Chapter 1
Framing the Questions

James Scouras

What are the risks of nuclear war in all its potential manifestations? This is not an easy question to answer, and I do not propose to answer it here. Rather, the more tractable question is whether the process of studying it could yield policy-relevant insights even if it is unlikely to lead to a precise determination of the actual risks of nuclear weapons use. In this chapter, I summarize the current state of analysis regarding the likelihood of nuclear war, focusing on The Lugar Survey on Proliferation Threats and Responses*, the* Bulletin of the Atomic Scientists' *Doomsday Clock, and a sampling of analysts' estimations of the likelihood of interstate nuclear war and nuclear terrorism. These estimations differ widely and are all of questionable validity because they are either fundamentally intuitive or based on very simple—even simplistic—analyses. Can we improve on this state of analysis by using more structured and more comprehensive approaches to provide a sounder basis for policies that will inevitably be based on imperfect analyses of the likelihood of nuclear war?*

What are the risks of the use of nuclear weapons? Through what paths might these risks arise? How best might these risks be reduced? Throughout the Cold War, American strategists seemingly had relatively clear answers to these questions, or at least a clear understanding of the policies we should follow to mitigate nuclear risks. During the near half-century between the end of World War II and the fall of the Berlin Wall, the United States and its allies were focused on the existential[1] threat posed by the Soviet Union. Broadly speaking, the overarching US national security policy was containment of Soviet expansionism, and the cornerstone of US strategy for preventing such expansion by military means was deterrence. Deterrence threatened "unacceptable consequences," code for nuclear devastation, to the Soviet Union should it attack the United States or its allies.

The capability to inflict unacceptable consequences was embodied in a nuclear weapons triad of intercontinental ballistic missiles, submarine-launched ballistic missiles, and long-range bombers. The adequacy of this capability was evaluated in the context of what was viewed as the most stressing scenario: a massive surprise attack on US nuclear forces and associated command and control. The redundancy embedded in both the US and Soviet nuclear triads, and the inability of either side to strike preemptively and simultaneously against all three legs of the other side's nuclear triad with confidence of success, even in the context of this so-called "bolt from the blue" scenario, was thought to provide robust deterrence for both sides and to nearly eliminate the incentives to be the first to use a nuclear weapon. Neither side would have the temptation to strike first because massive retaliation was virtually certain; nor did either side have to fear a disabling first strike by the other. The residual risk of nuclear weapons use, in this view, was best reduced through arms control agreements that further limited the incentives for a nuclear first strike, through reduced expenditures and dangers associated with an otherwise unconstrained nuclear arms race, and perhaps eventually through missile defenses should such defenses become technologically and economically feasible.

In contrast, the unanticipated and abrupt end of the Cold War initiated a period of additional uncertainty about the most likely paths to the use of a nuclear weapon and the best means of addressing them. One source of this confusion is the unresolved question of the extent to which Russia still did, or might once again, pose a mortal threat to the United States. In the decade-long afterglow following the end of the Cold War, the nuclear threat from Russia was largely dismissed as "Cold War thinking." This attitude was reinforced by the terrorist attacks of September 11, 2001, after which all things nuclear (except nuclear terrorism) took a back seat in national security planning. By the time of this writing, some two decades after 9/11, Russia has reemerged as a nuclear threat to be taken very seriously. Russia remains the one country other than the United States with a nuclear triad that under the latest arms control agreement will still have more than 1500 nuclear weapons of global range. Yet it is also true that the mutual hostility and mistrust that characterized much of the US–Soviet relationship are today far reduced, and the robustness of mutual deterrence that held during the Cold War still appears to apply.

