NUREG/CR-6372
UCRL-ID-122160
Vol. 1

Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts

Main Report

Prepared by
Senior Seismic Hazard Analysis Committee (SSHAC)
R. J. Budnitz (Chairman), G. Apostolakis, D. M. Boore, L. S. Cluff, K. J. Coppersmith, C. A. Cornell, P. A. Morris

Lawrence Livermore National Laboratory

0/1
0F02

Prepared for
U.S. Nuclear Regulatory Commission
U.S. Department of Energy
Electric Power Research Institute

AVAILABILITY NOTICE

Availability of Reference Materials Cited in NRC Publications

Most documents cited in NRC publications will be available from one of the following sources:

1. The NRC Public Document Room, 2120 L Street, NW., Lower Level, Washington, DC 20555-0001

2. The Superintendent of Documents, U.S. Government Printing Office, P. O. Box 37082, Washington, DC 20402-9328

3. The National Technical Information Service, Springfield, VA 22161-0002

Although the listing that follows represents the majority of documents cited in NRC publications, it is not intended to be exhaustive.

Referenced documents available for inspection and copying for a fee from the NRC Public Document Room include NRC correspondence and internal NRC memoranda; NRC bulletins, circulars, information notices, inspection and investigation notices; licensee event reports; vendor reports and correspondence; Commission papers; and applicant and licensee documents and correspondence.

The following documents in the NUREG series are available for purchase from the Government Printing Office: formal NRC staff and contractor reports, NRC-sponsored conference proceedings, international agreement reports, grantee reports, and NRC booklets and brochures. Also available are regulatory guides, NRC regulations in the *Code of Federal Regulations*, and *Nuclear Regulatory Commission Issuances*.

Documents available from the National Technical Information Service include NUREG-series reports and technical reports prepared by other Federal agencies and reports prepared by the Atomic Energy Commission, forerunner agency to the Nuclear Regulatory Commission.

Documents available from public and special technical libraries include all open literature items, such as books, journal articles, and transactions. *Federal Register* notices, Federal and State legislation, and congressional reports can usually be obtained from these libraries.

Documents such as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings are available for purchase from the organization sponsoring the publication cited.

Single copies of NRC draft reports are available free, to the extent of supply, upon written request to the Office of Administration, Distribution and Mail Services Section, U.S. Nuclear Regulatory Commission, Washington, DC 20555-0001.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at the NRC Library, Two White Flint North, 11545 Rockville Pike, Rockville, MD 20852-2738, for use by the public. Codes and standards are usually copyrighted and may be purchased from the originating organization or, if they are American National Standards, from the American National Standards Institute, 1430 Broadway, New York, NY 10018-3308.

DISCLAIMER NOTICE

NUREG/CR-6372
UCRL-ID-122160
Vol. 1

Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts

Main Report

Manuscript Completed: April 1997
Date Published: April 1997

Prepared by
Senior Seismic Hazard Analysis Committee (SSHAC)
R. J. Budnitz (Chairman), G. Apostolakis, D. M. Boore, L. S. Cluff, K. J. Coppersmith, C. A. Cornell, P. A. Morris

Under Contract to:
Lawrence Livermore National Laboratory
Livermore, CA 94550

Prepared for
Division of Engineering Technology
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001
NRC Job Code L2503

U.S. Department of Energy
Engineering and Operations Support Group
Office of Defense Programs
19901 Germantown Road
Germantown, MD 20874-1290

Electric Power Research Institute
3412 Hillview Avenue
Palo Alto, CA 94304-1395

ABSTRACT

Probabilistic Seismic Hazard Analysis (PSHA) is a methodology that estimates the likelihood that various levels of earthquake-caused ground motion will be exceeded at a given location in a given future time period. Due to large uncertainties in all the geosciences data and in their modeling, multiple model interpretations are often possible. This leads to disagreement among experts, which in the past has led to disagreement on the selection of ground motion for design at a given site.

In order to review the present state-of-the-art and improve on the overall stability of the PSHA process, the U.S. Nuclear Regulatory Commission (NRC), the U.S. Department of Energy (DOE), and the Electric Power Research Institute (EPRI) co-sponsored a project to provide methodological guidance on how to perform a PSHA.

