Risk Analysis, Vol. 18, No. 4, 1998

Use of Technical Expert Panels: Applications to Probabilistic Seismic Hazard Analysis*

Robert J. Budnitz,[1,8] George Apostolakis,[2] David M. Boore,[3] Lloyd S. Cluff,[4] Kevin J. Coppersmith,[5] C. Allin Cornell,[6] and Peter A. Morris[7]

Probabilistic Seismic Hazard Analysis (PSHA) is a methodology that estimates the likelihood that various levels of earthquake-caused ground motions will be exceeded at a given location in a given future time period. Due to large uncertainties in all of the geosciences data and in their modeling, multiple model interpretations are often possible. This leads to disagreements among the experts, which in the past has led to disagreement on the selection of a ground motion for design at a given site. This paper reports on a project, co-sponsored by the U.S. Nuclear Regulatory Commission, the U.S. Department of Energy, and the Electric Power Research Institute, that was undertaken to review the state-of-the-art and improve on the overall stability of the PSHA process, by providing methodological guidance on how to perform a PSHA. The project reviewed past studies and examined ways to improve on the present state-of-the-art. In analyzing past PSHA studies, the most important conclusion is that differences in PSHA results are commonly due to process rather than technical differences. Thus, the project concentrated heavily on developing process recommendations, especially on the use of multiple experts, and this paper reports on those process recommendations. The problem of facilitating and integrating the judgments of a diverse group of experts is analyzed in detail. The authors believe that the concepts and process principles apply just as well to non-earthquake fields such as volcanic hazard, flood risk, nuclear-plant safety, and climate change.

KEY WORDS: Experts; probabilistic; seismic; hazard.

1. INTRODUCTION

This paper, which is a summary of a longer report[1] that contains much more detail than is feasible in a journal article, discuss some of the results of a project whose objective was to provide methodological guidance on how to perform a probabilistic seismic hazard analysis (PSHA). Both technical (earth-sciences) guidance and process (for multiple-expert elicitation and aggregation) guidance were developed by the project, with a strong emphasis on the latter.

This paper provides an overview of the process guidance. For technical guidance on the earth-science aspects, the reader is referred to the full project report.[1]

* This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for any third party's use or the results of such use, of any information, apparatus, product, or process disclosed in this report, or represents that its use by such third party would not infringe privately owned rights. The views expressed in this paper are not necessarily those of the U.S. Nuclear Regulatory Commission.

[1] Future Resources Associates Inc., 2039 Shattuck Avenue, Suite 402, Berkeley, California 94704.
[2] Massachusetts Institute of Technology.
[3] U.S. Geological Survey.
[4] Pacific Gas and Electric Company.
[5] Geomatrix.
[6] Stanford University.
[7] Applied Decision Analysis, Inc.
[8] To whom all correspondence should be addressed.

0272-4322/98/0800-0463$15.00/1

We believe that the procedural principles outlined here are of much more general applicability than just for seismic-hazard analysis. When we use the term "PSHA" here, it can frequently be simply replaced by another subject of inquiry, be it "dose–response relationships" or "global warming."

2. BACKGROUND

Probabilistic seismic hazard analysis (PSHA) is an analytical methodology that provides an estimate of the likelihood that various levels of earthquake-caused ground motions will be exceeded at a given location in a given future time period. The results of such an analysis are expressed as estimated probabilities or estimated frequencies per unit time (such as expected number of events per year.) The ground motion is characterized in any of several different ways (peak acceleration, response spectral ordinates, various velocity measures, etc.).

Unfortunately, this objective of estimating the probability of exceedance of earthquake ground-motions can be attained only with significant uncertainty. Despite extensive advances in seismic knowledge in recent years by a large and active community of researchers around the world, there are still major limits to our understanding of the mechanisms that cause earthquakes, and of the processes that govern how an earthquake's energy propagates. This incomplete understanding means that, when a PSHA analysis is performed, there are inevitably significant uncertainties in the numerical results.