In any event, the risk of nuclear war with Russia has clearly receded relative to the Cold War in terms of both its likelihood and potential consequences. As a result, there is a tendency to discount the residual Russian threat relative to other threats that appear to be more immediate, growing, or more likely to result in nuclear use, even if their consequences might be orders of magnitude less severe than those of an unconstrained nuclear war with Russia. The most significant examples of such threats include North Korea's nuclear and missile program, under which a nascent nuclear deterrent has been established, and Iran's uranium enrichment program, which by its scale and nature appears to be aimed at developing, or at least having the capability to develop, a nuclear weapon. Cascading regional proliferation, especially if Iran becomes a nuclear state, is not implausible. Meanwhile, China continues to increase its nuclear capabilities, extend the range of its missiles, and diversify its means of delivering nuclear weapons. India tested a nuclear weapon in 1998, after having foregone any additional tests since its first in 1974, and Pakistan followed suit with its first nuclear test shortly thereafter; both states are considered to have dozens of nuclear weapons or more. Israel is widely believed to possess scores, if not hundreds, of nuclear weapons, and many more states have the technological ability to produce nuclear weapons should they decide to do so. The present number of nuclear weapons states is not as large as President Kennedy and others predicted in the early 1960s, but it has grown, and each additional nuclear power, including long-term US allies Britain and France, presents an added set of risks that nuclear weapons will be used someday against someone. One singular concern is the possibility that terrorist organizations willing to carry out mass-casualty attacks will eventually get their hands on a nuclear weapon by buying, stealing, or building one. The present and future number of such organizations and the likelihood of their obtaining a nuclear weapon are even harder to assess than the future number of nuclear weapons states.

In retrospect, the apparent simplicity and robustness of nuclear deterrence during the Cold War were neither as simple nor as robust as they seemed at the time. The gradual release of historical evidence has made clear that the actual risks of nuclear weapons use during the Cold War, and the most likely paths through which nuclear use could have been realized, were quite different from the scenario of large-scale and intentional use of nuclear weapons that preoccupied American and Soviet leaders and

national security analysts. The considerable number of close calls, accidents, incidents, misunderstandings, and false alarms that we now know arose during the Cold War were arguably more likely to have resulted in the use of nuclear weapons than the intentional calculation that the use of such weapons could advance some strategic purpose (and presumably there were additional close calls that are not publicly known). Indeed, perhaps the most serious incident was revealed only in 2002: during the Cuban missile crisis, Captain Valentin Savitsky, commander of a Soviet submarine, reportedly ordered his crew to prepare to launch a nuclear-armed torpedo against the American ships that were dropping depth charges to force his submarine to the surface. Fortunately, Soviet procedures required the consent of three top officers on the submarine for a nuclear weapon to be used, and another senior officer, Vasili Arkhipov, succeeded in convincing Savitsky to surface for orders from Moscow instead of launching a nuclear-armed torpedo without higher authorization.[2]

The divergence between contemporary impressions of nuclear risks and the accumulating historical evidence on actual close calls regarding nuclear weapons warrants caution in the assessment of current risks and humility in estimating future ones. Yet while precise and confident estimates of nuclear risks are not possible, the task of assessing and addressing the most pressing risks of nuclear weapons use is too important to forgo.

Objective

The nuclear threat space is clearly more complex today than during the Cold War. Interrelationships, obvious and obscure, abound among the myriad facets of the nuclear threat and policies intended to address them. A systems perspective, currently lacking, would help clarify how various aspects of nuclear policy affect all elements of the nuclear threat and thereby reduce the likelihood of unintended consequences. But even more basically, we need to establish policies informed by the risks of the various dimensions of the nuclear threat. What are the risks of nuclear use in all its potential manifestations?

This is clearly a difficult question to answer. In fact, it might not be answerable at all with a useful degree of certainty or consensus, and I do not propose to answer it here. Rather, my more limited purpose is to address whether the question is or is not analytically tractable and whether

the process of studying it would yield policy-relevant insights even if it is unlikely to lead to a precise determination of the actual risks of nuclear weapon use. As President Eisenhower often said, "Plans are worthless, but planning is everything." In other words, although we cannot fully anticipate our adversaries or the future, the attempt to do so will leave us in a better position than if we failed to plan at all. More specifically with regard to the question of nuclear risks, what analytical approaches have been or could be utilized, and what are their prospects for providing a degree of enlightenment on the risk of nuclear use?

