

are subject to different interpretations. Theory or fundamental principles about the behavior of nations and groups of people are inadequate and lack sufficient validation to augment the sparse historical record with authoritative information. For such data-poor problems, analysts rely heavily on knowledge from experts.

While everyone can have an opinion, not everyone is an expert. Experts are recognized by their peers as knowledgeable in a subject-matter field and qualified to solve problems and to answer questions.

Formal elicitation serves two purposes in considering the problem of assessing the risk of nuclear deterrence failure:

- 1.

Biases

Biases are a slanting, adjusting, or filtering of an expert's thinking and original knowledge due to their needs (motivation) and through cognitive processing. Biases degrade the quality of elicited knowledge through distortion. To counter these deleterious effects, formal elicitation includes bias minimization methods for monitoring and/or controlling common biases.

Table 3.1 lists names and descriptions of common biases. While names of biases may vary in different subject areas, their descriptions and effects are common across problems. For example, near-miss bias can be described as a combination of overconfidence and availability biases.

Nuclear war and deterrence are highly emotional topics, and factions exist on multiple sides of associated issues. Experts tend to place undue importance on the few facts available to them, be wishful about outcomes

is valid; therefore, anchoring bias is difficult to detect and to overcome for the deterrence failure problem.

Another anchoring bias is that humans inherently assume that others think and behave in the same way they think and behave. The close call in the NATO Able Archer 83 exercise is one such example. From the Soviet perspective, and consistent with its military doctrine, a nuclear exercise was a useful pretext for a nuclear surprise attack. Soviet leaders, assuming that US leaders think like they do, surmised that a US surprise attack could be the true purpose of Able Archer.

These biases require monitoring and understanding through formal techniques such as probing the experts for explanations, clarifications, and thought processes. Likewise, these techniques aid in distinguishing biases from expertise and experience.

Table 3.1. Common biases

	Name	Definition
Cognitive Biases	Anchoring	An expert's failure to sufficiently adjust from their first, long-held, or unchallenged impression in solving a problem—the expert anchors to first, long-held, or unchallenged impression. Sometimes this bias is explained in terms of Bayes' theorem as the failure to adjust knowledge in light of new information as much as it should

	Name	De nition
Motivational Biases	Group think	The tendency to modify knowledge and/or information so that it agrees with that of the group or of the group leader. Individuals are generally unaware that they have modi ed their thinking and responses to be in agreement. This bias stems from the human need to be accepted and respected by others. Individuals are more prone to group think if they have a strong desire to remain a member, if they are satis ed with the group, if the group is cohesive, and if they are not a natural leader in the group.
	Impression management	Resulting from social pressure, this bias occurs when the expert responds to the reactions of those not physically present. For example, the expert answers survey questions in a way that maximizes approbation either from society in the abstract or from the administrator of the elicitation in particular.
	Misinterpretation of the expert	The altering of the expert's thoughts as a result of the methods of elicitation and documentation.
	Social pressure	An effect that induces individuals to slant their responses or to silently acquiesce to the views that they believe the interviewer; their group, supervisors, organization, or peers; or society in general will accept. This altering of an individual's thoughts can take place consciously or unconsciously. The social pressure can come from those physically present or from the expert's internal evaluation of how others would interpret their responses. People's need to be loved, respected, and recognized induces them to behave in a manner that will bring af rmation. Political correctness is an example.
	Training bias	The tendency of the data gatherer, analyst, or both to misinterpret data/information from others for their own purposes (for example, choosing quotations, references, or events that suit the interviewer's purposes).
	Wishful thinking or con ict of interest	A tendency 3 140.514 Tm [(o)-3.992 /P <<.9 (4 T r)-12.2-USp

Elicitation Setting

The quality of elicited knowledge depends on the interviewer's ability to question experts about the assumptions they use, the heuristics and cues involved in their thinking, and their problem-solving processes. These details are best elicited in face-to-face elicitation sessions, making

include multiple choice, open-ended essay, continuous numerical scale, odds ratio, range of values, comparison, ranking, and likelihood. Some of these are described in the next section on structuring. Likelihood may be a concept consistent with the way many experts think, and it is general enough to encompass definitions used by specific communities of practice. In contrast, probability is only rarely appropriate to very specific communities.

a probability of 0.90 for that event to not occur. More difficult-to-detect violations of the axioms of probability include a sum of multiple mutually exclusive and exhaustive probabilities that is not 1.0 and improper estimates of conditional or dependent probabilities.

Often surveys interchange the terms *probability* and *percentage*.