The uncertainties arise for a host of reasons, but the most important is that even in the regions where earthquakes occur fairly frequently so that scientists have a basic understanding of the tectonic regime—such as in coastal California—the scientific data base (specific fault locations, orientations, slip rates, energy dissipation mechanisms, etc.) is still limited. In regions where large earthquakes are very uncommon—such as along much of the U.S. Eastern seaboard or in the American Great Plains—the database is even less able to support scientific understanding of the likelihood of large seismic ground motions striking a given site, because not even the sources or mechanisms of earthquakes can yet be well understood.

This limited understanding has operational implications for the analyst charged with performing a PSHA. Although some of the key inputs to a PSHA can be determined reasonably well from observations or experiments, other key inputs require the judgment of experts, and often the best judgments of the most informed experts differ markedly, meaning that there is significant uncertainty about some of the technical issues at hand. Specifically, the limited information from actual earthquakes, either observed by humans (with or without modern instruments) or inferred from the paleoseismic record, can be—and often is—interpreted quite differently by different experts. This diversity of interpretations translates into important uncertainties in the PSHA's numerical results. Operationally, a PSHA analyst is faced with how to use these different interpretations properly, incorporating the diversity of expert judgments into an analytical result that appropriately captures the current state-of-knowledge including its uncertainties.

For this project, addressing this situation has meant having to face two different (although related) tasks: (i) to develop *technical* guidance, drawn from the earth sciences, concerning the scientific issues involved in performing a PSHA; and (ii) to develop *process* guidance, drawn mostly from disciplines outside the earth sciences (although anchored in the specific details of PSHA and based largely on PSHA experience), concerning how to cope with the diversity of opinion among the experts about the technical issues.

The project itself emphasized the process guidance because of the authors' conclusion that it is often more difficult to execute the process aspects properly (including how expert interpretations are aggregated) than the technical aspects, and because there exists far less guidance in the literature on all of the process subjects except individual-expert probability elicitation, which was not emphasized by the project because there already exists a large literature on this subject.[2–4] For brevity, we will also not focus here on the multiple-expert process guidance found in several prior studies that formed a foundation for our guidance.[4–6] Past experience with expert panels in PSHA was also a key starting point.

As part of the project, the integration/aggregation methodology was tested in the seismic-ground-motion area through a special workshop, and found to be very useful. The details about this trial application can be found in Ref. 1. The process has since been applied to the assessment of volcanic hazard at a Nevada site.[7]

3. PHILOSOPHICAL APPROACH

Although there is general concurrence among PSHA practitioners regarding the purpose and goal of a PSHA, experience has demonstrated the importance of establishing a sound philosophical approach for developing the rules and guidance that are provided here.

As a central element of the project's specific methodological guidance, three elements of our philosophy merit discussion: (i) the task of technical integration; (ii) the use of "experts" and the meaning of "consensus"; and (iii) performing a PSHA using different levels of effort.

3.1. The Task of Technical Integration

One of the major points of emphasis is methodological guidance on *technical integration.* As we use this term, it represents a function that every PSHA project requires. We have defined two different types of integrator, which we have called the Technical Integrator (TI) and the Technical Facilitator/Integrator (TFI), respectively. These differ in that the latter includes the function of "facilitator" of a group of experts (see below), whereas the former does not.

The TI or TFI is a single entity who has the responsibility and is empowered to represent the composite state of information of the technical community regarding a technical issue. Because the TI or TFI must have the stature and expertise to deal convincingly with the multiplicity of disciplines and individuals involved, it would be unusual that one individual could be identified with all the needed qualities. It is more reasonable to anticipate that the TI or TFI will consist of a small group of individuals, typically two or three, covering different areas of expertise including not only the substantive (earth-sciences) and the analytical (PSHA-analysis) areas but also the process (expert-elicitation) area.

The scope of the TI's or TFI's work and the interaction with representatives of the technical community varies greatly with the complexity of the study. However, the task of "technical integration" always has the same thrust, regardless of the scale of the PSHA study. In our view, that thrust is to represent the center, the body, and the range of technical interpretations that the larger informed technical community would have if it were in its entirety to conduct the study or be brought in as a part of the study. Thus, "technical integration" requires assimilation of the full range of informed technical views at large. Here the word "informed" implies not only the necessary prior technical training and experience but also in-depth knowledge of the specific issue at hand.