Risk Terminology and Analysis

Risk is exposure to danger due to the likelihood and consequences of an adverse event. In our case the adverse event is nuclear weapon use, which is defined as the detonation of one or more nuclear weapons, except for nuclear weapons tests, whether intentionally or accidentally, anywhere in the world. The reason for such an expansive definition is that any nuclear use could directly or indirectly cascade to involve the major nuclear states.

Likelihood can be described in qualitative terms (e.g., unlikely, highly likely, a remote possibility) or quantitatively, such as in probabilistic terms. To be meaningful, a time frame must be specified (e.g., "There is a moderate likelihood of nuclear use within the next ten years.") In some risk assessments, frequency is used to portray likelihood ("We can expect two attacks over the course of the next decade."). However, this is not appropriate for nuclear attacks, because the original nuclear use and the reaction to it can be expected to significantly affect the likelihood of a subsequent use.

Consequences include fatalities, injuries, physical and economic damage, social and psychological impacts, and all other forms of harm. They can be immediate or can unfold over the course of decades. As with likelihood, a time period should be specified. Important, but often overlooked, consequences include those that would result from the reaction to nuclear weapon use.[3] Because its consequences would be extreme, even the remote likelihood of nuclear weapon use may well motivate policy changes because remote possibilities can accumulate to worrisome levels when aggregated over the long term. Put another way, if there is a nonzero constant risk of nuclear use each year, given enough time, the probability

of eventual nuclear use will approach 100 percent no matter how small the risk is in any given year.

For both likelihood and consequences, it is important that an uncertainty be associated with any estimation. Research in psychology suggests that humans are not very good at estimating "confidence levels" for their estimations.[4] For example, people are generally poor at tasks like the following: "Set the upper and lower bounds on the probability that a nuclear weapon will be used in the next ten years, in such a way as to be 90 percent confident that the actual probability will fall between your lower and upper bounds." Yet if we fail to press experts to associate a confidence level or uncertainty with their estimates, we can easily fall into the trap of assuming that the uncertainty is negligible.

Finally, the common practice of multiplying likelihood and consequences, which would result in an expected risk, is inappropriate for characterizing the risk of nuclear war. There are policy-relevant differences between the combination of low likelihood and high consequences (perhaps interstate nuclear war) and the combination of high likelihood and low consequences (perhaps terrorist nuclear weapon use). This critical distinction is lost when the product is used.

The Influence of the Perceived Risk of Deterrence Failure on Policy

Detailed consideration of the likelihood and consequences of nuclear war is not usually explicit in developing national security strategy. Yet implicit assumptions on these questions have a strong influence on nuclear policy, and explicit generalizations of them have been invoked to justify major new directions in policy. The entire nuclear arms control enterprise—from the "hotline" memorandum of understanding through the Anti-Ballistic Missile Treaty and Strategic Arms Limitation Treaty (SALT), the Intermediate Nuclear Forces (INF) Treaty, the Strategic Arms Reduction Treaty (START), the Strategic Offensive Reductions Treaty (SORT), and the current New START Treaty—was motivated principally by fear of nuclear war. In the 1950s the dominant theoretical concern involved a disarming surprise attack, but fear of nuclear war was made all too tangible by the Cuban missile crisis of 1962 and has been reinforced throughout the Cold War by the nuclear arms race and nuclear posturing.

Similarly, arguments for national missile defenses depend in no small part on the judgment that deterrence is unreliable. As expressed by President Reagan in his oft-quoted Strategic Defense Initiative speech of 1983:[5]

> Tonight . . . I am directing a comprehensive and intensive effort to define a long-term research and development program to begin to achieve our ultimate goal of eliminating the threat posed by strategic nuclear missiles. This could pave the way for arms control measures to eliminate the weapons themselves. We seek neither military superiority nor political advantage. *Our only purpose—one all people share—is to search for ways to reduce the danger of nuclear war.* [Emphasis added.]

President George W. Bush invoked the inadequacy of deterrence and the consequences of nuclear use by "rogue" states and terrorists to justify preemptive attack as a critical element of national security strategy:[6]

> It has taken almost a decade for us to comprehend the true nature of this new threat. *Given the goals of rogue states and terrorists, the United States can no longer solely rely on a reactive posture as we have in the past.* [Emphasis added.] The inability to deter a potential attacker, the immediacy of today's threats, and the magnitude of potential harm that could be caused by our adversaries' choice of weapons, do not permit that option. We cannot let our enemies strike first.

