Chapter 8
Knowledge Integration

Jane M. Booker

For multifaceted problems such as assessing the risk of nuclear deterrence failure, data, information, and knowledge can emerge from many different sources involving diverse subject areas and in myriad qualitative or quantitative forms. Often the amounts of data, information, and knowledge are limited, apply to rare events or events that have never occurred, or both, necessitating the combined use of all sources. For example, sources include historical data on past events; expertise from authorities in different subject areas; and knowledge about past and current cultures, human behaviors, sociology, politics of people and states, as well as the theory or rules governing politics. Regardless of source and form, available knowledge has uncertainty attached. Some uncertainties can be significant, and the uncertainties themselves can be of different types. Depending on the type of uncertainty, quantification may not be feasible or the appropriate mathematical theory for it may be difficult to apply. Nonetheless, decision- and policy-makers need a final or top-level answer about nuclear deterrence failure accompanied by an understandable uncertainty. Knowledge integration methods address these needs and provide ways to tackle other difficulties encountered when combining all available data, information, and knowledge and their associated uncertainties to produce an assessment of risk. Some of the integration principles and methods are described in this chapter, especially those related to the challenges in assessing the risk of nuclear deterrence failure—a problem of significant uncertainties and poor data, information, and knowledge.

The first step in laying the foundation of the concepts for the knowledge integration required to assess the risk of failure of deterrence is to distinguish among data, information, and knowledge.

Data are observations of knowledge that are measured, recorded, enumerated, described, or numerically or symbolically represented. Data are generally considered analyzable, which implies a numerical, ordinal, or categorical form.[1]

Information is commonly defined as facts provided or learned and that which is subsequently conveyed or represented.[2]

Knowledge refers to a body of facts gathered by study, experience, or observation or inferred from those. Knowledge implies processing through the human brain. Thus, cognition, experience, memory, and mental processing are involved in the formulation of knowledge.

Simple examples may serve to clarify these concepts. For the problem of determining the risk of nuclear deterrence failure, data might be in the form of frequency counts of categories of historical events. For example, how many times was the nuclear alert level raised during an international crisis? An example of information would be an intelligence report on a missile test by a foreign state. A knowledge example would be the physics theory necessary to develop a nuclear weapon.

Inherent in the study of data, information, and knowledge is the uncertainty associated with it, for which probability is the most commonly used uncertainty theory with the longest history of use.[3] An uncertainty in the data example above is the possibility of an unknown raising of the nuclear alert level. For the information example, an uncertainty exists in the accuracy of the intelligence report. Even theory, such as nuclear theory, may not be exactly known, creating knowledge uncertainty.

Because the lines among data, information, and knowledge are often blurred, the term *knowledge* is used in this chapter title to represent all three. Additionally, knowledge is the most general, and it is what would be assessed in determining the risk of deterrence failure. Any necessary distinctions among the three, such as differences in applicable methods, are noted. The smallest unit or singleton of knowledge is referred to as a piece of knowledge, which could be a number, a word, or a phrase or statement sufficient to contain the fundamental knowledge.

Knowledge can be considered as having two forms: quantitative and qualitative. The distinction between these two categorizations is not precise (i.e., it is fuzzy).[4] Qualitative knowledge often involves the use of linguistics, some of which can be quantified. Examples include ordinal linguistics such as small, medium, or large and relative comparisons such as worse, the same, or better. The principles and methods for the mathematical combination and/or summarization of all forms of knowledge under uncertainty are referred to as knowledge integration.

With this terminology in hand, the first section of this chapter provides a history of better-known methodologies leading to knowledge integration. The second and third sections address the various aspects of the deterrence assessment problem that require combination and integration principles and methods, respectively. Because of its centrality to the problem, uncertainty combination is treated separately in the fourth section. Following that, the fifth section describes the challenges and benefits associated with knowledge integration. Finally, the sixth section summarizes key issues for the assessment of the risk of deterrence failure.

A Brief History of Knowledge Integration

One of the earliest data combination techniques—multiple frame sampling—came from statistical sampling theory.[5] A sampling frame can be thought of as a partial list of the entire population of interest. Data gathered from different sampling frames are combined, and uncertainties are combined by using probability theory. For example, a telephone survey uses the frame of the phone directory. The telephone-based sample could be combined with a mailed survey using the frame of addresses from the Census Bureau or local government records. The goals of multiple frame sampling are to ensure proper coverage of the population by using data gathered from different frames and to decrease the uncertainty (as variability) in the inferences made about the population. The latter is accomplished by an increased sample size from the combined frame samples.

This sample combination idea was later extended to combining entire studies. Because gathering data from studies (especially human studies) is expensive, meta-analysis was developed to combine different studies. Meta-analysis also has its foundation in statistics, again characterizing uncertainty with probability theory. Combining quantitative data from various sources (i.e., studies) has several advantages:[6]

1. Results from one study can be confirmed by others.
2. The sample size is increased, which reduces variability.
3. Additional effects of varying conditions can be determined.
4. The strengths of relationships between associated quantities can be determined.

As these reasons suggest, more than just data are considered in the combination. Information, in the form of relationships among quantities, is also combined from separate studies to achieve the fourth listed benefit. Application areas for meta-analysis include biology, medicine, social sciences, and education. Credit is given to Glass for pioneering work in this area.[7]

Data fusion[8] and data integration generally refer to the process of combining data from different locations, such as a sensor network or multiple geographical sites. The computational community is part of this field because of the need for managing data with databases and data structures. Data integration is sometimes associated with combining data from different studies, overlapping with meta-analysis. However, methods for data integration are not limited to statistical ones, as with meta-analysis. These integration methods also involve mathematical logic and computational algorithms.

Extending meta-analysis and data integration methods to a more general knowledge or information integration methodology began in the late 1990s with the PREDICT reliability methodology.[9] Before that, an enhancement of probabilistic risk assessment toward an information integration methodology was done in the NUREG-1150 study by merging expertise with data.[10] Four additional major extensions were developed for PREDICT, the first information integration methodology:

1. Formally elicited knowledge was integrated with sparse data and poorly validated theoretical calculations.
2. Knowledge and its uncertainties were quantified and combined using different theories (probability theory and fuzzy sets).
3. Expert knowledge was used to provide structure for a problem not suited to contemporary structuring methods.
4. Validation of the integration methodology was achieved by a simultaneous application to a problem where vital data eventually became available.

Knowledge integration extends the combination beyond data, studies, and relationships to all knowledge, information, and data and their associated uncertainties. For example, elicited knowledge from experts is integrated with what sparse data may be available. Knowledge integration enhances and generalizes many of the techniques from meta-analysis and data fusion. In particular, different types of uncertainties are analyzed using general information theories rather than just probabilistic uncertainty.[11] Knowledge integration methods are useful for combining the different quantities described in the following section.

Multiple Integrations

There is a need for integration methods that would be used to combine (1) all available knowledge from different sources and (2) different types of uncertainties associated with these. However, this is not the complete list of what requires integrating in order to assess the risk of deterrence failure.

Integrations for Risk and Deterrence

Assessing the risk of failure of nuclear deterrence involves assessing the two constituents of risk (likelihood and consequence) and the two constituents of deterrence (credibility and capability). Individually, each of these four constituents presents a difficult combination problem, involving multiple integrations of diverse knowledge sources, only some of which are quantifiable. In addition, integrating tools are lacking and uncertainties of various types are large, primarily because of lack of knowledge. None of these four constituents for assessing the risk of deterrence failure has sufficient knowledge available to use statistical risk analysis methods.

Traditionally, likelihood has been represented as probability. Probabilistic risk assessment has a long history of this practice. Consequence evaluation involves difficult-to-assess quantities such as the value of human life and property and the chaos or damage from destruction. Often multiplication is the mechanism for combining likelihood and consequence. The difficulty with multiplication stems from the fact that the same value of risk can result from a low-consequence–high-probability event as from a high-consequence–low-probability event.

Credibility and capability should be evaluated from the perspective of the party being deterred. Because perspective is involved, evaluation and determination of associated uncertainties of these two constituents are challenging.