3.2. The Use of "Experts" and the Meaning of "Consensus"

In writing the guidance, we gave careful attention to the role of "experts" in the PSHA process. We have identified several different *types* of experts and *roles* for experts, ranging from the narrow type (a substantive expert in a very specific technical subject) to the very broad type (an expert with experience across a technical field); and also ranging from the role of proponent of a particular interpretation to that of an evaluator of the full state-of-knowledge of the subject at issue. Consider the following possible types of consensus:

Consensus Type 1. Each expert believes in the same deterministic model or the same value for a variable or model parameter.

Consensus Type 2. Each expert believes in the same probability distribution for an uncertain variable or model parameter.

Consensus Type 3. All experts agree that a particular composite probability distribution represents them as a group.

Consensus Type 4. All experts agree that a particular composite probability distribution represents the overall scientific community.

We seek Type 4 consensus, which is potentially the easiest type of consensus to achieve. There is reason to be far more optimistic that a multi-expert process can achieve legitimate Type 3 or 4 representational consensus than that such a process can achieve the more traditional Type 1 or 2 technical consensus. In our methodology, the issue is not consensus on scientific issues, which may be almost impossible to achieve for some issues at any given time. We have found that it is far easier for a group of experts, when they have legitimate scientific disagreements, to agree on how to represent the informed community's legitimate diversity of opinion about a scientific issue, than it is for the experts to agree on specific technical issues.

3.3. Performing a PSHA Using Different Levels of Effort

We have concentrated our methodology-development work on guidance for a sponsor and analysis team whose financial and personnel resources would be sufficiently large that they would not significantly limit the scope of the PSHA analysis. However, some sponsors may not be able to devote vast resources to a PSHA project, or may not even require a PSHA assessment of the extreme potential ground motions that would be associated with very rare events. In these cases, a scaled-down approach may be appropriate. To accomplish a range of scaled-down analyses, we have identified four different "levels," from a simple literature review to the highest level employing the TFI. While the full report[1] provides guidance on all of these levels, in this paper we will not discuss the three lower levels but will concentrate on the highest (TFI) level.

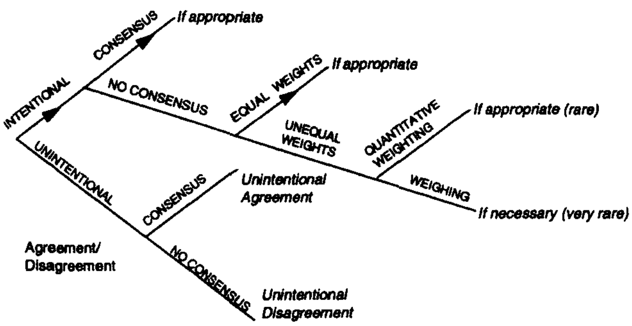

Fig. 1. Overview of the TFI Process. The objective is to create conditions for exiting the top half of the tree at the earliest possible stage that is a legitimate stopping point.

4. PROCESS GUIDANCE: OVERVIEW

Since this article summarizes more extensive guidance provided elsewhere,[1] it is not possible to do justice to the full range of our process guidance. However, the following brief summary will outline the highlights.

Figure 1 provides a "roadmap" for the procedural logic. Reading left to right, the tree indicates increasingly less desirable final process outcomes. Paths with an arrowhead indicate desirable (and expected) process outcomes. The TFI's job is to organize a process that will exit the tree at the earliest possible point, while at the same time making sure that this is a legitimate stopping point.

The process is founded on a specific objective: to use the expert panel to represent the overall scientific community's state of knowledge. The underlying premise is that the primary objective for public policymaking is not capturing the composite judgment of any specific subset of experts (including the panel), but rather capturing as best one can the composite judgment of the overall scientific community of informed experts.

Of course, it is impractical to engage an entire scientific community in any meaningful interactive process. Thus, the panel is viewed as a sample of the overall expert community.