More recently, the prospect of nuclear use motivated President Obama's call for a nuclear-free world:[7]

> Today, the Cold War has disappeared but thousands of those weapons have not. In a strange turn of history, the threat of global nuclear war has gone down, but the risk of a nuclear attack has gone up . . . So today, I state clearly and with conviction America's commitment to seek the peace and security of a world without nuclear weapons.

Clearly, quite different policies have been motivated by the concern that deterrence might fail. However, assertions that deterrence cannot be relied on are based on intuition and limited perspectives rather than syntheses of

the broadest expertise and most appropriate analytic methods that can be brought to bear.

The State of Analysis

Consider the current state of analysis. In 2005 the office of Senator Richard Lugar published *The Lugar Survey on Proliferation Threats and Responses* (hereinafter, the Lugar survey),[8] which addresses the risk of nuclear use and has been widely cited on the internet and in the academic literature. Among the questions asked in the survey was, "What is the probability (expressed as a percentage) of an attack involving a nuclear explosion occurring somewhere in the world in the next ten years?" The distribution of replies from seventy-nine respondents is shown in Figure 1.1.

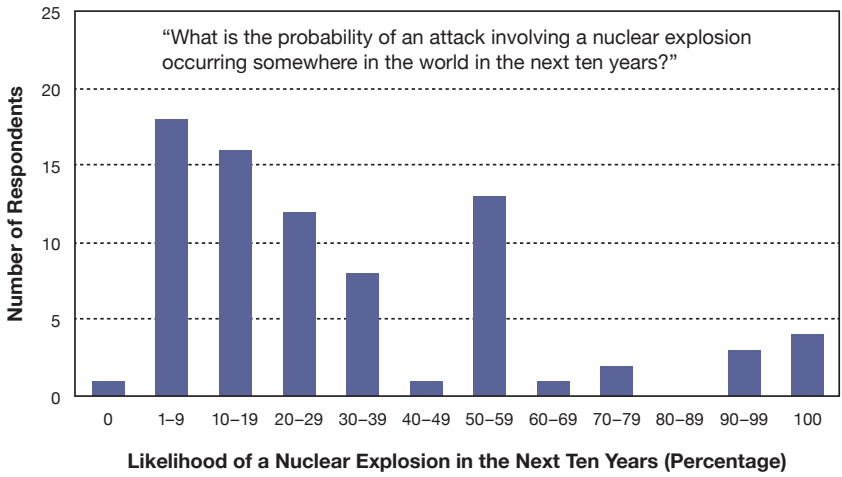

Figure 1.1. The Lugar survey, question 5.

What is most striking about Figure 1.1 is the breadth of opinion, which spans the full spectrum from 0 to 100 percent. From a classical statistics perspective, the true probability lies in only one bin. The fact that most experts' answers missed that value, whichever bin it lies in, means that most experts must necessarily be wrong. There are a number of possible explanations for this. One reason for the wide variation could be the lack of control of biases in the elicitation of the answers. Without bias control, experts can interpret and think differently about how to answer the

question, resulting in wide variability. Even if biases are controlled, wide dispersion can still occur because of high uncertainty in the current state of knowledge. In any event, the most significant conclusion to be drawn from Figure 1.1 is that there is no consensus on the answer to the question. In contrast, the Lugar survey report highlights the mean (29 percent) of these data as the most relevant finding. If it had also reported the standard deviation (which is approximately 26 percent) with this mean, then the high variability in Figure 1.1 would have been more apparent.