Neither combining likelihood and consequence to assess risk nor combining credibility and capability to assess deterrence may be feasible because of these difficulties. To approach achieving a risk analysis capability for the problem of nuclear deterrence, what little knowledge that exists must be collected and put together, necessitating the use of knowledge integration methods for each constituent. Any connections or relationships discovered among the four during their individual assessments should be noted. However, the risk assessment may address the four constituents without an integration of them.

Integration of Experts' Knowledge

For risk assessment of the failure of nuclear deterrence, a major source of knowledge is going to be provided by experts from multiple and diverse subject areas. Drawing conclusions about or providing top-level answers for the four constituents is likely to require resolving differences among these various experts. While differences can originate from different subject areas, differences among experts in the same field are also to be expected. Any group of experts not exhibiting any disagreement is a warning sign that something is amiss with the elicitation or with the selection of experts.

Expert resolution elicitation techniques permit understanding of why disagreements occur and provide ways to resolve many forms of and reasons for differences.[12] As noted in chapter 3, reasons for differences include experts answering slightly different questions than the one posed, experts making different assumptions, and experts using different problem-solving processing in arriving at their responses. Any unresolved differences represent the inherent uncertainty in the current state of knowledge.

Integration of Scenarios, Conditions, and Problem Dimensions

An expert may be asked to provide knowledge about different versions of the problem. For example, different background conditions and/or scenarios leading up to a nuclear attack or war can be posed. The expert would then provide different (conditional) responses depending on the specified conditions and/or scenarios. An unconditional response is formed from conditional responses by using an aggregation method from the "Conditional Combination" subsection of the "Knowledge Integration Principles and Methods" section.

In addition to aggregating conditional responses for each expert, it may be necessary to combine responses across multiple experts from different subject areas. This is particularly true for the problem of determining the risk of deterrence failure, where no single individual has expertise in all subject areas and all problem dimensions. An example of ignoring different subject area expertise is the Lugar survey, in which respondents from different fields answered survey questions covering multiple subjects.[13]

Integrating Uncertainties

Portions of a complex problem that have less available knowledge tend to have larger uncertainties. In addition, different types of uncertainty may have been identified for specific parts of the problem. Those different uncertainties can have different mathematical theories to characterize them. Thus, knowledge integration involves the combining of uncertainties and uncertainty theories. Methods and issues for combining uncertainties are described in the "Uncertainty Combination" section.

Recomposition

In formal elicitation, it is vital to decompose a complex problem or system into manageable pieces for an expert. Such a decomposition must eventually be recombined to obtain an overall top-level answer, which a decision-maker or policy-maker expects. Recombination of different decomposed portions of the problem can involve integrating over different levels of detail, combining the specific with the general.

Difficulties with recomposition of the problem occur when different amounts and types of data, information, and knowledge exist for different parts. For example, historical data may exist for certain (more common) categories of events but not for unique (one of a kind) or never-observed events. The tendency (bias) is to focus on pieces of the problem that have more data, information, and knowledge because analysis and integration methods are easier to implement and uncertainties are easier to quantify. This biased activity is called pearl polishing.

An example of pearl polishing in nuclear deterrence during the Cold War is the US focus on exchange analysis and first-strike stability because those issues were better understood at the time and more analytically tractable. More attention should have been paid to the broader, more difficult issues regarding culture, psychology, and politics. Attention, analysis, and decisions based only on better-known portions of the deterrence problem will not make up for the difficulties resulting from, and lack of attention to, pieces of the problem that are less known or understood.

Recomposition involves utilizing the framework or structure of the problem. Structures can be networks, trees, diagrams, and specialized structures (such as might be supplied by experts consistent with their thinking). Methods for determining that structure are discussed in chapter 3 and include methods from risk analysis, decision analysis, graph theory, logic, and complex systems. Whatever the form of the structure and the nature of the connections between problem parts, it can provide the mechanism and/or rules for integration of parts to reach the top-level answer, such as the four constituents of risk and deterrence.

Knowledge Integration Principles and Methods

Although some tools and methods exist for knowledge integration, research and development of these continues, and many issues remain unresolved. Without presenting details of some of the complicated and esoteric methods, most of which apply only to data-rich problems, an overview of fundamental principles and methods follows.

Quality of Knowledge

Because data, information, and knowledge in the deterrence problem are sparse and/or have large uncertainty, quality of knowledge becomes an important issue. Quality of knowledge involves the use of established practices for gathering knowledge in a manner that properly represents the current state of knowledge.

Knowledge from validated theory, models, or computation is of good quality. Unfortunately, validation (matching theory, models, and computation with reality) requires sufficient data, information, and knowledge. Validation is highly improbable for the risk of deterrence failure problem.

Much of the knowledge for the deterrence problem will come from experts. The difference between using and not using formal elicitation methods is the quality of knowledge. Experimentalists are taught to practice good protocols to ensure the quality of their data, such as implementing controlled conditions and documenting every detail. Formal elicitation serves the same purpose, using techniques to monitor and minimize biases.[14]

In data fusion, care must be taken to understand the quality of the data before combination. Data from biased studies or data that are unrepresentative of the population are of poor quality and are not equivalent to data gathered under controlled conditions. Likewise, knowledge must be gathered that properly represents the population (the current state of knowledge). With experts being the major source of knowledge, expert selection becomes important, as discussed in chapter 3.

The risk of deterrence failure is one of the problems for which a portion of the current state of knowledge resides in the classified community. Studies conducted in an unclassified environment deal with the unclassified population, and results and conclusions are conditional on that population. If classified knowledge becomes accessible, the population is broadened. Whatever the population, the quality of knowledge depends on adequate representation of it, and the conclusions are conditioned on it.

Source Inventory and Evaluation

Integrating knowledge from different sources requires an inventory and some evaluation of the sources and types of knowledge that are available or accessible to anticipate integration difficulties and needs. Table 8.1 shows a simplified listing of knowledge sources (rows) for four discrete cases along a continuum of problems (columns). A common (and somewhat related) example is given for each case: case I concerns how to determine the energy yield of TNT explosive, case II is for determining the electrical power generation capability and safety of nuclear energy reactors, case III concerns the determination of nuclear weapon yield under a testing ban, and case IV is the risk of failure of deterrence problem.

The types and sources of knowledge are listed in the far-left column of Table 8.1. The availability and evaluation of these sources is shown in interior cells of the table for each of the four cases. In the Data row, the cases range from data rich (case I) to data poor (case IV). The History row describes the length of time of experience, including the number of realized events. The state of knowledge about theory or first principles that apply is represented in the third row. How much inference is required to interpret the knowledge, draw conclusions, and make decisions or policy is depicted in the next row. TNT explosive energy output is so well known that it is a National Institute of Standards and Technology standard, requiring no inference. The need for experts as a major source of knowledge is evaluated in the next row. How well calculations or models can be formulated and validated to make predictions is addressed in the Model row. Finally the degree of uncertainty present is listed in the last row. One might argue that uncertainty is not a source of knowledge; however, it is important in the evaluation process and in the knowledge integration.

According to the evaluations in Table 8.1, the risk of deterrence failure problem is the worst and most difficult on all accounts—not a surprising result. Available knowledge is biased because some sources (e.g., states) are not as open as others, and that bias is a form of uncertainty. Other sources are also poor, resulting in heavy reliance on experts and contributing to large uncertainty.

While experts do formulate their knowledge from these other poor sources, they also incorporate their own cognitive processing ability as a primary source. Understanding experts' thinking becomes even more important when combining expert knowledge with whatever other meager knowledge is available from other sources.

Table 8.1. Knowledge Source Evaluation

Knowledge Source	Case I: TNT Explosive	Case II: Nuclear Power	Case III: Nuclear Weapons	Case IV: Deterrence Failure
Data	Large	Moderate	Small	Sparse
History	Long	Short	Short	Short
Theory	Solid	Good	Moderate	Poor
Inference	None	Little	Much	Very much
Experts	Not needed	Some use	Greatly needed	Heavy reliance
Model	Great	Good	Some	Poor
Uncertainty	Small	Moderate	Moderate	Large

The paucity of knowledge in the deterrence failure problem necessitates a waste-nothing assessment approach that utilizes knowledge integration methods. The Table 8.1 evaluation also indicates the need for:

- formal expert knowledge elicitation methods, including techniques for understanding experts' thinking and problem-solving processes;
- integration methods that combine elicited knowledge with any data, information, and knowledge from other sources, such as history; and
- integration methods that focus on making inferences and dealing with large uncertainties.