The project has been carried out by a seven-member Senior Seismic Hazard Analysis Committee (SSHAC) supported by a large number other experts.

The SSHAC reviewed past studies, including the Lawrence Livermore National Laboratory and the EPRI landmark PSHA studies of the 1980's and examined ways to improve on the present state-of-the-art.

The Committee's most important conclusion is that differences in PSHA results are due to procedural rather than technical differences. Thus, in addition to providing a detailed documentation on state-of-the-art elements of a PSHA, this report provides a series of procedural recommendations.

The role of experts is analyzed in detail. Two entities are formally defined—the Technical Integrator (TI) and the Technical Facilitator Integrator (TFI)—to account for the various levels of complexity in the technical issues and different levels of efforts needed in a given study.

Disclaimer

This work was supported by the United States Nuclear Regulatory commission under a Memorandum of Understanding with the United States Department of Energy, and performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

CONTENTS: VOLUME 1

FIGURES: VOLUME 1

TABLES: VOLUME 1

CONTENTS: VOLUME 2

Note: Turn to the Contents section of Volume 2 for a complete list of tables and figures in Volume 2.

SPONSOR'S PERSPECTIVE

Probabilistic Seismic Hazard Analysis (PSHA) has become an increasingly important tool for aiding design and decision making at all levels in both the private sector and government. The level of sophistication applied to PSHA has increased dramatically over the past 27 years since the technique was first introduced in the literature. As more and more people and groups implemented and used PSHA in different forms, it became clear to the sponsors of the Senior Seismic Hazard Analysis Committee (SSHAC) report that the time had arrived to establish more uniform and up-to-date guidelines for future PSHA studies.

The need for such guidelines is threefold:

1. As the situation stands today, it is often the case that multiple PSHA studies are available for the same geographic region. However, due to differences in implementation, results of these studies often differ by substantial amounts for the same physical location. Further, because of the amount of technical information and complex combination of techniques utilized, it is not always simple to determine the source of these differences and which answer should be used.

2. Potential sponsors of a PSHA study are faced with the difficulty of determining the appropriate level of a proposed PSHA to ensure stable results that meet the sponsor's needs.

3. The cost to perform a PSHA study can be quite large. The sponsors of this report expected that a suitable set of guidelines could be developed to assist the potential user in choosing the appropriate level of analysis consistent with the overall goals and resources available. Given the need to conserve resources, issuing such guidelines to optimize future PSHA studies in accordance with the sponsor's need takes on added importance.

Overall, the sponsors saw a need for more stability in the PSHA process, both for nuclear and non-nuclear applications, in dealing with future needs for using PSHA to establish seismic hazard levels throughout the United States.

Comparative evaluations have shown that the differences between PSHA studies are often not technical, but due to the information gathering and assembly process used in the study. The integration of the different types of information required in a PSHA (geologic, seismotectonic, probability and statistics, information theory, and decision making) presents significant inter-disciplinary challenges and requires a project structure and process that assure proper integration. The skills required to be a good integrator and evaluator are not necessarily the same skills needed to be a good scientist. Our observation is that although many PSHA practitioners are trained experts in one or more fields, the PSHA divergence issue can partly be explained by a lack of integration and evaluation skills so important to the PSHA product. We believe this is true at all levels of PSHA, and these skill requirements may be most acute at the simpler levels of seismic hazard analysis not associated with critical facility assessments where typically the PSHA analysts must complete their work.

This report addresses the integration and evaluation issues that should be considered and focuses on the process of integration required in a PSHA. The SSHAC's investigations have led to the conclusion that technical facilitation and integration is a necessary component for the proper implementation of a PSHA in some instances. In most of these cases, it is anticipated that following the approaches outlined in the report will bring about more consistent interpretations that are supported by the data or bulk of scientific thought. However, if an outlier interpretation persists, it is our firm belief—in agreement with the SSHAC—that the approaches outlined will allow for essential downweighting of that interpretation. This is

preferable to the stiff adherence to an equal weighting scheme, which can result in the final seismic hazard being driven by a single outlier input.