In other respects as well, the Lugar survey did not make use of best practices in elicitation and analysis.[9] While each survey respondent was an expert in some aspect of nuclear policy, arguably no single person is truly an expert on *all* the factors that must be considered when answering broadly phrased questions such as that depicted in Figure 1.1.[10] Additionally, the survey provides no information about the experts' assumptions, reasoning, and uncertainties. Such information could, for example, be useful in understanding the apparently anomalous peak at 50–59 percent. The cumulative impact of these and other deficiencies is that the survey falls short of what could be achieved through a survey using best practices in expert elicitation. Yet references to the Lugar survey are almost uniformly uncritical, even in the academic literature, and policy advocates have used its results to argue for important decisions. Clearly, a more scientific survey could be conducted that would improve on the reliability of the Lugar survey. Nevertheless, the fact that the survey was undertaken and that it was extensively cited demonstrate that the question of the likelihood of deterrence failure is relevant to policy-makers, analysts, and the public.

Another exercise in characterizing the likelihood of nuclear war has been ongoing since 1947, when the Doomsday Clock first appeared on the cover of the *Bulletin of Atomic Scientists*.[11] The setting of the clock is intended to represent how close the world is to nuclear war, metaphorically midnight. The clock was originally set at seven minutes to midnight and has been reset periodically every several (one to seven) years. As shown in Figure 1.2, the time of greatest danger, two minutes to midnight, was set in 1953 after the US and Soviet hydrogen bomb tests, while the time of least danger, seventeen minutes to midnight, was set in 1991 after the START Treaty was signed and unilateral initiatives on both sides removed many nuclear weapons from "hair-trigger" alert.[12]

There are multiple problems with taking the clock seriously as an assessment of the likelihood of nuclear war. There could be motives in setting the clock beyond accurately characterizing the nuclear threat, such as to promote certain policies, especially with respect to arms control treaties, or simply to draw attention to the *Bulletin of the Atomic Scientists*. The process by which the clock is set is obscure, although brief summaries of the reasons for changing the clock's setting have been provided.[13] No attempt has been made to define the clock's scale, which is almost certainly nonlinear. Does ten minutes to midnight indicate half the probability of five minutes to midnight? And finally, the clock is unable to reflect the risks associated with short-duration, high-risk episodes, such as the Cuban missile crisis of 1962 and the coup attempt against Gorbachev in August 1991.[14] Ironically, the former occurred during a period of reducing risk, according to Figure 1.2, and the latter occurred during the period of least risk.

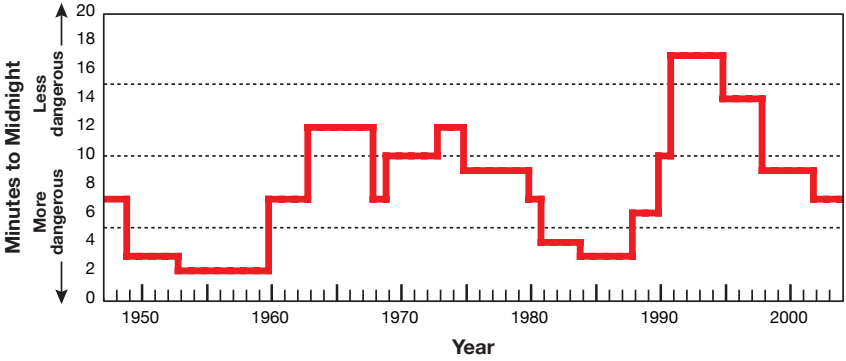

Figure 1.2. The Doomsday Clock, 1947–2004. The clock indicates then current perspectives of the *Bulletin of the Atomic Scientists* on the dangers of nuclear war. Since 2007, dangers associated with climate change and developments in the life sciences have been added.

Notwithstanding these points, the Doomsday Clock does seem to have captured the broad trends in the nuclear threat as it derives from the international political climate. Gaining a better understanding of the processes by which the clock has been set could prove useful in developing more scientific approaches. Unfortunately, the clock's future utility as an indicator of the risk of nuclear war has been diminished since 2007 by the inclusion of climate change and developments in the life sciences as additional harbingers of doomsday.

Several individuals have also estimated the likelihood of interstate nuclear war or nuclear terrorism. These estimates are summarized in Table 1.1. Most are subjective judgments (Kennedy,[15] Bundy,[16] Allison,[17] Perry,[18] Albright,[19] and Garwin[20]) without a formal underlying analysis, while others are based on a specific analysis (Hellman,[21] Bunn,[22] and Mueller[23]).