Common Quantity

A fundamental principle of integration is to combine data, information, and knowledge having common units, common definition, common representation, or common structure. This principle is designed to avoid combining "apples with oranges," meaning unlike or disparate things. The common quantity is usually the quantity of interest in the study, such as risk or its constituents.

Often it is possible to transform or convert dissimilar quantities so that they have a common scale or definition. For example, one can use conversion factors to change foreign currency to US dollars or to establish common units (e.g., measuring every quantity in feet rather than a mixture of units). These conversions are well established and straightforward but will occur infrequently when assessing the risk of deterrence failure. Other less obvious, more frequently occurring, and more important conversions may require querying subject-matter experts. Term definitions must also be verified for consistency of use, especially when dealing with knowledge elicited from experts in different fields. For example, one expert may view consequence of a nuclear attack in terms of loss of lives and property, while another may also include changes in stability among states.

In some methodologies such as probabilistic risk assessment, reliability, or decision analysis, common quantities are well established, and all knowledge is transformed to those. In probabilistic risk assessment, probability is the common quantity, in reliability it is reliability, and in decision analysis it is utility. Experts may have to agree on formulae or functions (e.g., a utility function) to supply the mechanism for transformation between quantities.

In general risk assessment, likelihood is one common quantity that may or may not be defined as a probability. However, consequence is also a component of risk, and there have been various attempts at determining common definitions for consequence, such as providing equivalency of the dollar value for the loss of human life. It is difficult to assess the dollar value of things like physical damage; political stability; and emotional, cultural, and lifestyle changes of peoples as a result of deterrence failure.

Sometimes establishing commonality can be accomplished by changing the level of detail of the knowledge. In the apples with oranges example, although an apple is not the same thing as an orange, they are both fruit. If the common quantity level is broadened to be the more general fruit rather than the specific apple, one can combine apples with oranges. An example would be to categorize different types of weapons of mass destruction threats according to weapon type (e.g., biological, chemical, nuclear, etc.) rather than using specifics, such as a nuclear device manufactured by a terrorist group.

The sacrifice made by changing to a more general level is that detailed information is lost. Loss of detail induces a nonspecificity uncertainty when or if such detail is ever needed in the future. For example, it may become important later to know whether the more general fruit was originally an apple or an orange. If that original detail is lost, it is uncertain which fruit it was. Documentation of the original knowledge avoids this kind of nonspecificity uncertainty when transforming to a more general common quantity.

Obviously, one could carry the idea of generalization to a ridiculous extreme, losing all content and meaning of the original knowledge. Finding the appropriate level of generality to establish common quantity may require a group elicitation, including resolving differences among experts; this is especially true for experts in multiple and diverse subject areas.

Weighting Schemes

Any combination, aggregation, or integration procedure can be considered as the implementation of some sort of weighting scheme. Using this general definition makes weighting schemes the backbone of knowledge integration.

Knowledge integration of an established common quantity can essentially be accomplished by the choice of the appropriate weighting scheme.

The "Multiple Integrations" section described the various kinds of combinations necessary for the risk of deterrence failure problem; however, the means of formulating those aggregations was not addressed. Weighting schemes are the primary mechanism for those combinations. For example, most of the knowledge will come from experts. If more than one expert provides expertise for an issue or question, then a weighting scheme is required to combine their knowledge. Likewise, knowledge from other sources (e.g., historical data) would be combined with that from experts. Finally, knowledge about issues or portions of a problem is combined using a weighting scheme to form the top four constituents of risk and deterrence.

Before application of a weighting scheme (or any combination method), differences, inconsistencies, and disparities among the pieces of knowledge to be combined must be resolved. Chapter 3 provides guidelines on how to resolve these differences among experts, and most of those methods are applicable to resolving other differences. For example, if two pieces of knowledge are contradictory and nothing can be found to explain this, then the resolution becomes a matter of determining the combined uncertainty from the two pieces. Specifically, the uncertainty in both pieces of knowledge is so large that both realizations are possible. To illustrate, suppose a state leader claims that they will attack an adversary on one day, but their next speech talks about peaceful coexistence. This leader keeps alternating between these outcomes in other speeches and documents, for no apparent reason. The result is that the uncertainty regarding the leader's course of action is so large that their adversary must simultaneously prepare for both actions. Should the adversary prepare equally for both outcomes or favor one as more likely over the other? That answer is a matter for establishing the weights.

The weights for an integration scheme may be numerical, including ranks, or ordinal, including linguistic qualifiers or rules. Likewise, the knowledge being integrated, its uncertainty, and the uncertainty of the weights can be quantitative or qualitative. The quantitative schemes are introduced here, and the qualitative schemes are described in the "Logic and Rule-Based Combinations" subsection.

How to determine the quantitative weights is the first challenge. The easiest and most common choice is to consider the pieces of knowledge to be combined as having equal validity and applicability. This is the equal weights combination. Equal weights are recommended for combining knowledge from different experts, unless there are definitive reasons for weighting some experts more than others.[15] That same recommendation can be applied to any integration process for the same reason: differential weights require good reason and justification.

Assigning a weight of zero means that the piece of knowledge (or even the expert) is deemed incorrect, irrelevant, or inapplicable. Reasons for eliminating knowledge from combination should be documented in case that piece becomes relevant later. Elimination should be a rare occurrence.

A weight of zero is often calculated for weighting schemes based on an event's frequency of occurrence when the event has never happened. This is true for percentages, proportions, weight of evidence[16] and other ratio-based weights. For never-observed events, such as the number of times terrorist groups have used nuclear weapons, these weight calculations become meaningless. However, these weight formulations can be used when multiple data, information, and knowledge sources are combined if any source has a nonzero numerator. The section on "Bayesian Integration" illustrates.

The human brain assimilates knowledge in cognitive processing by using its own weighting scheme. Each of us determines the relevance and importance of the knowledge we acquire and how to combine new knowledge with existing pieces from our experience. This is why eliciting expert thinking is useful for determining weights and weighting schemes.

Weights, including equal weights, have uncertainty. The simplest way of capturing that uncertainty is to select a range of values or ordinal descriptions for each weight. For example, an expert comparing events may explain that event A is two or three times more important than event B. If the weight for event B is 1, then the weight for event A is somewhere between 2 and 3. The uncertainty for the weight of event B must next be determined. In doing so, the expert may also have to expand the interval for event A to maintain the factor of 2 to 3 between A and B.

The most fundamental weighting scheme is the average or mean. In calculating the mean of two or more pieces of knowledge, the combination is the sum of equally weighted pieces. The weights are defined as the fraction, 1 divided by the number of pieces.

Bayesian Integration

Because weights and their uncertainties are a challenge to determine, an automatic weighting process is desirable, which is one reason for the application of Bayes' theorem. No supplied weights are required to implement Bayesian integration because the mathematics within the theorem generates them automatically from the information contained in the supplied knowledge sources. Bayesian integration is important because many analysts consider it the premier data combination methodology; however, it has disadvantages and limitations of applicability.[17]

Bayes' theorem[18] is a convenient mathematical combination or weighting scheme for combining two sources of knowledge quantitatively expressed in functional form, called the prior distribution and the likelihood function.[19] The resulting combination of these two functions is another function called the posterior distribution.

Integrating expert knowledge with experimental or observational data using this theorem has been done for many decades and remains popular today. The expert-supplied knowledge is formulated as the prior distribution, and what little data may exist are formulated as likelihood functions.[20] This expert-with-data integration is useful for problems with phenomena that have not yet occurred, such as a failure. Thus, it is applicable to the risk of deterrence failure problem, where prior distributions could be formulated from experts to combine with the sparse historical record.