The issues that are raised and discussed in the SSHAC report, especially but not exclusively the process issues, apply in varying degrees to any PSHA project, and should be at least considered by sponsors and analysts before undertaking a PSHA. While the primary focus of SSHAC was on siting critical facilities, it is believed that all PSHA projects should attempt to achieve several primary objectives: 1) proper and full incorporation of uncertainties, 2) inclusion of the range of diverse technical interpretations that are supported by available data, 3) consideration of site- specific knowledge and data sets, 4) complete documentation of the process and results, 5) clear responsibility for the conduct of the study, and 6) proper peer review. Regardless of the level of the study, the goal in the various approaches is the same: to provide a representation of the informed scientific community's view of the important components and issues and, finally, the seismic hazard.

For these reasons, the sponsors believe that the SSHAC report is complete in terms of outlining the process a principal investigator should follow to complete a PSHA. Indeed, the report provides for technical flexibility where such flexibility is needed and, at the same time, encourages standardization of technical approaches and procedures as much as is feasible.

The future utility of PSHA in decision making depends to a large degree on our ability to implement the process in a meaningful and cost-effective way. Development of the SSHAC guidelines was planned with this goal in mind.

EXECUTIVE SUMMARY

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. The results of such an analysis are expressed as estimated probabilities per year or estimated annual frequencies. The objective of this project has been to provide methodological guidance on how to perform a PSHA. The project, co-sponsored by the U.S. Nuclear Regulatory Commission, the U.S. Department of Energy, and the Electric Power Research Institute, has been carried out by a seven-member Senior Seismic Hazard Analysis Committee (SSHAC), supported by a large number of other experts working under the Committee's guidance, who are named in the following "Acknowledgments" section.

The members of the Senior Seismic Hazard Analysis Committee (SSHAC) are:

Dr. Robert J. Budnitz (Chairman)	President Future Resources Associates, Inc.
Professor George Apostolakis	Massachusetts Institute of Technology previously at University of California, Los Angeles
Dr. David M. Boore	Seismologist U.S. Geological Survey
Dr. Lloyd S. Cluff	Manager, Geosciences Department Pacific Gas & Electric Company
Dr. Kevin J. Coppersmith	Vice President Geomatrix
Dr. C. Allin Cornell	C. A. Cornell Company
Dr. Peter A. Morris	Applied Decision Analysis, Inc.

The scope of the SSHAC guidance is intended to cover both site-specific and regional applications of PSHA (more broadly, applications in both low-seismicity and high-seismicity regions) in both the eastern U.S. and western U.S. Although the sponsors' primary objective is guidance for applications at nuclear power plants and other critical facilities, the methodological guidance applies in whole or in part, on a case-by-case basis, to a broad range of applications.

The SSHAC guidance involves both technical guidance and procedural guidance, with a strong emphasis on the latter for reasons explained below. Therefore, the audience for the report includes not only analysts who will implement the methodology and earth scientists whose expertise will support the analysts, but also PSHA project sponsors—those decision-makers in organizations such as private firms or government agencies who have a need for PSHA information and are in a position to sponsor a PSHA study.

Note that our guidance is not intended to be "the only" or "the standard" methodology for PSHA to the exclusion of other approaches; there are other valid ways to perform a PSHA study. Likewise, our formulation should not be viewed as an attempt to "standardize" PSHA in the sense of freezing the science and technology that underlies a competent PSHA, thereby stifling innovation. Rather, our guidance is intended to represent SSHAC's opinion on the best current thinking on performing a valid PSHA.

The most important and fundamental fact that must be understood about a PSHA is that the objective of estimating annual frequencies of exceedance of earthquake-caused ground motions can be attained only with significant uncertainty. Despite much recent research, major gaps exist in our understanding of the mechanisms that cause earthquakes and of the processes that govern how an earthquake's energy propagates from its origin beneath the earth's surface to various points near and far on the surface. The limited information that does exist can be—and often is—legitimately interpreted quite differently by different experts, and these differences of interpretation translate into important uncertainties in the numerical results from a PSHA.