Arguably, the most compelling assessments are those of crisis managers who experienced a nuclear close call firsthand: President John F. Kennedy and his national security advisor, McGeorge Bundy. Not long after the Cuban missile crisis, Kennedy told Ted Sorenson, special counsel to the president, that he believed during the crisis the chances that the Soviets would go to war were between one in three and even, while Bundy, reflecting twenty-six years after the crisis, came to the dramatically lower estimate of up to one in one hundred. Of course, the crisis occurred almost a half-century ago, and even with the additional information now available it is hard to estimate its risks retrospectively. For example, depending on one's interpretation of the probabilities associated with the Soviet submarine incident discussed above, and the risks one should attach to other "close-call" incidents during the Cuban crisis,[24] one could argue for either Kennedy's estimate or Bundy's. Moreover, neither Kennedy nor Bundy knew at the time they made their estimates that a Soviet submarine had come close to launching a nuclear torpedo, but they could have imagined this and other scenarios as part of their risk estimates, so it is unclear whether either of them would have raised or lowered their estimates if they had known at the time of their estimates everything we know now. Of course, beyond the question of what the actual risk was at the time of the Cuban crisis is the problem of the relevance of that information to the assessment of future risks.

Recently, Martin Hellman assessed the risk of a future "Cuban missile-type" crisis that results in nuclear use as between two in one thousand and one in one hundred per year. Note that this is only one of three estimates in Table 1.1 that provides a range, a useful approach to addressing uncertainty. Hellman also points to a dearth of analyses of the risk of deterrence failure and proposes that "several prestigious scientific and engineering bodies undertake serious studies to estimate its failure rate."[25]

Table 1.1. Individual estimates of the probability of nuclear war

	Question	Estimate	Author	Year
War	Probability that the Cuban missile crisis could have escalated to (nuclear) war?	Between 1 in 3 and even (war)	John F. Kennedy	1962
		As large as 1 in 100 (nuclear war)	McGeorge Bundy	1988
	Probability of a future Cuban missile-type crisis that results in at least one nuclear weapon being used?	2 in 1,000 to 1 in 100 per year	Martin Hellman	2008
Terrorism	Probability that terrorists will detonate a nuclear bomb?	More likely than not (on America)	Graham Allison	2004
		50–50 odds within the next decade	William Perry	2004
		Less than 1 percent in the next 10 years	David Albright	2005
		29 percent probability within the next decade	Matthew Bunn	2007
		10–20 percent per year against a US or European city	Richard Garwin	2007
		Less than 1 in 1,000,000 (per attempt)	John Mueller	2008

Not surprisingly, a number of post-2001 estimates have focused on the probability of nuclear use by terrorist organizations. Of the subjective estimations (i.e., those not based on a specific analysis), Richard Garwin's estimate of 10–20 percent per year against a US or European city is the highest; it equates to a probability of approximately up to 90 percent within a decade assuming that the probability remains constant over that period. In the middle of the range of subjective estimates are Graham Allison and William Perry, who independently judge this probability to be 50 percent

within a decade. At the low end is David Albright, who estimates less than 1 percent over ten years. These subjective assessments span almost the complete range of possibility from near 0 to 90 percent.

Two estimates in Table 1.1 are based on specific analyses. Matthew Bunn estimates 29 percent within the next decade and John Mueller estimates less than one in one million per attempt. This large difference in estimates is not an encouraging indicator that analysis will facilitate convergence on a consensus estimate, but at least it provides valuable insights into the basis for each estimate.

In summary, the principal insights I take from the estimates in Table 1.1 are the same as for the Lugar survey: (1) they differ widely, and (2) they are all of questionable validity because they do differ widely and because they are fundamentally either intuitive or based on simple analysis. Also, subjective judgments appear to gravitate to either 1 percent or 50 percent as an estimate, which suggests that the resolution of human intuition is relatively coarse on this question.