However, before the 1990s, Bayesian analysts did not concern themselves with formal elicitation methods until those methods were developed and applied. One of the first applications was NUREG-1150, probabilistic risk assessment studies of several nuclear reactors.[21] The Nuclear Regulatory Commission sponsored this massive study to replace the previous one, WASH-1400, in which formal elicitation methods were not used.[22]

For the risk of deterrence failure problem, Bayesian integration could be applied for combining different knowledge sources. For example, results from a previous study could serve as the prior for the results of a new study. This example also illustrates another advantage of using Bayes' theorem: updating or integration can be done on a continual basis. Whenever new knowledge becomes available, the previously combined knowledge (the posterior distribution) then becomes the prior distribution to be updated with the new knowledge. This updating feature would be useful for the deterrence problem because of the constant changes that occur in the available knowledge.

One disadvantage of the combination capability of Bayes' theorem is that the available knowledge must be captured and quantified in the form of functions—distribution functions for the prior and posterior, and likelihood functions and quantities (called parameters) associated with those. Critics of Bayesian methods cite this reason to argue against its use. Reverend Bayes' original form of the theorem contained probabilities instead of functions. However, this formulation is no easier to use for expert knowledge because humans are not good probabilistic thinkers.

Because of the way the mathematics of the theorem operates to weigh the two sources of knowledge, there are cases where the resulting combination (the posterior distribution) does not make sense. This is another disadvantage, limiting the utility of the theorem. One case arises when a large amount of knowledge—a body of evidence—is formulated into a prior that is inconsistent with a small amount or piece of knowledge formulated into the likelihood. For example, at the time of the Cuban missile crisis, military experts had considerable knowledge to support the idea for a land invasion of Cuba. However, their prior information would have been inconsistent with a new piece of knowledge—that the Soviet Union had placed tactical nuclear weapons in Cuba to repel an invasion—had such information been available. But Bayesian combination would have still supported invasion because the large amount of prior information would have outweighed the single new piece, as depicted in the top of Figure 8.1. Basing a decision to attack on the Bayesian combination would have been a bad idea. Instead, the single new piece of evidence in this example should outweigh all prior knowledge and drive the decision not to invade.

Another undesirable result from Bayesian integration arises when prior knowledge conflicts with near equal amounts of likelihood knowledge, as shown in the bottom of Figure 8.1. For this situation, Bayes' theorem produces a combination that lies between the two, in a region where knowledge from neither source is found. Returning to the state leader example in the "Weighting Schemes" subsection, the leader's first statement (prior) supported attack, while the second (likelihood) supported peace. Bayes' theorem produces a combination of half attack and half peace, an indeterminate result.

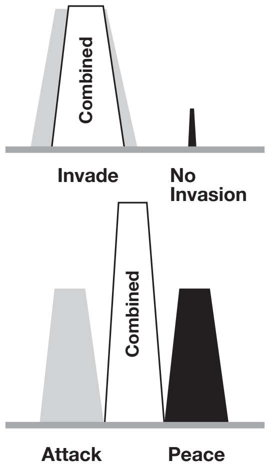

Figure 8.1. Bayesian integration issues. The top image depicts Bayesian integration ("combined" white trapezoid) of a large amount of prior knowledge (gray trapezoid) with new knowledge of nuclear weapons (black line). The bottom image depicts equal amounts of prior (gray) and likelihood (black) knowledge, resulting in a combination ("combined" white trapezoid) falling between, where no knowledge resides. Shapes are for illustration purposes and are not drawn to exact dimensions.

Redundancy and Dependency

Dependency between and redundancy among knowledge sources leads to double counting of the same knowledge, unless the overlap is identified and remedied. Figure 8.2 illustrates this knowledge integration principle. The top portion represents two independent sources of knowledge, perhaps information from two different experts. The bottom portion shows some of the same knowledge provided by both A and B in the white overlap area. If the A and B circles represent knowledge from different experts, then not recognizing the white overlap results in counting the same knowledge twice. A simple example of double counting occurs when gathering historical events on a particular subject and the same event is described in different documents. The event only happened once, regardless of how many times it is cited.

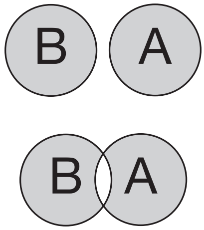

Figure 8.2. Double counting redundant knowledge. Nonoverlapping, independent information from A and B (top); double counting of information from A and B in the white area (bottom).

Recognizing dependency and redundancy is difficult, especially during or after an elicitation. Just because two experts provide the same answer to a question does not necessarily mean they are completely dependent or overlapping. Studies have shown that experts who use similar problem-solving processes produce similar answers—a correlation of cognition and responses.[23] However, correlation is not necessarily dependence, and small degrees of dependence do not result in significant overlap or double counting.[24]

A simple example of dependent experts that does matter is when expert A learns of expert B's answer and decides to copy it, rather than providing A's original answer. In this case, there is only one independent answer even though two experts responded. The dependence or redundancy among experts, as illustrated by expert A's response, primarily comes from expert A's deliberate decision to provide the party-line or community-established response. This social pressure bias can be detected and mitigated through formal elicitation methods.

Dependent relationships are similar to conditional relationships, which are common in most knowledge integration, as described next.

Conditional Combination

Underlying conditions attached to knowledge must be identified before knowledge integration to avoid mixing "apples" with "oranges." Assumptions and dependencies are conditions. Conditions are important because the knowledge can change when conditions change.

For example, history shows that no nuclear weapons were used since World War II. Thus, one might reasonably conclude that the likelihood of such use during that period (i.e., the end of the war to present day) was low.

However, that statement is conditioned on the assumption that close calls have been nonexistent or irrelevant. Because there have been a number of crises in which nuclear use was contemplated (i.e., close calls), and if nuclear use was a high-probability outcome of at least one of these crises, then one might reasonably conclude that the likelihood of nuclear use since World War II was high. The knowledge changed from "low" to "high" when the condition of close calls was considered.

Accounting for all assumptions, conditions, and caveats attached to knowledge is challenging. Conditions inherent in the knowledge may not be readily identifiable. Even an expert supplying the knowledge may not be fully aware of the conditions attached or the assumptions being made. Integration requires care not to combine knowledge having differing conditions. A simple example comparing estimates for the likelihood of nuclear war published in the literature illustrates this (see chapter 3). The authors supply their estimates in different units (different conditions): some provide a per-decade value, some a per-event or scenario value, and some a value without description. These different values cannot be compared or combined until they are all based on common ground, that is, common units. In addition, some of these authors may not be experts. Being an expert in the relevant subject area is a condition for combining expert knowledge.

This example illustrates only a couple of conditions encountered in the risk of deterrence failure problem. Others include scenario description, groups of people or nations involved, time frames, event sequences, subject areas involved, political environments, socioeconomic factors, and human factors.

Some conditions may not be influential and hence do not have to be considered; however, making that determination in a knowledge-poor environment is difficult. The degree of influence or effect of some conditions may not be determinable. In that case, risk assessment is done with a caveat stating that it is unknown what effect, if any, this condition has on the results.

A risk assessment can be done with every quantity conditioned on a particular assumption, such as a chosen scenario. Often risk assessments list these caveats as a caution that the results may differ if the conditions are changed. A simple example of one such condition is a specified time frame for the risk analysis, such as the risk within the next decade.

An implicit condition for all analyses is that the results depend on the knowledge and analysis method used therein. This knowledge includes how the problem was structured, its scope, what knowledge was used in the analysis, how uncertainties were handled, what theory or first principles were applied, and what analysis methods or models were chosen. However, as important as these conditions are to understand, one rarely sees such a detailed statement accompanying a risk assessment.

Inconsistency

Inconsistencies can be found in any form of knowledge. Inconsistencies must be identified and understood before integration to decrease uncertainty and to correct any errors or mistakes. Some inconsistencies are easily detected because they make no sense and are simply errors. For example, the number of member nations in the nuclear club is not one hundred, but it might be ten.