The existence of these differences of interpretation translates into an operational challenge for the PSHA analyst who is faced with (1) how to use these different interpretations properly, and (2) how to incorporate the diversity of expert judgments into an analytical result that appropriately captures the current state-of-knowledge of the expert community, including its uncertainty.

The SSHAC studied a large number of past PSHAs, including two landmark studies from the late 1980s known as the "Lawrence Livermore (LLNL)" study and the "Electric Power Research Institute (EPRI)" study, both of which broke important new methodological ground in attempting to characterize earthquake-caused ground motion in the broad region of the U.S. east of the Rocky Mountains. Most important, the mean seismic hazard curves presented in the reports for most sites in the eastern U.S. differed significantly. However, the median hazard results did not differ by nearly as much. We now understand that differences in both the inputs and the procedures by which the two studies dealt with the inputs were among the key reasons for the differences in the mean curves. At the time this was not understood, and the differences between the mean curves caused not only considerable consternation, but launched several efforts to understand what might underlie the differences and attempts to update the older work.

Ultimately, the inability to understand all of the differences between the LLNL and EPRI hazard results—and the concomitant need for an improved methodology going beyond the late-1980s state-of-the-art—led directly to the formation of the SSHAC to perform this project. However, although the Committee studied both the LLNL and EPRI projects carefully to obtain methodological insights (both positive and negative), it did not undertake a forensic-type review to identify past "errors." Rather, it attempted to draw more broadly upon the entire body of PSHA literature and experience, including of course the LLNL and EPRI projects along with many others, to formulate the guidance herein.

In the course of our review, we concluded that many of the major potential pitfalls in executing a successful PSHA are procedural rather than technical in character. One of the most difficult challenges for the PSHA analyst is properly representing the wide diversity of expert judgments about the technical issues in PSHA in an acceptable analytical result, including addressing the large uncertainties. This conclusion, in turn, explains our heavy emphasis on *procedural* guidance.

This also explains why we believe that *how a PSHA is structured* is as critical to its success as the technical aspects—perhaps more critical because the procedural pitfalls can sometimes be harder to avoid and harder to uncover in an independent review than the pitfalls in the technical aspects. Finally, this also explains why *one of the key audiences for this report is the project sponsor*, who needs to understand the procedural/structural aspects in order to initiate and support the desired PSHA project appropriately.

This Executive Summary will conclude with a brief overview of what the SSHAC believes are its most important findings, conclusions, and recommendations in the procedural area. Because we recognize that several very important pieces of technical guidance concerning the earth-sciences aspects of PSHA will not be discussed in this Executive Summary, the SSHAC requests that readers turn to the full report to review the technical guidance. The key procedural points follow:

1) SSHAC identifies and describes several different *roles for experts* based on its conclusion that confusion about the various roles is a common source of difficulty in executing the aspect of PSHA involving the use of experts. The roles for which SSHAC provides the most extensive guidance include the expert as *proponent* of a specific technical position, as an *evaluator* of the various positions in the technical community, and as a *technical integrator* (see the next paragraph).

2) SSHAC identifies four different types of consensus, and then concludes that one key source of difficulty is failure to recognize that 1) there is not likely to be "consensus" (as the word is commonly understood) among the various experts and 2) no single interpretation concerning a complex earth-sciences issue is the "correct" one. Rather, SSHAC believes that the following should be sought in a properly executed PSHA project for a given difficult technical issue: (1) a representation of the legitimate range of technically supportable interpretations among the entire informed technical community, and (2) the relative importance or credibility that should be given to the differing hypotheses across that range. As SSHAC has framed the methodology, this information is what the PSHA practitioner is charged to seek out, and seeking it out and evaluating it is what SSHAC defines as *technical integration*.

3) SSHAC identifies a hierarchy of complexity for technical issues, consisting of four *levels* (representing increasing levels of participation by technical experts in the development of the desired results), and then concentrates much of its guidance on the most complex level (level 4) in which a panel of experts is formally constituted and the panel's interpretations of the technical information relevant to the issues are formally elicited. To deal with such complex issues, SSHAC defines an entity that it calls the Technical Facilitator/Integrator (TFI), which is differentiated from a similar entity for dealing with issues at the other three less-complex levels, which SSHAC calls the Technical Integrator (TI). Much of SSHAC's procedural guidance involves how the TI and TFI functions should be structured and implemented. (Both the TI and TFI are envisioned as roles that may be filled by one person or, in the TFI case, perhaps by a small team).