Study Scope

Based on this review of the current state of analysis, two alternative courses of action are apparent. The first is to make the case that the risk of nuclear weapon use is so analytically intractable that even the most careful and comprehensive assessment of this risk would not be relevant to policy-making. The other option is to improve current approaches to assessing nuclear risk in order to provide a sounder basis for policies that will inevitably be based on imperfect analyses of such risks. Either course of action would represent an improvement over the current state of analytic affairs in which individual judgments are offered, usually without a clear trail of assumptions and reasoning, and simple analyses and surveys are undertaken that rely on unsound elicitation practices.

As a first step toward both of these ends, this book tackles the somewhat more modest objective of addressing whether assessing the risk of deterrence failure is *feasible*. We have examined the potential utility and limitations of four of the more promising approaches to the question of likelihood. Case studies of nuclear weapon use and historical close calls in which nuclear weapon use was contemplated or could have occurred, discussed by Andrew Bennett in chapter 2, provide a unique window into

past nuclear risk. Jane M. Booker addresses challenges and best practices for utilizing elicited expert knowledge, which underpins almost all analytic approaches, in chapter 3. Probabilistic risk assessment, which has been ever more successfully applied to complex engineered systems including those with human components, is assessed by Martin Hellman in chapter 4. Edward T. Toton analyzes the potential applicability of complex systems theory to the question of the risk of deterrence failure in chapter 5.

In chapter 6, Michael J. Frankel, George W. Ullrich, and I address the consequence dimension of the risk, focusing on the state of knowledge and tools to support anticipation of the physical consequences of nuclear weapon use. Dallas Boyd provides a complementary analysis of the intangible consequences of nuclear use in chapter 7. The challenge of integrating knowledge from these disparate approaches to both likelihood and consequences is discussed by Jane M. Booker in chapter 8. In the final chapter, I conclude with some thoughts on the fundamental questions, Is a risk assessment of deterrence failure worth pursuing, and, if so, what is the most promising path forward?

Other approaches to assessing the risk of deterrence failure also hold promise but have not been included in this study. For example, for likelihood assessments, we have not examined the utility of statistical analyses of the historical record of warfare and terrorism. Nor have we studied the potential contributions of the humanities, particularly human psychology, or the social sciences, notably organizational psychology, anthropology, and the emerging discipline of strategic culture.[26] Perhaps our work will motivate others to pursue these omissions.

Notes

1. Existential to the United States as a constitutional republic, not to the human race.
2. Aleksandr Mozgovoi, *Kubinskaya Samba Kvarteta Fokstrotov* [Cuban Samba of the Foxtrot Quartet] (Moscow: Voennyi Parad, 2002). See also William Burr and Thomas Blanton eds., "The Submarines of October: U.S. and Soviet Naval Encounters During the Cuban Missile Crisis," *National Security Archive Electronic Briefing Book No. 75* (Washington, DC: National Security Archive, GWU, October 2002), https://nsarchive2.gwu.edu/NSAEBB/NSAEBB75/; and Svetlana V. Savranskaya, "New Sources on the Role of Soviet Submarines in the Cuban Missile Crisis," *Journal of Strategic Studies* 28, no. 2 (2005): 233–259, http://dx.doi.org/10.1080/01402390500088312.