Sometimes an apparent inconsistency is not an actual one because conditions or assumptions have changed. For example, an expert may respond that there are two ways to construct a weapon and then later state there is only one way. After probing, it is discovered that the expert was assuming a certain material was available for the first case but not the second.

Sometimes an apparent inconsistency comes from the failure to recognize the effect of high uncertainty. For example, one expert claims an event will almost surely happen, while another claims that event is nearly impossible. Both experts arrived at their responses using different problem-solving processes, but both responses are valid given the high degree of uncertainty about the likelihood of the event. This high uncertainty is a nightmare for the analyst when presenting results to a decision-maker, as well as for the decision-maker who has to determine a course of action when none is clearly apparent.

In eliciting knowledge from experts, care must also be taken to query why and how an expert apparently switches a reason or response. For example, an expert may state that they cannot answer a particular question because they simply do not know about that subject but then may supply information about that subject later, even to the point of answering the original question. Resolving this and other inconsistencies is done using formal elicitation methods.

Whatever the reason or source of inconsistencies, they must be understood, resolved or remedied, and noted before knowledge integration.

Categorization and Enumeration

Categorization and enumeration can be used for quantification and subsequent integration, analysis, and assessment. When the majority of knowledge for a topic is in qualitative form and involves linguistics, as may occur in the risk of deterrence failure, it is difficult to combine verbal descriptions. However, in some cases essay responses from experts or historical records can be categorized. If that is possible, then counting the pieces of knowledge for each category is a form of quantification. This activity is also an integration method. In addition, it provides numerical results for analysis because enumerations can produce percentages or proportions relative to all categories. Categorization and enumeration are commonly used in the data analysis of surveys.

Both categorization and enumeration activities can involve uncertainty. Figure 8.3 shows how two kinds of uncertainty are involved in making a decision about a terrorist state: is it manufacturing weapons of mass destruction or not? The preponderance of existing knowledge about this state tips the scale and seesaw toward the weapons of mass destruction side. However, a new piece of knowledge emerges about this state acquiring uranium. Because the amount and which isotope(s) were obtained are unknown, it is not certain where the white block belongs: on the weapons of mass destruction (WMD) scale or the No WMD scale. The white block may partially belong on each of the scales, as illustrated with dashed arrows at the top of Figure 8.3. The uncertainty involved in making the categorization can be handled by probability. The probability that the uranium acquisition is for the WMD scale is 0.7, and the probability that it is for other purposes is 0.3. In probability theory, those two assignments sum to 1.0; however, for other uncertainty theories, such as possibility, that is not a requirement.

The second kind of uncertainty is illustrated at the bottom of Figure 8.3—the uncertainty in forming distinctively concise categories. The continuous seesaw at the bottom of Figure 8.3 depicts the inability to precisely distinguish between activities pertaining to the manufacture of a weapon of mass destruction and legitimate related activities (e.g., nuclear power and research reactors). The indeterminacy of the exact (crisp) boundaries of the categories WMD or No WMD makes those two categorizations fuzzy sets.[25]

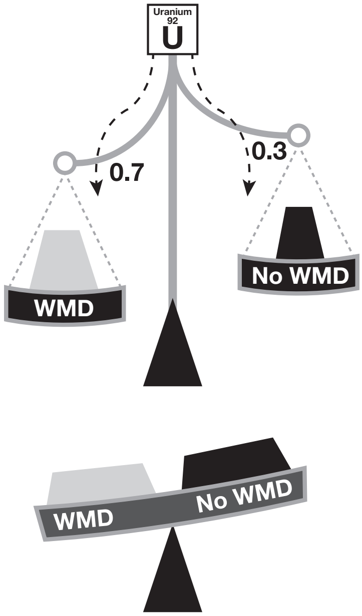

Figure 8.3. Enumeration and categorization uncertainties. Illustration of uncertainty in assignment of knowledge (white block) to two crisp sets (top) and determining the boundaries of two fuzzy sets (WMD and No WMD) (bottom).

Logic and Rule-Based Combinations

Alternatives to mathematical integration formulae are logic and rule-based combinations. Rules and logic are the ways to combine qualitative and linguistic knowledge. These describe the relationships existing among the issues, events, knowledge, experts, etc., involved.

Related to conditional integration logical rules are if–then rules. For example, the statement “If A occurs then B does not occur” is an if–then relationship that describes how to combine A and B. A string or series of if–then rules dictates which items or statements coincide and which are unrelated.

Other logical rules offer the same benefits—providing guidance on how pieces of knowledge are or are not related and dictating how they are combined. For example, logic dictates that there are some minimal requirements for constructing and delivering a weapon of mass destruction. Each of these steps or acquisitions must be included; otherwise weapon construction is not possible. For the problem of constructing and delivering a weapon of mass destruction, knowledge must be gathered about each of those steps and combined according to how that weapon of mass destruction can be produced and delivered. This is an example of the use of the AND logic operator, where each step must be accomplished for an achieved goal.

An example of the OR logic is when alternatives or options are present such that any one is all that is necessary. For example, a dirty bomb requires that some radioactive material be dispersed. However, there are multiple types and sources of radioactive material that can be used, and only one is minimally necessary.

Logic operators such as AND, OR, and NOT are used to connect and combine. In addition to these common crisp logic operators, there are also fuzzy logic counterparts.[26] Fuzzy logic[27] is useful for relationships and combinations that are uncertain, usually because of a lack of knowledge. Often these relationships involve linguistic descriptions rather than numbers. For example, an expert may answer a question about a terrorist group as follows: "Well, *if* this group gets more radical in its beliefs than it is now, *then* it would provide sufficient funding for making a nuclear device." The words *more radical* and *sufficient funding* are fuzzy quantities in this if–then statement. Fuzzy sets and logic provide the mechanisms for how to quantify linguistic statements and descriptions and how to compare and/or combine statements or rules from multiple experts.[28]

Inference-Based Combination

In keeping with the theme of providing some fundamental integration methods, inference-based combination provides a way to combine multiple sources of knowledge at any level of detail. The sources being combined are related to each other by some degree of inference, such as a similar problem or scenario (analogical inference); a related quantity (proxy inference); a relevant model, theory, or computation (validation inference); or a prediction (prediction inference).[29]

A simple example, using estimates and data for nuclear weapon use, illustrates how to combine knowledge from multiple sources. Figure 8.4

depicts four sources of knowledge available: two sources from history (top row) and two sources of estimates made by experts or authors found in the literature (bottom row).

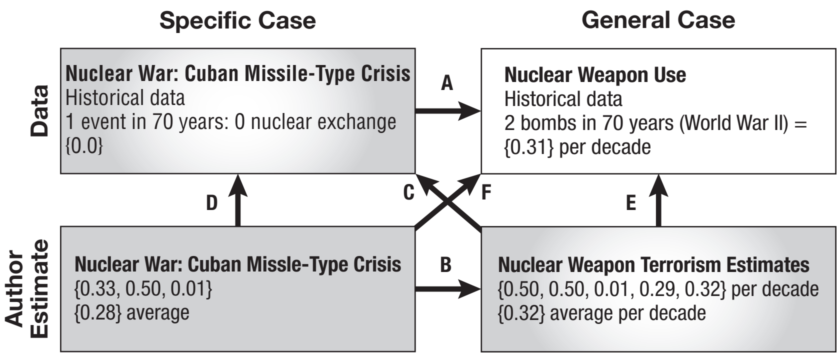

Figure 8.4. Four-box inference technique for nuclear weapon use: combining specific and general historical data with author estimates.

The boxes in the right column refer to the general case of nuclear weapon use in war or terrorism, while the boxes in the left column contain information for a specific event leading to a nuclear exchange—a Cuban-missile-type crisis. The three gray boxes contain knowledge that is “similar” but not identical in quality or relevance to the sparse data in the white box. The white box contains what little, if any, knowledge exists for the problem at hand. If sufficient amounts of knowledge were available for this box, there would be no need to combine that knowledge with the sources in the other boxes. Thus, using the knowledge in the gray boxes to represent the white box is making an inference about its degree of applicability to the white box. The arrows A–F indicate the inference being made and point toward the more important or relevant box from the supporting boxes.