Among the most important guidance is that, regardless of the level of a PSHA study, certain key attributes are essential. These include a reasonable representation of the diversity of interpretations that exist in the informed scientific community regarding particular technical issues; a clear description of the technical basis for all assessments; a quantification of the uncertainties associated with the models and parameters in the analysis; an effective peer review that is keyed to the overall project scope; and documentation of all of the above.

The TFI process is centered on using extensive expert interaction as the principal mechanism for integration. In particular, the TFI process methodically creates conditions under which equal weights on the individual expert evaluators are appropriate. Based on Bayesian expert-aggregation theory, equal weights make sense only under certain conditions, namely that the experts either be completely independent or (roughly speaking) equally interdependent, and that they be equally informed and credible. Since it is unrealistic to expect a set of experts with similar backgrounds and information bases to be independent, the TFI process through intensive, highly structured interaction attempts to maximize the overlap in expert databases, models, and reasoning processes. Some of the principles emphasized and more carefully explained in Ref. 1 are described below.

Experts as Evaluators Not Proponents. Viewing the experts as evaluators who provide interpretations of a range of models and data is an attractive alternative compared to viewing the experts as proponents, advocating their own models or assessments. Although the TFI might sometimes ask a panel expert to act temporarily as a proponent, this is solely for the purpose of explaining a particular model.

Emphasis on Expert Interaction. The TFI conducts structured, facilitated discussions among the experts in which the focus is on underlying models and hypotheses, not on individual experts. The process evolves in stages, and in each stage there are intensive face-to-face interaction workshops preceded and succeeded by TFI interaction and elicitation with individual experts.

Isolation of Sources of Disagreement. Experts may disagree about underlying scientific hypotheses; about interpretations of data sets; about the values of model parameters; and about the ranges of uncertainties that affect seismic hazard. Paradoxically, isolating and focusing discussion on the different potential types of disagreement may actually move the group toward agreement on-scientific issues.

Active Listening. This is a useful facilitation tool borrowed from the field of education, in which a person's reasoning is not considered fully understood unless each listener, whether or not that listener agrees with the reasoning, can explain it back to the person who is making the point. The TFI summarizes points of agreement and disagreement, encouraging active listening and frequently playing back a clear summary of the conversation during the meeting.

Tone of the Interaction. It is critical for the TFI to set the right tone in two dimensions: (i) by establishing that the purpose is not to choose the best model or answer (the premise is that there is no one correct model or answer, no single "winner" or "loser"), (ii) by establishing that the purpose is not to achieve consensus of either Types 1 or 2. Consensus may occur, but it is important psychologically for the participants not to feel that the process is failing if everyone does not agree.

Concerning *responsibility for the process*, it is important to note that final responsibility for the process of obtaining the aggregated product rests with the TFI. Both the TFI and the expert evaluators have *intellectual responsibility for the results.*

In fact, our guidance is structured so that the expert evaluators are asked to represent both their own interpretations and uncertainties, and then their view of the entire informed community's composite interpretation. In the latter activity, each of them is acting as an integrator (see below) in evaluating these views.

5. PROCESS GUIDANCE: THE TECHNICAL FACILITATOR/INTEGRATOR (TFI) APPROACH

The Technical Facilitator/Integrator (TFI) approach involves the formal elicitation and aggregation of multiple experts/evaluators. In developing the TFI guidance, we drew both on the literature and on past experience with major PSHA studies involving multiple experts. However, we determined that both the underlying concepts and the implementation could benefit from improvements.

These innovations are necessary because the broadest and most difficult PSHA studies using multiple experts—those with extensive scopes, controversial applications, large budgets, important implications for public safety, etc.—are most in need of process guidance, yet (i) little process guidance exists, especially for multi-expert group elicitation and aggregation, and (ii) practical experience with documented-methods seems to be very limited.[8]