Third, humans (even statisticians) are not well calibrated for estimating probabilities. As a rule, they cannot accurately express their perceived likelihood or frequency of their experiences as probabilities. For extremely rare events, poor calibration of probability estimation can be magnified. For example, it is difficult to distinguish between a probability of 0.000001 and 0.0000001. This is why in some subject areas, orders of magnitude (e.g., the Richter scale for earthquakes) are used. However, if experts are not experienced in thinking in such scales, it is difficult to teach or train them. In general, it is difficult to train experts to accurately estimate probability.

Unless an expert is used to dealing with and thinking in terms of probability, it is best to avoid asking for probability as a response. Other response modes and descriptions are advised, such as odds (betting odds), likelihoods, ratios, ranks, or other comparisons. The choice should be consistent with the expert's community of practice. At the very least, the interviewer should thoroughly define any unfamiliar response mode for the expert.

In those special cases in which probability is appropriate to characterize uncertainty, it should be noted that there are at least two modern interpretations of probability that are equally valid within its theory. The first is what most are taught as probability—the number of event occurrences divided by the total number of outcomes. This is the frequentist or relative frequency interpretation of probability. For example, the probability of drawing a red marble from a jar containing one hundred marbles of which twenty are red is $20/100 = 0.05$.

The second is the personalistic interpretation, often referred to as the Bayesian interpretation, the centerpiece of Bayesian analysis. Personalistic

occur within three months, the expert should be willing to stake \$0.90 in exchange for \$1.00 if the attack occurs within three months. If the attack occurs within three months, the expert wins the \$1.00, for a net gain of \$0.10. If the attack does not occur, the expert loses \$0.90. To prevent cheating, the expert should also be willing to make the opposite bet, where they are willing to stake \$0.10 in exchange for \$1.00 if the event does not occur. This two-sided bet is depicted in Table 3.2. In terms of betting odds, this example demonstrates odds of 9 to 1.

Table 3.2. An example of a two-sided bet

Bet	Attack Occurs, $p = 0.90$	Attack Does Not Occur, $p = 0.10$
Expert stakes \$0.90 in exchange for \$1.00 if attack does occur	Expert's net gain is \$0.10	Expert's net loss is \$0.90
Expert stakes \$0.10 in exchange for \$1.00 if attack does not occur	Expert's net loss is \$0.10	Expert's net gain is \$0.90

An expert who believes the probability of attack is 90 percent should be willing to take either side of this bet.

Regardless of whether or not an uncertainty is probabilistic, the interviewer should elicit it along with the responses to the questions asked of experts during an elicitation. The form or format for noting uncertainties should be consistent with the way the experts think and the available knowledge.

One of the recommended forms for eliciting uncertainties is to request a range of answers after eliciting the expert's response. To avoid introducing ambiguous uncertainty in the analysis of experts' ranges, it is necessary to define what the requested range represents. For example, the range could represent absolute highest and lowest values. Unless experts are familiar with percentiles (and most are not), tying range limits to percentiles (e.g., 5th and 95th) is not recommended. To minimize anchoring bias, the expert should be encouraged to consider their range in conjunction with their response, making any necessary adjustments.

The decomposition process includes specifying definitions, conditions, scenarios, assumptions, timelines, quantities, and parties involved. Usually, several preliminary questions that provide these specifications are asked to set the stage for the questions of interest. A structure or framework of the problem provides guidance on how to do the decomposition.

The decompositions and operating conditions of physical systems can be easily represented because of their structure. However, decompositions of complexities of human behaviors, timelines, or event sequences—all of which are applicable to assessing the risk of failure of deterrence—may not be so obvious or conducive to common structures such as fault trees. The nuclear deterrence failure problem currently lacks a systems perspective (and hence structure) or model, making decomposition difficult. Even establishing initial or boundary conditions may pose challenges because of all the facets and factors involved. It may be possible for experts to contemplate some specifically defined scenarios or special cases and begin decomposing the problem by using those.

Risk analysis has two aspects: likelihood and consequence. Risk studies

in chapter 1, it shows the varied responses of seventy-nine experts to the question, “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?” While this question may sound specific, the geopolitical conditions leading up to such an event were not specified, assumptions about the attacker were absent, and what constitutes an “attack” was not defined, leaving each respondent free to decide what these factors might be. The wide variety of responses suggests that different experts answered differently based on their assumptions and what they were free to specify in their thought processes (but were not asked to report). As noted in chapter 1 and in the bias subsection above, such lack of specifics provided to the experts opens the door for biases to dominate, adding to the wide dispersion seen in Figure 3.1.