The goal is to combine the knowledge in the three gray boxes with the white box, accounting for the inferences and their uncertainties. This is done using a weighting scheme where the weights for the knowledge in each box are determined based on the degree of inference (the arrows) between boxes. Experts are usually the resource used to determine the degree of inference (the similarity of each box relative to another) for each of the six arrows. Experts assign a value for the degree of relevancy using a numerical comparison scale modified from the pairwise comparison scale developed by Saaty.[30] A simple example for how to apply Saaty’s method to determine the weights for combining the boxes is provided in Langenbrunner et al.[31] Quantifying and combining the corresponding uncertainties for these arrows and for the knowledge in the boxes are more complicated.

Expert-Supplied Integration

Subject-matter experts are often the best sources for determining how knowledge is combined, what combination approach is appropriate, what weights should be used, and what uncertainties apply. However, experts may not be aware of some of the methods for these determinations. Thus, it is the responsibility of the interviewer and analyst to inform experts and to recognize what methods may apply, based on the experts' descriptions. For example, an expert may be thinking about a complicated functional combination method but is unable to write down the formula. The analyst or interviewer recognizes this and provides the expert with some formulations and explanations to determine whether any of these are consistent with the expert's thinking.

It is common for experts to be unaware that they are expressing an uncertainty, especially a nonprobabilistic one. Again, it is the job of the interviewer or analyst to recognize the uncertainty and to clarify its meaning with the expert.

Knowledge integration often requires the cooperation and coordination of different experts: subject-matter experts, experts on elicitation, experts on knowledge integration methods, and experts on uncertainties. Previously noted integration efforts among experts include agreement on common quantities necessary for risk assessment and on how to transform various forms of knowledge into those quantities. Experts may also have to provide and agree on the conditions for the problem structure, such as a given scenario, and the types of uncertainties inherent in the knowledge and in relationships among problem issues.

For quantities or issues that have indeterminate relationships yet require combination, it is possible that no expert is be able to identify an appropriate integration method. In cases in which no one knows how to combine things, a decision can be made to assume some simple combination scheme with the realization that a better integration approach may be available in the future. The inability to find an integration method may be due to a lack of knowledge about the subject or to a lack of good choices of integration methods currently in existence. Either way, this shortcoming and the assumptions made to circumvent it should be clear caveats accompanying any final answers or conclusions presented to decision-makers.

Uncertainty Combination

Throughout this chapter, different types of uncertainty have been mentioned that arise when doing knowledge integration: imprecision, inconsistency, nonspecificity, probability, the unknown, fuzzy, and likelihood. Because uncertainty is common to all knowledge throughout a risk assessment problem, it can be considered a common quantity. The precedent for this is in probabilistic risk assessments where probability (probabilistic uncertainty) is a common quantity for determining the likelihood component of risk.

In the past decade, risk assessment tools have expanded to include other mathematical theories of uncertainties. For example, possibility theory has been used to assess the risk of terrorism.[32] The advantage of using possibility theory over probability theory is that the axioms for possibility are more general, and less restrictive, than those for probability theory. Possibility is better suited to rare-event estimation, as evidenced by the common expression "That is possible but not probable." The disadvantage of using an alternative to probability theory is that most experts and decision-makers will be unfamiliar with it and how to interpret it. For all its faults, probability theory has a long history; many experts and decision-makers have at least heard about it, and some even understand it (although far fewer truly understand it than those who think they do).

Ideally, analysts would be able to work with experts to quantify each type of uncertainty with its appropriate mathematical theory, propagate and combine these uncertainties for an uncertainty estimate attached to the final or top-level answer, and then explain what it means to a decision-maker. However, insufficient research has been done to understand how to mix and match different uncertainty theories, let alone how to explain them to experts and policy-makers. An example of one such difficulty is when the integrated result of different uncertainty theories is desired to be in a familiar form, such as probability, for conveyance to a policy-maker. An uncertainty from a general uncertainty theory combined with an uncertainty from a more restricted one (e.g., probability) can force the combination to follow the more restricted theory. That result changes the interpretation and reason for using the more general theory in the first place.[33] In other words, a false sense of precision may be imposed on the integrated uncertainty that is not warranted given its constituent uncertainties.

Currently, research and experience of application is available for linking fuzzy membership functions with probability distribution functions.[34] Short of other uncertainty theory mixing techniques and experience, the familiar probability theory continues to be used for every uncertainty as is done in probabilistic risk assessment.

An alternative strategy for handling different kinds of uncertainties would be to select one of the most general uncertainty theories, such as imprecise probability,[35] and to characterize every uncertainty and integration by using that theory. An advantage of choosing imprecise probability is that this theory has a probabilistic nature, meaning it can be explained to experts and decision-makers. However, such use of imprecise probability theory would be breaking new application ground. Another theory to consider applying to the entire problem is information gap decision theory.[36] This theory has some history of application and can mathematically accommodate the use of multiple uncertainties within its framework, including probability.[37] Either of these two general uncertainty theories would be worth considering for integrating the different types of uncertainties inherent in the risk assessment for the failure of deterrence problem.

Managing Uncertainties

With no clear solution about how to combine uncertainties, the key may be to manage uncertainty.[38] The first step to managing uncertainties is becoming aware of the uncertainty types; of what knowledge is available; and of the limitations of the experts, analysts, and decision-makers in dealing with uncertainties. One tool that is currently being used is creation of an uncertainty inventory.[39] A quick uncertainty inventory for the risk of failure of nuclear deterrence problem could well reveal something like the following:

- Uncertainties of many types will exist.
- The most common uncertainty will be lack of knowledge—that is, "we just don't know." Unfortunately that type of uncertainty is not the kind that probability theory is designed to quantify.

- Another common uncertainty would be fuzzy from the use of linguistic terms.
- Applying available knowledge to the problem will require making inferences.[40]
- Reliance will be placed on experts as a source of knowledge, uncertainty, and integration methods.
- Not every uncertainty can be quantified, even using ordinal measures.
- Clear choices for handling and combining uncertainties do not exist, and application experience of the more exotic uncertainties is lacking. However, uncertainties cannot be ignored.

The goal is to get an integrated answer to the top-level question: What is the risk of deterrence failure? An integrated answer that ignores uncertainties will be incorrect, as shown in Figure 8.5. The real answer, denoted by an asterisk, is captured only when the uncertainties of the seven data points are considered. An integrated answer using overly large (e.g., anything is possible) uncertainties will be indeterminate. The risk would be anywhere from zero to doomsday. The best that can be done is to make every attempt to utilize all available knowledge and document how and why uncertainties were determined. This is managing uncertainty, and there are some simple methods and ideas for management.

Most humans (experts and decision-makers) can understand the uncertainty involved as expressed in an interval of values and in relative comparisons. Interval arithmetic can be used for combining intervals.[41] Combining comparisons is not as straightforward, but techniques from decision analysis may be useful, such as the pair-wise comparisons used in Figure 8.4.

Using defined words and concepts such as likelihood (rather than probability) prevents tying the expert or the analyst to any particular uncertainty theory. In probabilistic risk assessments and other probability-based analyses, combinations are often done with simulations of probability distribution functions. Similarly, simulations can be used to combine likelihood functions.

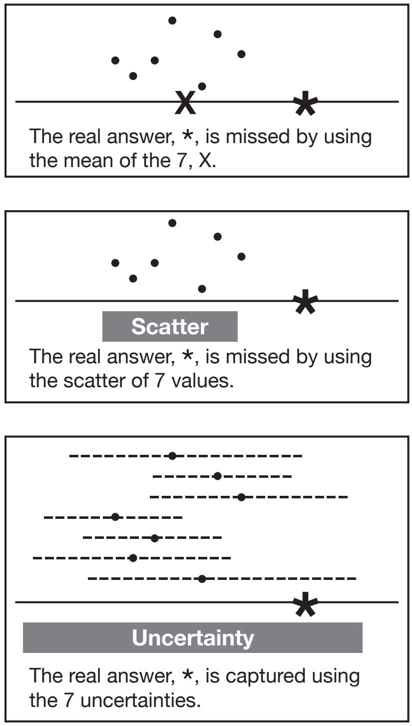

Figure 8.5. An integrated answer must consider uncertainties. Top box: no uncertainty for seven values. Middle box: use scatter of seven values for uncertainty. Bottom box: use uncertainties for all seven values.