Our analysis of earlier large multi-expert PSHA studies revealed to us that some of the processes followed could be improved, and led us directly (although only after considerable deliberation and several false starts) to the TFI concept. The two PSHA studies analyzed in the most depth are described in Refs. 9 and 10; a sample of other relevant studies and observations may be found in Refs. 4–6. Among the problems we identified in earlier studies were overly diffused responsibility for the composite results; insufficient face-to-face interaction among the experts; confused roles for the experts, who were sometimes asked to be what we have called proponents, sometimes evaluators, and sometimes integrators, but in many cases with confusing or overlapping roles being played simultaneously; the use of inflexible aggregation schemes, decided on *a priori* by the project's sponsors; imprecise or overly narrow objectives as to what the experts were being asked to provide, and what the final analysis was being asked to accomplish; awkward treatment of so-called "outlier" interpretations, including dealing with these in an *ad hominem* way ("outlier experts") instead of based on the technical merits of the interpretations; and inadequate use of diagnostic and feedback tools. The TFI process has been explicitly structured to overcome or mitigate each of these problems.

For example, the issue of outlier experts has been especially contentious in past multiple-expert studies and deserves extra attention here. For our purposes, an outlier expert is defined by two conditions: (i) the expert makes an interpretation *far* different than that of the rest of the experts, and (ii) the expert cannot support the interpretation with solid data or reasoning from the point-of-view of the other experts. A past PSHA study included an expert who attached a probability of unity to future Modified Mercalli Intensity XII earthquakes throughout the Northeastern U.S. If the objective were limited to developing a composite representation of, say, a five-person panel, then the TFI is in a logical "trap" because the outlier expert does, in fact, represent one-fifth of the panel. Common sense says that the expert should be downweighted, but how can this be justified without superimposing the TFI's own judgment on the process? The perspective of developing a composite representation of the overall scientific community affords a way out of the trap. *The panel is a sample of this community and may not be representative statistically.* When asked to identify other supportive experts, the panel, even including the outlier, may agree that he or she is the only one out of a hundred seismicity experts who would attach significant probability to such a large earthquake everywhere. To represent the overall community, if the TFI wishes to treat the outlier's position as equally credible to the other panelists, he might properly assign a weight of one in a hundred to the outlier's position, not one in five. (In different terms, the panel is treated as a "stratified sample.")

In the TFI approach, multiple experts, empaneled together, act *not* as proponents, each of one specific viewpoint, but as informed *evaluators* of a range of viewpoints. (These individual viewpoints or models may be defended by proponent experts invited to present their views and "debate" the panel.) Separately, the experts on the panel also play the role of *integrators*, providing advice to the TFI on the appropriate representation of the composite position of the community as a whole.

In contrast with the classical role of experts on a panel acting as individuals and providing inputs to a separate aggregation process, in the TFI approach the panel is viewed as a team, with the TFI as the team leader, working together to arrive at (i) a composite representation of the state-of-knowledge of the group, and then, (ii) a composite representation of the knowledge of the community at large.

Furthermore, the process is structured to be transparent to the experts at all stages, in contrast to some previous PSHA studies in which experts have complained that the aggregation process was a "black box."

The TFI conducts both individual elicitations and group interactions, and with the help of the experts themselves the TFI integrates data, models, and interpretations to arrive at the final product: a full probabilistic characterization of the seismic hazard at a site, including the uncertainty. Together with the experts acting as evaluators, the TFI "owns" the study and defends it as appropriate.

The TFI's special role only comes into play for an issue that is complex and controversial enough to warrant the challenge and expense of a suite of multiple integrators. The advantages come at a cost of having to aggregate or in some way represent the judgments of a set of diverse experts, a problem that has been a source of major difficulty in some past PSHA projects.

Our TFI guidance distinguishes between two TFI roles, the TFI as *technical facilitator* and the TFI as *integrator*. In this latter role, the TFI is responsible for developing a composite characterization of the state-of-knowledge of the panel and then of the expert community, based on input from the panel of experts/evaluators. The TFI-Integrator role is *not* that of a "super-expert" who has the final say on the weighting of the relative merits of a set of (proponent) models and positions; rather, the TFI attempts to characterize both the commonality and the diversity in a set of panel estimates, each of which may itself represent a weighted combination of models and positions. In our opinion, *the TFI-should be viewed as performing an integration assisted by a group of experts who provide integration advice.*