3. Dallas Boyd and James Scouras, "The Dark Matter of Terrorism," *Studies in Conflict & Terrorism* 33, no. 12 (2010): 1124–1139, https://doi.org/10.1080/10576 10X.2010.523863.
4. Mary A. Meyer and Jane M. Booker, *Eliciting and Analyzing Expert Judgment: A Practical Guide* (Philadelphia: Society for Industrial and Applied Mathematics, 2001).
5. Ronald Reagan, "Address to the Nation on Defense and National Security," March 23, 1983, http://www.reagan.utexas.edu/archives/speeches/1983/32383d. htm.
6. George W. Bush, *The National Security Strategy*, September 2002, http:// georgewbush-whitehouse.archives.gov/nsc/nss/2002/index.html.
7. Remarks by President Barack Obama, The White House Office of the Press Secretary, April 5, 2009, http://www.whitehouse.gov/the_press_office/Remarks-By-President-Barack-Obama-In-Prague-As-Delivered/.
8. Richard G. Lugar, *The Lugar Survey on Proliferation Threats and Responses* (Washington, DC: US Senate, 2005), https://irp.fas.org/threat/lugar_survey.pdf.
9. Meyer and Booker, *Eliciting and Analyzing Expert Judgment*. See also Bilal M. Ayyub, *Elicitation of Expert Opinions for Uncertainty and Risks* (Boca Raton, FL: CRC Press, 2001).
10. According to the Lugar survey report, "Many of these men and women have dedicated their professional careers to the study and practice of preventing weapons of mass destruction and materials from falling into unauthorized hands. Others have been national security leaders within their countries. As a group, they possess enormous experience in the fields of non-proliferation, counter-proliferation, diplomacy, military affairs, arms inspection, intelligence gathering, and other national security fields relevant to the questions asked." The fault of the survey is to confuse the expertise of the group as a whole, if it could be brought to consensus, with the sum of individual expertise within the group.
11. "Doomsday Clock Overview," *Bulletin of the Atomic Scientists*, http://thebulletin. org/overview.
12. "Doomsday Clock Overview," *Bulletin of the Atomic Scientists*.
13. "Timeline," *Bulletin of the Atomic Scientists*, http://www.thebulletin.org/content/ doomsday-clock/timeline.
14. Peter Vincent Pry, *War Scare: Russia and America on the Nuclear Brink* (Westport, CT: Praeger Publishers, 1999).
15. Theodore C. Sorensen, *Kennedy* (New York: Harper & Row, 1965), 705. The exact date of Kennedy's estimate is not specified in this source, but the estimate appears to apply to Kennedy's belief in the midst of the crisis. According to Sorenson, "The odds that the Soviets would go all the way to war, he [Kennedy] later said, seemed to him then 'somewhere between one out of three and even.'" Note that Kennedy's estimate refers to the likelihood of war but does not explicitly specify *nuclear* war. Nevertheless, it seems clear that if the Soviets

initiated a conventional war (in Berlin, perhaps), the likelihood of escalation to nuclear conflict was high.

16. McGeorge Bundy, *Danger and Survival: Choices about the Bomb in the First Fifty Years* (New York: Random House, 1988), 461. See also p. 453 for a discussion of Kennedy's estimate.
17. Graham Allison, *Nuclear Terrorism: The Ultimate Preventable Catastrophe,* 1st ed. (New York: Times Books/Henry Holt and Company, 2004), 15.
18. Nicholas D. Kristof, "An American Hiroshima," *New York Times*, August 11, 2004, https://www.nytimes.com/2004/08/11/opinion/an-american-hiroshima.html.
19. Corine Hegland and Greg Webb, "The Threat," *National Journal* 37, no. 16 (2005): 1138–1145, https://www.nationaljournal.com/s/420480/cover-story-threat/.
20. Federal News Service, *Hearing of the Energy and Water Subcommittee of the House Appropriations Committee*, March 29, 2007, http://www.cdi.org/PDFs/Energy and Water Subcommittee Hearing.pdf.
21. 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.
22. Matthew Bunn, "Guardians at the Gates of Hell: Estimating the Risk of Nuclear Theft and Terrorism—And Identifying the Highest-Priority Risks of Nuclear Theft" (doctoral thesis, MIT, 2007), https://scholar.harvard.edu/files/matthew_bunn/files/guardians_at_the_gates_of_hell_thesis_2007.pdf.
23. John Mueller, "The Atomic Terrorist: Assessing the Likelihood," presented at the Program on International Security Policy, Chicago, IL, January 15, 2008, https://politicalscience.osu.edu/faculty/jmueller/APSACHGO.pdf.
24. Scott D. Sagan, *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons* (Princeton, NJ: Princeton University Press, 1993).
25. See the preface for further information on this proposal.
26. On strategic culture, see Jeffrey S. Lantis, "Strategic Culture and National Security Policy," *International Studies Review* 4, no. 3 (2002): 87–113, https://doi.org/10.1111/1521-9488.t01-1-00266. See also Kerry M. Kartchner, Jeannie L. Johnson, and Jeffrey A. Larsen, *Strategic Culture and Weapons of Mass Destruction: Culturally Based Insights into Comparative National Security Policymaking* (New York: Palgrave Macmillan, 2009).