Figure 3.1. The Lugar survey, question 5.

While it is important to select a diverse group of experts to ensure the state of knowledge is represented, such a dispersion of responses could also indicate that some respondents did not know how to answer because of lack of expertise so thof eo answe4.9 (, s)1.3 (u)2.2 (c)-3.1 (h a d)-20.re f 265.2acr

bias and getting experts to expand their thinking beyond their anchored views is what bias minimization elicitation is all about.

The nuclear terrorism and war literature contain some examples of decomposing the complex and ill-posed deterrence problem. Bunn, Hellman,²⁸ and Mueller²⁹ decomposed the problem into separate events for evaluation. Each provided their own problem structuring for the conditions and assumptions of the events they chose. Each then provided their own estimates of the likelihoods of these events and descriptions of how to combine or propagate those estimates to obtain the final answers.

Their analyses of their versions of the problem could be called self-elicitations. Self-elici. Sle6 (a)-306 Tw 10.5 0 00 10.10.5 0 0. onid asepage.1 (-)-

such experts. To ensure proper representation, a random sample or other

Experts' Problem Solving and Cognition

Much of the wide dispersion (a type of uncertainty often measured by a variance) in the responses in Figure 3.1 could be understood if the experts had recorded their thoughts and problem-solving processes while answering the question. These activities are part of formal elicitation design. Querying the experts about their thinking and problem-solving processes is conveniently done in a face-to-face interview. It can also be done during the feedback process to clarify responses.

Probing into cognitive and problem-solving processes is important for determining whether an expert is answering the posed question or some modified or misinterpreted version. Often experts think about conditions, assumptions, cues, and experiences and use problem-solving methods that affect their responses, but these thoughts and methods may not be recorded. Changing one or more of these could significantly change an expert's response. If the analyst does not know details about how the experts answered a question, the analyst will not be able to draw proper conclusions or resolve disagreements among experts.

A simple example illustrates the importance of eliciting cognitive and problem-solving processes. Experts A and B both respond with high likelihoods of nuclear weapon use within the next ten years. However, after eliciting their problem-solving processes, it is discovered that expert A assumes a terrorist use while expert B assumes an interstate war. Further probing reveals that expert A considers the interstate war an unlikely situation for nuclear use and expert B considers nuclear terrorism unlikely. Thus, without knowing what the experts were assuming when responding to the nuclear use question, their apparent agreement is not the correct conclusion. Experts A and B were actually providing different answers based on different assumptions and cognitive processing.

The analyst is often faced with determining the degree of dependency among experts. This is important if experts' responses need to be aggregated (e.g., reporting an average response as done in the Lugar report). Experts who are highly dependent are expressing the same knowledge and cannot be counted as independent sources. It is difficult to determine the extent of overlapping or double-counted knowledge from a group of experts. Without details about how experts arrived at their responses, dependency determination becomes untenable. Experts who solve problems by using similar

Drawing Conclusions

Usually the reasons for analysis are to summarize the elicited responses and to draw conclusions from them, often to inform decision- and policy-makers. Even though elicited expert knowledge is not a substitute for experimental, historical, or observational data, it can be analyzed and conclusions can be drawn from it. If there is ever a time when data might become available, elicited and analyzed expert knowledge can be considered a placeholder for those future data and can be compared and combined with the future data.

For highly qualitative responses, there may be little opportunity to analyze the information elicited by using statistical or data analysis methods. While qualitative knowledge can sometimes be grouped or categorized, this is subject to misinterpretation bias. If the responses are continuous numeric quantities, integers, ordinal, or categorical, then statistical analysis methods are useful for providing defensible conclusions inferred from experts' responses.

Decision-makers may be accustomed to seeing a central aggregated response from all the experts—a mean (the average of numerical values), median (the middle of the range), or mode (most frequent or common value). For example, the mean for the question in Figure 3.1 is 31 percent, which falls in the 30–39 percent bin. The median of seventy-nine values is the fortieth value, which falls in the 20–29 percent bin. The mode is the bin with the largest count, the 1–9 percent bin. Because of how these three differ, the conclusion is that these data are not distributed symmetrically around a central value. Figure 3.1 visually confirms the lopsided loading of the data in the lower percentages. The wide dispersion of responses in Figure 3.1 is summarized by the large standard deviation—an uncertainty metric for dispersion—of 28 percent. Another common measure of dispersion uncertainty is the range, which is 100 percent.