4) The role of *technical integration* is common to the TI and TFI roles. What is special about the TFI role, in SSHAC's formulation, is the *facilitation* aspect, when an issue is judged to be complex enough that the views of a panel of several experts must be elicited. SSHAC's guidance dwells on that aspect extensively, in part because SSHAC believes that this is where some of the most difficult procedural pitfalls are encountered. In fact, the main report identifies a number of problems that have arisen in past PSHAs and discusses how the TFI function explicitly overcomes each of them.

5) For most technical issues that arise in a typical PSHA, the issue's complexity does not warrant a panel of experts and hence the establishment of a TFI role. Technical integration for these issues can be accomplished—indeed, is usually best accomplished—by a TI. In fact, SSHAC has structured its recommended methodology so that even the most complex issues *can* be dealt with using the less expensive TI mode, although with some sacrifice in the confidence obtained in the results on both the technical and the procedural sides.

6) One special element of the TFI process is SSHAC's guidance on sequentially using the panel of experts in different roles. Heavy emphasis is placed on assuring constructive give-and-take interactions among the panelists throughout the process. Each expert is first asked, based on his/her own knowledge (yet cognizant of the views of others as explored through the information-exchange process), to act as an *evaluator*; that is, to evaluate the range of technically legitimate viewpoints concerning the issue at hand. Then, each expert is asked to play the role of *technical integrator*, providing advice to the TFI on the appropriate representation of the composite position of the community as a whole.

Contrasting the classical role of experts on a panel acting as individuals and providing inputs to a separate aggregation process, the TFI approach views the panel as a team, with the TFI as the team leader, working together to arrive at (i) a composite representation of the knowledge of the group, and then (ii) a composite representation of the knowledge of the technical community at large. (Neither of these representations necessarily reflects panel consensus—they may or may not, and their validity does *not* depend on whether a panel consensus is reached.)

The SSHAC guidance to the TFI emphasizes that a variety of techniques are available for achieving this composite representation. SSHAC recommends a blending of behavioral or judgmental methods with mathematical methods, and in the body of the report several techniques along these lines are described in detail. A key objective for the TFI is to develop an aggregate result th[illegible]iorsed by the expert panel both technically and in terms of the process used.

7) The TFI's integrator role should be viewed not as that of a "super-expert" who has the final say on the weighting of the relative merits of either specific technical interpretations or the various experts' interpretations of them; rather, the TFI role should be seen as charged with characterizing both the commonality and the diversity in a set of panel estimates, each representing a weighted combination of different expert positions. SSHAC thus sees the TFI as performing an integration assisted by a group of experts who provide integration advice.

8) Thus, the TFI as facilitator structures interaction among the experts to create conditions under which the TFI's job as integrator will be simplified (e.g., either a consensus representation is formed or it is appropriate to weight equally the experts' evaluations of the knowledge of the technical community at large). In the rare case in which such simple integration is not appropriate, additional guidance is provided. In the main report, guidance is presented on two possible approaches involving (i) explicit quantitative but unequal weights (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 fixed numerical weights to be artificial, and when it is appropriate to represent the community's overall distribution in a less rigid way.

9) The SSHAC guidance gives special emphasis to the importance of an independent peer review. We distinguish between a participatory peer review and a late-stage peer review, and we also distinguish between a peer review of the process aspects and of the technical aspects for the more complex issues. We strongly recommend a participatory peer review, especially for the process aspects for the more complex issues. This paper details the pitfalls of an inadequate peer review.