One of the tenets of formal expert knowledge elicitation is to reveal uncertainty in the terms used and understood by the expert. That rule applies to integration with the addition that integration across experts may require thorough understanding of any subtle differences in definitions of terms describing uncertainty.

Experts are useful for providing a reality check for an integrated answer and corresponding uncertainty. Either or both may end up unreasonably distorted if an inappropriate integration was done. Often a large uncertainty for the integrated answer can be traced back to one or a few large individual uncertainties. Large uncertainties preclude definitive decisions and result in broad ranges of values of risk and its constituents. When sources of large uncertainties are identified, the decision-maker can be informed of where invested resources can reduce them, and hence reduce the large uncertainty in the final answer. It helps to show decision-makers how this can occur through a what-if demonstration: *What if* a dominant uncertainty is reduced by half? *Then* the final uncertainty is reduced by one-third; therefore, investing time and money for this reduction is valuable.

Managing uncertainty also involves understanding how it relates to making predictions and to inconsistencies among knowledge sources. Some mathematical relationships—trade-offs—between uncertainties and prediction have been established and may prove useful for combining and managing uncertainties.[42]

Rather than quantifying potentially large uncertainties, it may be prudent to assume some reasonable value (usually provided by an expert). This assumption is then identified as a placeholder unless and until more knowledge becomes available that would provide a better uncertainty estimate. That assumption is also a condition (caveat) on which the entire analysis and conclusions rest.

Regardless of the care and documentation implemented in characterizing and combining uncertainties, there will be criticisms and questions about them. Constructive criticisms are welcome because they offer a source of additional knowledge and potential alternative methods. Questions should be answerable from the complete and traceable documentation that has been created.

Knowledge Integration Challenges and Benefits

Some of the challenges and benefits of knowledge integration applied to the assessment of risk of deterrence failure are described below.

Challenges

One of the biggest challenges is that knowledge is constantly changing; today's prediction is tomorrow's historical data. Complete documentation of the knowledge and analysis of it are important for updating that documentation when new knowledge becomes available, necessitating a new integration.

Another challenge for assessing the risk of deterrence failure is the heavy reliance on experts as a primary source of knowledge, including using them to determine how to integrate that knowledge and how to characterize its uncertainty. However, formal elicitation methods are available to aid in this endeavor. Because of the lack of data for this problem, statistical methods are of limited use for analysis, summarization, integration, prediction, and drawing conclusions.

Uncertainties of different types abound for the risk of deterrence failure problem, with lack of knowledge being a major type. Although probability is the common theory for uncertainty, it is not appropriate for many uncertainty types and is not suggested when eliciting uncertainty from experts. Alternative uncertainty theories exist; however, application for many of these is limited.[43] Additional research is needed to provide methods for mixing different theories within a problem.

Care must be taken to identify and accommodate the conditioning factors of the knowledge when applying integration methods or principles. Included in these conditionings are dependence and double counting of the same knowledge.

Existing knowledge integration methods and studies that have been developed were applied to problems involving physical systems. The extent to which these methods are applicable to the ill-posed structure of the deterrence problem has yet to be determined.

Benefits

The goal of a risk assessment is to convey the risk and its uncertainty to decision-makers. That goal does not require the integration of likelihood with consequence, per se; however, it necessitates integration of the likelihoods and consequences over all the parts of the problem. For the deterrence failure problem, where data are sparse, all sources of data, information, and knowledge must be utilized and combined, requiring integration methods. The unique challenges presented for this problem make it difficult to apply traditional risk methods such as probabilistic risk assessment; however, the principles and methods presented here offer some solutions for assessing risk. It should not be too difficult to explain these fundamental methods to the decision- or policy-maker.

As with any thorough assessment process, the risk assessment for this problem will provide the opportunity to learn about aspects of the problem not obvious from a cursory examination. Lessons learned about how to manage uncertainties should produce insights about making decisions in the sparse knowledge environment. Risk analysts may need to work more closely with the decision-maker in order to convey the impact of these uncertainties on conclusions and decision choices.

The integration required for this problem may also require experts from different subject areas to work together or at least understand how their knowledge fits inside the larger problem covering various disciplines. Insights are gained through this process as well.

Understanding where the gaps in knowledge and placeholders are is a planning tool for the experts, analysts, and decision-makers. These "holes" are areas where improvements can be made, perhaps with investment of resources.

Having an updatable integration methodology permits demonstration of how results (and decisions) can change if/when new knowledge surfaces. The benefit of careful and complete documentation is that the knowledge can be used in the future and all the participants (e.g., the experts, the decision-maker, the analyst) can have a productive experience and speak favorably of their involvement in a well-designed and implemented, defensible study. Some examples of integration methodologies exist, for problems with uncertainties of different types, where heavy reliance is placed on expertise as a knowledge source and where different sources of data, information, and knowledge are combined. The methods and principles presented in this chapter have their origins in those studies:

- Reliability methodology, Performance and Reliability Evaluation with Diverse Information Combination and Tracking (PREDICT), 1999[44]
- Yield estimation prediction protocol, 2006[45]
- Inference uncertainty integration methodology (the four-box approach), 2010[46]

Summary

Once problem structure(s) have been determined and knowledge-gathering activities are ongoing or have been completed, knowledge integration becomes the critical step for achieving the goal of providing top-level answers, summaries, and conclusions to policy- and decision-makers.

Knowledge integration extends data-based and multiple study combination analysis methods in new directions. One extension is to combine all forms of data, information, and knowledge. Another is to characterize and combine different types of uncertainties, most of which are not appropriate for probability theory. Finally, a risk assessment for the deterrence failure problem involves additional combinations such as different experts (providing knowledge), different subject areas, and different scenarios or problem formulations.

Accommodating all these integrations brings new challenges not previously addressed by traditional risk assessment methods such as probabilistic risk assessment. Some of the research necessary to do these integrations has yet to be developed. Yet, timely integrations are necessary and must also be conveyed to policy- and decision-makers, as well as to experts involved in providing the knowledge. Therefore, some fundamental principles and methods are provided for present use.

Among the principles involved is the use of formal expert knowledge elicitation methods because experts are valuable resources for providing the knowledge, characterizing the uncertainties, and determining appropriate integration rules or schemes. Another principle is to waste nothing—gather and utilize all available data, information, and knowledge because of its sparseness and high uncertainty.

Following the integration principles and methods should provide the desired top-level or problem solution in terms conveyable to a decision-maker. In addition, these methods are designed to permit the necessary traceability to answer inquiries and update as new knowledge becomes available.

Acknowledgments: Thanks go to James Scouras, Andrew Bennett, Edward Toton, and James Langenbrunner for their constructive thoughts, knowledge, and comments regarding the planning and drafting of this chapter.

Notes

1. Mary A. Meyer and Jane M. Booker, *Eliciting and Analyzing Expert Judgment: A Practical Guide* (Philadelphia: Society for Industrial and Applied Mathematics, 2001).
2. Information also has a technical definition in information theory, founded by C. E. Shannon, dealing with communication, signals, computation, storage, and transmission of information. See Claude E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal* 27 (July, October 1948): 379–423, 623–656, https://doi.org/10.1002/j.1538-7305.1948.tb01338.x.

3. E. T. Jaynes, “Information Theory and Statistical Mechanics,” *Physical Review* 106, no. 4 (1957): 620–630, https://doi.org/10.1103/PhysRev.106.620.

4. L. A. Zadeh, “Fuzzy Sets,” *Information and Control* 8, no. 3 (1965): 338–353, https://doi.org/10.1016/S0019-9958(65)90241-X.

5. H. O. Hartley, “Multiple Frames Surveys,” presented at the Annual Meeting of the American Statistical Association, Minneapolis, MN, September 1962.

