lecture (Stockholm: The Nobel Foundation, 2005), https://www.nobelprize.org/prizes/economic-sciences/2005/schelling/lecture/.