Statistical methods can be used to determine whether the experts responded uniformly across the percentage scale as might be suspected in Figure 3.1. The answer here is no; significantly fewer than expected experts

discussion of nuclear war tends to divide viewpoints into factions based on the emotional response that concept evokes. That emotion translates to inducing bias as experienced from decades of elicitation efforts on sensitive and taboo topics, including nuclear weapons and war. The deterrence community also appears to be divided into factions regarding the effectiveness, or lack thereof, of nuclear weapons. The well-documented debate of two such factions can be found in the works of Sagan and Waltz.

Analysts should be aware of such perspectives and should question experts about their preexisting (i.e., anchored) positions. Along with that, other questions about the experts' specific areas of research and experience provide information about how their responses may be biased. Statistical analysis may be able to determine whether or not these biases affect responses, by comparing responses among experts whose preexisting positions are established and whose problem-solving processes have

and two are specific to the time during the Cuban missile crisis. It should be noted that lack of specificity is a type of uncertainty because the analyst looking at this table faces the conundrum of how to compare results from unspecified conditions to the results from specified ones.

Table 3.3. 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)	George Bundy	1988
	Probability of a future Cuban missile-type crisis that results			

problem, and gaining insights into the author's cognition by reading the author's papers. Without such conditioning information, the analyst can only compare "apples to apples." The four estimates that terrorists will detonate a nuclear bomb in the next decade have a large, unexplained, range of 1 percent to 50 percent. The other estimates cannot be included with these unless and until the conditional factors inherent in them are known, putting them in the same terms as the first four. The Bundy and Kennedy estimates can be compared to each other but not the rest.

Informing Decision-Makers

Quantifying or summarizing results from elicitation and analysis should be done in a form useful for and understandable to decision- and policy-makers. Determining that format may involve an elicitation with the decision-maker. While top-level managers rely on executive summaries, details should be made accessible for their staff and for future updates as knowledge changes.

Returning to the data in Figure 3.1, quoting the mean response of 31 percent to a decision-maker without the uncertainty does not convey an adequate summary of these data. In this particular case, the histogram in Figure 3.1 does provide an appropriate summary. However, a decision-

Jane M. Booker

of assessing the risk of nuclear deterrence failure may suffer from the same difficulties.

However, even for difficult, amorphous, or ill-posed problems, experts tend to think in terms of some sort of problem structure or framework based on the logic behind their understanding. That structure may be loosely defined, choppy, disjoint, approximate, general, vague, and difficult to record on paper, a whiteboard, or a computer pad/tablet. Detailed probing into the expert's thinking may be required to elicit a rough draft that mimics the expert's thoughts about their portion of the problem. During the elicitation reasu4-6.30.8..30.8 (t)-27hlt6-23.2 (i5-25.6 (l)-26.-4 (v

structure. Neural networks, factor analysis, cluster analysis, statistical covariance, and correlation structures are some commonly used techniques to uncover data structures. Although many of these require large amounts of numerical data, some can still be used for smaller amounts of more general knowledge.

For example, an expert examining the results of a neural network or factor analysis of historical events data might be able explain the data structure found from this analysis by seeing an association or reason that was previously not considered. That reason or association would then be an added feature to the problem structure.

Analysis for structure in the data, information, or knowledge (e.g., historical record) is recommended, when possible, because understanding the data/knowledge structure often provides insights into the problem structure. Even organizing all the available data, knowledge, and information into files, spreadsheets, or perhaps databases reveals problem structure. For the nuclear deterrence problem, it is unlikely that much analysis would be possible because of the sparse amount of data, information, and knowledge available. However, some collection and organization of the applicable data, information, and knowledge will be necessary for simple bookkeeping. This effort can reveal structure in the knowledge, which might, in turn, be useful for considerations about the problem structure. If the structure in the knowledge is inconsistent with the problem structure, the reasons for this conflict should be understood.

Eliciting a Structure

The formal elicitation principles from the first section of this chapter have been applied to eliciting a structure from experts. Eliciting a problem

Such interactions often reveal new understandings that cut across different aspects of the problem.