The detailed guidance in Ref. 1 covers the TFI's facilitator role and integrator role separately. Guidance is provided on issues such as dealing with different degrees of expertise; handling "outlier" interpretations; dealing with non-independent experts; equal-weights vs. non-equal weights; and the level at which aggregation should be done. The guidance emphasizes that the TFI must have a basic understanding of expert-aggregation issues in order to steer the process to the simplest possible integration procedure. The guidance also emphasizes that the TFI does not need to use a prescribed, rigid combination formula, such as a fixed-weighting scheme. Nevertheless, mathematical expert-aggregation models have an important supporting role in the TFI process, and several simplified expert-aggregation models are presented for possible use. Some of the models described in the larger report[1] are presented in Refs. 11–17. The TFI is urged to use these models to check the implications of various assumptions, so that the ultimate aggregation, even if purely nonmathematical in nature, will be sound and defensible.

Many earlier multiple-expert processes have had a single objective such as "achieve consensus" or "elicit and then equally weight individual judgments" or "have the principal investigator choose the 'best judgment' or even the 'best model.'" In contrast, the TFI process does not operate with a single pre-set objective but rather proceeds through a pushdown list of objectives, attempting to achieve the simplest end-state possible. Specifically, while "consensus" and "equal weights" are highly desirable, we have come to recognize that they are only appropriate under certain conditions (described in detail in Ref. 1). However, these conditions can be controlled, and we believe that equal weights on experts as evaluators (not on models), at least, can usually be attained with sufficiently structured intensive expert interaction. Also, as discussed above, there are different types of consensus, each of which has an *a priori* different likelihood of being achievable.

In the rare case in which such simple integration is not appropriate, additional guidance is provided. In Ref. 1, guidance is presented on two possible approaches involving (i) explicit quantitative but unequal weights (especially, as discussed above, when it becomes obvious that using equal weighting misrepresents the community-as-a-whole); and (ii) "weighing" rather than "weighting," in cases when the experts themselves, acting as evaluators and integrators, find that simple fixed numerical weights are inadequate because explicit weighting of different models is understood to be artificial and a more *ad hoc* way to represent the community's overall distribution is appropriate. This approach is particularly effective if the issue under consideration can be reduced to a scalar quantity of uncertain numerical value.

It is noteworthy that in the TFI process, each expert documents and takes technical responsibility for his/her own interpretation. Each expert is also asked to act as an integrator to estimate the community's distribution. However, the TFI is ultimately responsible for ("owns") the composite representation of the expert community, which is based on the individual expert evaluations as well as the various expert-acting-as-integrator estimates of the community's distribution. The TFI is also responsible for documenting and defending how the composite representation was developed, be it by equal weighting of the individual expert estimates of the community's distribution, or, if necessary, by means more appropriate to the particular circumstances, including downweighting or removing an expert under certain carefully controlled circumstances.

Thus, the TFI as *facilitator* structures interaction among the experts to create conditions under which the TFI's job as *integrator* will be easy (e.g., either a consensus representation is formed or equal weights are appropriate).

Guidance on our structured TFI process[1] comprises two stages, the first with expert panelists acting as independent evaluators representing themselves, and the second with these same panelists acting as integrators representing the overall community. A 7-step structured TFI process based on previous successful expert-elicitation studies[4–6,18] is outlined, involving selection of experts, interaction, elicitation training, expert elicitation, analysis, and documentation.[1]

ACKNOWLEDGMENTS

This project was carried out under the three-way sponsorship of the U.S. Department of Energy (DOE), the U.S. Nuclear Regulatory Commission (NRC), and the Electric Power Research Institute (EPRI). NRC's contribution was managed by Lawrence Livermore National Laboratory and DOE's by Sandia National Laboratories. To support the authors' work, a large number of experts on various technical subjects worked under the authors' direction on specific topics that were integral to the effort. These experts included Michael Bohn, Auguste Boissonnade, Martin McCann, Robin McGuire, Richard Mensing, Jean Savy, and Gabriel Toro. We also benefitted from the expertise of the representatives of the sponsors: Jeffrey Kimball and Ann Bieniawski (DOE), Ernst Zurfleuh and Andrew Murphy (NRC), and Carl Stepp (EPRI). The project report was reviewed both at an early stage and when complete by the "Panel on Seismic Hazard Evaluation" of the National Research Council/National Academy of Sciences, and we wish to thank that Panel for very helpful advice.