6. James Neill, “Meta-analysis Research Methodology,” last updated May 10, 2006, http://wilderdom.com/research/meta-analysis.html.

7. Gene V. Glass, “Primary, Secondary, and Meta-Analysis of Research,” *Educational Researcher* 5, no. 10 (1976): 3–8, https://doi.org/10.3102/0013189X005010003; and Gene V. Glass, “Integrating Findings: The Meta-Analysis of Research,” *Review of Research in Education* 5 (1977): 351–379, https://doi.org/10.3102/0091732X005001351.

8. Dave L. Hall and James Llinas, “An Introduction to Multisensor Data Fusion,” *Proceedings of IEEE* 85, no. 1 (1997): 6–23, https://doi.org/10.1109/5.554205.

9. J. M. Booker, T. R. Bement, M. A. Meyer, and W. J. Kerscher III, “PREDICT: A New Approach to Product Development and Lifetime Assessment Using Information Integration Technology,” in *Handbook of Statistics: Statistics in Industry*, ed. R. Khattree and C. R. Rao, vol. 22 (Amsterdam: Elsevier Science Press, 2003), 499–522; PREDICT is an acronym for Performance and Reliability Evaluation with Diverse Information Combination and Tracking.

10. *Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants*, NUREG-1150 (Washington, DC: US Nuclear Regulatory Commission, 1991), https://www.nrc.gov/reading-rm/doc-collections/nuregs/staff/sr1150/index.html.

11. Cliff Joslyn and Jane M. Booker, “Generalized Information Theory for Engineering Modeling and Simulation,” chap. 9 in *Engineering Design Reliability Handbook*, ed. Efstratios Nikolaidis, Dan M. Ghiocel, and Suren Singhal (Boca Raton, FL: CRC Press, 2004).

12. Meyer and Booker, *Eliciting and Analyzing Expert Judgment*.

13. Richard G. Lugar, *The Lugar Survey on Proliferation Threats and Responses* (Washington, DC: US Senate, 2005), https://irp.fas.org/threat/lugar_survey.pdf.

14. Meyer and Booker, *Eliciting and Analyzing Expert Judgment*.

15. Equal weights is called the *maximum entropy* solution because uniform weights produce the maximum uncertainty when compared to all other unequal weightings; Meyer and Booker, *Eliciting and Analyzing Expert Judgment*.

16. Legal and regulatory communities use weight of evidence, which aligns with statistical logistic regression and likelihood ratio. It is the natural logarithm of the ratio of the number of events to the number of nonevents.

17. James R. Langenbrunner and Jane M. Booker, *To Bayes or Not to Bayes—That Is the Question*, Report No. LA-UR-07-5269 (Los Alamos, NM: Los Alamos National Laboratory, 2007).

18. Thomas Bayes, "Essay towards Solving a Problem in the Doctrine of Chances," *Philosophical Transactions* 53 (1763): 370–418, https://doi.org/10.1098/rstl.1763.0053. Reprinted in *Biometricka* 45 (1958): 293–315.

19. Gudmund R. Iversen, *Bayesian Statistical Inference* (Newbury Park, CA, and London: Sage Publications, 1984).

20. Harry F. Martz and Ray A. Waller, *Bayesian Reliability Analysis* (New York: Wiley & Sons, 1982).

21. *Severe Accident Risks*.

22. *Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants*, WASH-1400 (NUREG 75/014) (Washington, DC: US Nuclear Regulatory Commission, 1975).

23. Jane M. Booker and Mary A. Meyer, "Sources and Effects of Interexpert Correlation: An Empirical Study," *IEEE Transactions on Systems, Man, and Cybernetics* 18, no. 1 (1988): 135–142, https://doi.org/10.1109/21.87061.

24. Correlation means dependence only in the special statistical case where the data are distributed according to a Gaussian or normal distribution.

25. Zadeh, "Fuzzy Sets."

26. George J. Klir and Bo Yuan, eds., *Fuzzy Sets, Fuzzy Logic, and Fuzzy Systems: Selected Papers by Lotfi A. Zadeh* (Singapore: World Scientific Publishing Company, 1996).

27. Fuzzy logic is an example of multi-valued logic, which generalizes binary outcomes restricted to 0 or 1.

28. Ronald E. Smith, Jane M. Booker, Thomas R. Bement, Mary A. Meyer, William J. Parkinson, Mary A. Meyer, and Mohammad Jamshidi, "The Use of Fuzzy Control System Methods for Characterizing Expert Judgment Uncertainty Distributions," in *Proceedings PSAM 4 International Conference* (September 1998): 497–502, https://www.osti.gov/biblio/296889.

29. James R. Langenbrunner, Jane M. Booker, François M. Hemez, and Timothy J. Ross, "Inference Uncertainty Quantification Instead of Full-Scale Testing," presented at the 10th Non-Deterministic Approaches Conference, Schaumburg, IL, April 7–10, 2008, https://doi.org/10.2514/6.2008-2157.

30. Thomas L. Saaty, *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation* (New York: McGraw-Hill, 1980).

31. Langenbrunner et al., "Inference Uncertainty Quantification."

32. John L. Darby, *Estimating Terrorist Risk with Possibility Theory*, Report No. LA-14179 (Los Alamos, NM: Los Alamos National Laboratory, 2004), https://www.osti.gov/servlets/purl/836683.

33. In general, an intersection (multiplicative) combination results in the more restrictive component dominating (the minimum), and the union (additive) combination results in the more general component dominating (the maximum). Timothy J. Ross, personal communication, August 2010.

34. Nozer D. Singpurwalla and Jane M. Booker, "Membership Functions and Probability Measures of Fuzzy Sets," *Journal of the American Statistical Association* 99, no. 467 (2004): 867–877, https://doi.org/10.1198/016214504000001196.
35. Peter Walley, *Statistical Reasoning with Imprecise Probabilities* (London: Chapman and Hall, 1991).
36. Yakov Ben-Haim, *Info-Gap Decision Theory: Decisions under Severe Uncertainty*, 2nd ed. (Oxford and Amsterdam: Academic Press, 2006).
37. François M. Hemez and Yakov Ben-Haim, "The Good, the Bad, and the Ugly of Predictive Science," presented at the 4th International Conference on Sensitivity of Model Output, Santa Fe, NM, March 8–11, 2004.
38. James R. Langenbrunner, Los Alamos National Laboratory, personal communication, April 2010.
39. James R. Langenbrunner, Jane M. Booker, François M. Hemez, Timothy J. Ross, and Isaac F. Salazar, "An Uncertainty Inventory Demonstration—A Primary Step in Uncertainty Quantification," presented at the 11th AIAA Non-Deterministic Approaches (NDA) Conference, Palm Springs, CA, May 4–7, 2009, https://doi.org/10.2514/6.2009-2297.
40. Research into inference uncertainty is currently under way: Timothy J. Ross, Jane M. Booker, and James R. Langenbrunner, "Quantification of Inference Uncertainty in Scientific and Social Modeling/Forecasting Applications," DTRA Basic Research Program, University of New Mexico, 2009–2010.
41. Scott Ferson and Vladik Kreinovich, "Modeling Correlation and Dependence among Intervals," Academia.edu, 2006, https://www.academia.edu/26477123/Modeling_Correlation_and_Dependence_Among_Intervals.
42. Hemez and Ben-Haim, "The Good, the Bad and the Ugly."
43. Fuzzy sets and logic are the exception, with a history of application since 1965.
44. Mary A. Meyer, Jane M. Booker, and T. R. Bement, "PREDICT—A New Approach to Product Development," 1999 R&D 100 Award winner, *R&D Magazine* 41 (September 1999): 161; *AMSTAT News* 267 (August/September 1999): 26; and Booker et al., "PREDICT."
45. J. R. Langenbrunner, *The Four-Star General Was Allocated 15 Nuclear Shots*, Report No. LA-UR-06-5077 (Los Alamos, NM: Los Alamos National Laboratory, 2006).
46. Ross, Booker, and Langenbrunner, "Quantification of Inference Uncertainty."