The first step in eliciting problem structure is to ask the expert(s) to simply write down some of the fundamental components, issues, or aspects of the problem. For nuclear deterrence experts, this would include eliciting their areas of expertise and experience. Defining the problem scope—what may or may not be included—also starts here. Usually this first set of items supplied is at a very general level of detail, representing the basic problem features, facets, subject areas, and historical record. For the deterrence problem, these items could include a time frame (past and future), participants involved (states, groups, leaders), socio(en-8)-5.6 (o

socioeconomic and political factors necessary for any state or terrorist group to consider when committing to the acquisition of a nuclear weapon.

and issues of the problem, including the first round of relationships and associations. Any difficulties in formulating or recording these should be noted and completion should be postponed. Likewise, focus on organization or logic flow is not necessary yet and may still be too ill posed. Organization and flow may become clearer as the elicitation progresses.

To distinguish details from general items, an iterative course in the elicitation is helpful. Start with the most general level of detail and then elicit more specific issues, facets, ideas, etc. However, getting specific can quickly burden and complicate the expert's thinking, resulting in inconsistency and in reaching knowledge voids or gaps. An alternative strategy is to stop drilling down in detail and generalize once more. Guide the expert, without fatiguing them, to iterate between thinking about the general to the specific and back again as often as required. The reason for this is to aid the expert in keeping the bigger picture in mind while decomposing the problem into details. For example, the bigger picture might be a particular assumed political environment, affecting the detailed issues, events, and outcomes within it.

Permit the expert to leave holes, blanks, and question marks as placeholders for things not easily characterized or known. These voids can be addressed in a later iteration or after the expert has had a chance to ponder, calculate, or research. Other experts may have to be used to fill in these gaps. Alternatively, these holes, blanks, or questions may never get completed because the knowledge simply does not exist. This lack of knowledge is part of the uncertainty inherent in the problem. The same is true of describing associations. Some may remain vague or ill defined. A simple notation succeeds such as "I know A is somehow related or important to B, but I just don't know what that relationship is."

The experts should not try to complete the structure in one elicitation session or even one day. Time between sessions gives the experts a chance to rethink and reorganize, preventing cognitive overload. It is not uncommon for the expert to return to the next elicitation session and completely start over. However, the previous work should not be discarded.

It may be possible to establish some major general features in one session and then develop the structures for each of these in subsequent sessions. The level of detail may not be the same for all features of the problem. Some aspects of the problem may be known in great detail. Others may be listed

at only the most general level, with nothing known in detail. For example, the actions of some newly formed terrorist faction would difficult to detail.

An expert may designate some issues, relationships, or portions of the problem for other experts to structure. Bringing in new experts brings in new knowledge, but it can also bring in disagreements about how to structure the problem. Resolution of disagreements between experts takes time; however, it usually provides valuable insights for the interviewer, analyst, and the experts. Some disagreements may not be resolved. These unresolvable differences reflect the large uncertainty in the state of knowledge for that issue.

Some Difficulties in Eliciting a Structure

A few difficulties involved for ill-posed problems such as the risk of nuclear deterrence failure are described below.

Experts may run into dead ends where their thoughts cannot be depicted because of complexities or lack of knowledge or because they have not thought about how to structure aspects of the problem before. Dead ends are legitimate. There is a difference between forcing experts to supply knowledge that does not exist and asking them to use their expertise beyond their personal experience or comfort zone. The former results in biased, contentious responses, whereas the latter minimizes anchoring bias. For example, asking experts to consider circumstances according to their knowledge for when a state leader might detonate a nuclear weapon on US soil may be uncomfortable but can be within the expert's capability. Demanding that the experts read the leader's mind is unreasonable.

The unknown or little known details (high uncertainty issues) can hinder thinking and even contribute to cognitive overload. The same is true for poorly understood relationships, such as degrees of association or dependency. For example, an expert may state something like "I just don't know why country A nearly always votes like country B in the United Nations, but it just does."

The expert may have to explore various ways of depicting the problem, which can be frustrating and time consuming. The expert may find it difficult to think aloud or record on paper their thoughts about the structure. These difficulties are not necessarily due to some inability of the expert, but they stem from the complexity, knowledge-poor nature, and high uncertainty inherent in the problem.

respectively, for likelihood and consequences, as shown in Figure 3.2. However, this representation actually depicts fuzzy sets for the risk constituents. For instance, the risk denoted by the X has degrees of both yellow and green but is mostly green. us, X partially belongs to the yellow (medium) set and more to the green (low) set. e risk at X cannot be precisely assigned to either the low or the medium sets. e same is true of the risk denoted by the asterisk, which has most membership in the red (high) set but some in the yellow (medium) set.




Figure 3.2. Fuzzy shades for the constituents of risk.

Uncertainty Perspective

would be asked to view the problem and its aspects in terms of uncertainties instead of the event/issue perspective. The challenge would be that experts should be comfortable with thinking about uncertainties, and most experts are not.