REFERENCES

1. R. J. Budnitz, D. M. Boore, G. Apostolakis, L. S. Cluff, K. J. Coppersmith, C. A. Cornell, and P. A. Morris, Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts, NUREG/CR-6372, U.S. Nuclear Regulatory Commission (1995).
2. R. L. Keeney and D. von Winterfeldt, "On the Uses of Expert Judgment on Complex Technical Problems," *IEEE Trans. Engin. Mgmt.* **36**, 83–86 (1989).
3. M. A. Meyer and J. M. Booker, Eliciting and Analyzing Expert Judgment, A Practical Guide, Report NUREG/CR-5424, Los Alamos National Laboratory for U.S. Nuclear Regulatory Commission (1990).
4. R. L. Keeney and D. von Winterfeldt, "Eliciting Probabilities from Experts in Complex Technical Problems," *IEEE Trans. Engin. Mgmt.* **38**, 191–201 (1991).
5. H. Otway and D. von Winterfeldt, "Expert Judgment in Risk Analysis and Management: Process, Context, and Pitfalls," *Risk Anal.* **12**, 83–93 (1992).
6. A. R. DeWispelare, L. T. Herren, M. P. Miklas, and R. T. Clemen, Expert Elicitation of Future Climate in the Yucca Mountain Vicinity, Report NRC-02-88-005, Center for Nuclear Waste Regulatory Analyses, San Antonio, Texas (1993).
7. Geomatrix Consultants, Inc. and TRW, Probabilistic Volcanic Hazard Analysis of Yucca Mountain, Nevada, Report BA-01717-2200-00082, June (1996).
8. A. Mosleh, V. M. Bier, and G. Apostolakis, Methods for the Elicitation and Use of Expert Opinion in Risk Assessment: Phase I, A Critical Evaluation and Directions for Future Research, Report NUREG/CR-4962, U.S. Nuclear Regulatory Commission (1987).
9. D. L. Bernreuter, J. B. Savy, R. W. Mensing, and J. C. Chen, Seismic Hazard Characterization of 69 Nuclear Plant Sites East of the Rocky Mountains, Report NUREG/CR-5250, in 8 volumes, Lawrence Livermore National Laboratory for the U.S. Nuclear Regulatory Commission (1989).
10. Electric Power Research Institute, Seismic Hazard Methodology for Central and Eastern United States, Report EPRI-NR-4726, in 10 volumes (1988).
11. R. T. Clemen, "Combining Overlapping Information," *Mgmt. Sci.* **33**, 332–340 (1987).
12. R. M. Cooke, *Experts in Uncertainty: Expert Opinion and Subjective Probability in Science* Oxford University Press, New York, 1991).
13. R. T. Clemen and R. L. Winkler, "Limits for the Precision and Value of Information from Dependent Sources," *Operat. Res.* **33**, 427–442 (1985).
14. R. L. Winkler, "The Consensus of Subjective Probability Distributions," *Mgmt. Sci.* **15**(2), B61–B75 (1968).
15. R. L. Winkler, "Combining Probability Distributions from Dependent Information Sources," *Mgmt. Sci.* **27**(4), 479–488 (1981).
16. S. Chibber, G. Apostolakis, and D. Okrent, "On the Use of Expert Judgments to Estimate the Pressure Increment in the Sequoyah Containment at Vessel Breach," *Nucl. Technol.* **105**, 87–103 (1994).
17. P. A. Morris, "Decision Analysis Expert Use," *Mgmt. Sci.* **20**(9), 1233 (1974).
18. R. Winkler, T. Wallsten, R. Whitfield, H. Richmond, S. Hayes, and A. Rosenbaum, "An Assessment of Chronic Lung Injury Attributable to Long-Term Ozone Exposure," *Operat. Res.* **43**, 19–28 (1995).