ACKNOWLEDGMENTS

This project has been jointly supported by the U. S. Department of Energy (DOE), the U. S. Nuclear Regulatory Commission (NRC), and the Electric Power Research Institute (EPRI). The project managers for the three sponsors, whose participation went far beyond managerial support to include technical input as well, and without whom the project could not have been successfully completed, were Jeffrey Kimball and Ann Bieniawski (DOE), Andrew J. Murphy and Ernst G. Zurflueh (NRC), and J. Carl Stepp and John Schneider (EPRI). They were ably supported administratively by Michael P. Bohn of Sandia National Laboratories (DOE) and Jean B. Savy of Lawrence Livermore National Laboratory (NRC), both of whom made technical contributions as well.

Although the intellectual responsibility for the report rests solely with the seven SSHAC members, the task could not have been accomplished without the diligent and very capable contributions of a large number of others who worked in a collaborative mode under the SSHAC Committee's overall direction. These other experts, some of whom ghost-wrote entire subsections or appendices of the report, were:

Don L. Bernreuter, Lawrence Livermore National Laboratory

Michael P. Bohn, Sandia National Laboratories (DOE Project Manager; lead author for the Glossary)

Auguste C. Boissonnade, Lawrence Livermore National Laboratory

Martin W. McCann, Jack Benjamin & Associates Inc. (lead author for Chapter 7)

Robin K. McGuire, Risk Engineering Inc. (author of Appendix G)

Richard W. Mensing, Logicon-RDA (lead author for Chapter 6; author of Appendix D; contributor to Chapter 4 and Appendix A)

Jean B. Savy, Lawrence Livermore National Laboratory (NRC Project Manager)

Gabriel R. Toro, Risk Engineering Inc. (major contributor to Chapter 5; author of Appendices F and I; contributor to Chapter 4 and Appendices A and B)

The SSHAC project organized four different workshops that were attended by a large number of experts representing a variety of disciplines. The participants, who are all listed in the workshop descriptions (see Appendices A, B, C, and H), deserve our thanks for contributing so significantly to the project.

We would also like to thank Norman Abrahamson for contributing an excellent piece of guidance that we have taken the liberty to incorporate herein as Appendix E.

The National Academy of Sciences/National Research Council organized a special "Panel on Seismic Hazard Evaluation" under its Committee on Seismology with the charter to review our report. This review was supported by the U.S. Nuclear Regulatory Commission. The Panel's review comments on our draft report of November 11, 1994 were especially helpful in focusing the SSHAC on key issues that needed extra attention. Besides the NAS/NRC Panel's review, we had the benefit of informal review comments on the November draft from about a dozen other specialists and organizations for which we are very grateful. The comments of the NAS/NRC Review Panel, which were published separately by the National Academy Press, are included in this report as an Appendix to Volume 1.

Finally, the logistical work of pulling the report together, based on input from many different authors typing on many different word processors, was accomplished in an outstanding manner by Rosa I. Yamamoto of LLNL, whose skill and dedication the SSHAC gratefully acknowledges.

1. INTRODUCTION

1.1 The SSHAC

In order to provide technical guidance on the subject of a *methodology for Probabilistic Seismic Hazard Analysis* (PSHA), a "Senior Seismic Hazard Analysis Committee" (SSHAC) was formed in early 1993 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).[1] The SSHAC has carried out this project as a working committee, and its members, the seven authors of this report, are jointly responsible for the report's contents.

To support the committee's work, a large number of experts on various technical subjects have been working under the committee's direction on specific topics integral to the effort.[2] These experts are listed in the Acknowledgments section.

The specific *objective* of this project, which will be discussed in more detail below, is to provide *methodological guidance on how to perform a PSHA*. Both technical guidance and procedural guidance are provided, with a strong emphasis on the latter. Why such guidance is necessary is discussed below.

1.2 Background

PSHA is an analytical 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. The results of such an analysis are expressed as estimated probabilities per unit time or estimated frequencies (such as expected number of events per year).

Unfortunately, this objective of estimating earthquake-caused ground-motion frequencies 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 gaps in our understanding of the mechanisms that cause earthquakes, and of the processes that govern how an earthquake's energy propagates from its origin beneath the earth to various points near and far on the surface. These gaps in understanding mean that, when a PSHA 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 setting—such as in coastal California—the scientific data base (specific fault locations, orientations, slip rates, energy dissipation mechanisms, etc.) is still limited. In fact, major new insights arise whenever there is another large earthquake. 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 data base is even less able to support scientific understanding of what might cause earthquakes, because, despite significant recent advances in knowledge, not even the sources or mechanisms of earthquakes are well understood.