Regardless of the problem structure used for assessing the risk of deterrence failure, managing the different types of uncertainties will be a challenge.

Information-Gap Structure and Triad Principles

The JASON study correctly concluded that rare events could not be predicted because of the lack of data and the lack of specific information about such events.⁵¹ This conclusion actually is based on a type of uncertainty called nonspecificity.

Nonspecificity is the uncertainty from relying on the general to determine the specific. For example, with so few and so varied kinds of attacks (events), the best an expert could predict for the future would be a rate or average time until the next attack but with no specification about what, where, or how it would happen. As the JASON study concludes,

is sought from similar and relevant problems. Using other knowledge sources induces additional uncertainty based on how close these other sources are to the problem of interest. For example, an important issue in the deterrence problem is assessing close calls. Another knowledge source that could be useful for understanding nuclear close calls would be to understand close calls in historical military attacks (see chapter 2). Knowledge source structuring can be considered a knowledge-integration-structuring approach and is described in chapter 8.

Assessing Risk with Expert Knowledge

Regardless of the structuring approach for the problem or for the knowledge, the high uncertainty and knowledge-poor nature of the risk of deterrence

must also be combined through the structure. Again, the source of these uncertainties may be solely from the experts' experience and knowledge.

Evaluating Consequences

The second constituent of risk is determining the consequences, as shown in Figure 3.2. A common form, quantity, or standard of these is less obvious because consequences stem from different subject areas: loss of life, damage to property, cost, time, and perception. A utility or utility function is often formulated to transform these different consequences to a common scale or measure of value or worth. Sometimes a dollar value is used as a common measure of utility.

Consequences of deterrence failure are particularly devastating—nuclear weapons exchange or nuclear war. While these are difficult to evaluate and estimate, comparative techniques, such as Saaty's Analytic Hierarchy Process, and formal elicitation techniques aid the expert in thinking about the unthinkable.



Assessing the risk of nuclear deterrence failure is a complex problem covering multiple subject areas. Common to these subject areas are sparse or lacking data, lacking theory or models, high uncertainty, and involvement of human behaviors and decisions. Because of these difficulties, analysts must rely on the use of experts and formally elicited expert knowledge. Established problem structuring and framework methods (e.g., logic or block diagrams) may not be appropriate and may be inconsistent with the way experts think about the problem or their portions of it.

An alternative approach for structuring, framing, and/or organizing the ill-posed deterrence problem is to elicit the structure from the experts. The same formal elicitation techniques briefly described in the first part of this chapter also apply to eliciting problem structure described in the second section. These bias minimization techniques help ensure that the knowledge gathered is of the best quality.

QMC /P <</Lang (en-US)/MCID 1538 >>BDu.8 (r)24.6 (n)-14.9 (i)-17.8 (zTw 10.

For the challenging problem of assessing the risk of deterrence failure, an analyst should rely on an expert-oriented structuring of the problem and should use all available sources of data, knowledge, and information. The integration approach necessary to analyze such a structured problem and to draw conclusions is discussed in chapter 8.

In summary, assessing the risk of nuclear deterrence failure relies on the existing state of knowledge of the experts in its subject areas. Eliciting that knowledge with established formalism for minimizing biases is feasible, as outlined in this chapter. What is described is an expert-oriented, expert-

27.

43. An example of a hand-drawn decision tree can be found at https://commons.wikimedia.org/wiki/File:Manual_decision_tree.jpg
44. Ronald A. Howard and James E. Matheson, eds., "Influence Diagrams on the Principles and Applications of Decision Analysis" (Menlo Park, CA: Strategic Decisions Group, 1984).
45. John von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior* (Princeton, NJ: Princeton University Press, 1944).
46. Bunn, "A Mathematical Model of the Risk of Nuclear Terrorism."
47. Jane M. Booker and Laura A. McNamara, "Expert Knowledge in Reliability Characterization: A Rigorous Approach to Eliciting, Documenting, and Analyzing Expert Knowledge," chap. 13 *Engineering Design Reliability Handbook*
48. Zadeh, "Fuzzy Sets."
49. James J. Buckley, Thomas Feuring, and Yoichi Hayashi, "Fuzzy Hierarchical Analysis Revisited" *European Journal of Operations Research* 120, no. 1 (2001): 48–64, [https://doi.org/10.1016/S0377-2217\(99\)00405-1](https://doi.org/10.1016/S0377-2217(99)00405-1)
50. John L. Darby,