This lack of understanding has operational implications for the analyst charged with performing a PSHA. Specifically, there often exist wide differences of legitimate scientific opinion on many of the key inputs into a PSHA. 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. These differences of interpretation translate into important uncertainties in the PSHA's numerical results, and make these results less useful for many potential applications of PSHA. Operationally, a PSHA analyst is faced with how to use these different interpretations properly,

[1] Some members of the SSHAC have been supported by NRC funds directly, some members by NRC through contracts with Lawrence Livermore National Laboratory, and other members by DOE funds through contracts with Sandia National Laboratories.

[2] Contractually, these experts have been supported variously by NRC, DOE, and EPRI.

incorporating the diversity of expert judgments into an analytical result that appropriately captures the current state-of-knowledge including its uncertainties.

For the Committee, addressing this situation has been a challenge. In developing guidance for performing a PSHA we have had to face two different (although related) tasks:

(i) developing technical guidance, drawn from the earth sciences, concerning the scientific issues involved in performing a PSHA; and

(ii) developing procedural 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.

Because this situation was recognized from the start, the three sponsors of this project (DOE, NRC, and EPRI) established a broad-based committee, supported by a broad-based group of other scientists and engineers, with expertise not only in all of the major earth-science disciplines but also in the other key areas. The resulting guidance in this report is comprised of a mix of both earth-sciences guidance and procedural guidance. If a successful PSHA project is to be carried out, there is heavy emphasis on the importance of the latter. This is because it is often more difficult to execute the procedural aspects properly (including how expert interpretations are used) than the technical aspects, and because there exists far less procedural guidance in the literature.

Note that our guidance is not intended to be "the only" or "the standard" methodology for PSHA to the exclusion of other approaches; there are other valid ways to perform a PSHA study. Likewise, our formulation should not be viewed as an attempt to "standardize" PSHA in the sense of freezing the science and technology that underlies a competent PSHA, thereby stifling innovation. Rather, our guidance is intended to provide not only up-to-date technical guidance for the analyst, but also procedural guidance that we believe is crucial to the successful execution of a PSHA project today and for the next several years.

Because our sponsors are interested in applications for siting and regulation of nuclear power plants and other nuclear facilities, we have considered their interests throughout the project. However, as discussed below, we believe that the resulting methodology should be useful, in whole or in part depending on the issues, for other PSHA applications as well. In Chapter 3, we distinguish among four different levels of study in a PSHA.

1.3 History

The discipline of PSHA has evolved over several decades. Early empirical statistical methods (for example, Milne and Davenport 1969) have been largely replaced by the analytical/numerical models initiated by Cornell (Cornell 1968), and further refined by many researchers in subsequent years.

Many site-specific and regional mapping applications have been made around the world. The need to consider the uncertainty in parameters and models was recognized early on. The SSHAC members have drawn on their extensive experience in such studies, both large and small in terms of the resources expended.

The systematic, explicit incorporation of the diversity of expert interpretations on a regional basis was pioneered by a Lawrence Livermore National Laboratory study (Bernreuter et al. 1981) that examined several U.S. sites with operating nuclear power plants. The methods therein were later applied to several DOE sites. The expert interpretation aspect of PSHA was then addressed more formally in two major PSHA projects in the mid-1980s, both breaking major new ground on several fronts. Today they remain significant landmarks. The "Livermore" and "EPRI" studies included a PSHA on a broad regional basis covering the entire central and eastern United States:

(i) The "Livermore" study (Bernreuter et al. 1989) was sponsored by the U.S. Nuclear Regulatory Commission and executed by a

team at the Lawrence Livermore National Laboratory. Its objective was to develop seismic hazard curves for the 69 sites in the eastern U.S. (east of the Rocky Mountains) at which nuclear power plants were then operating. It accomplished this by performing a broad regional study, and then extracting the 69 site-specific seismic hazards from the regional PSHA information. It called upon a large number of experts, whose interpretations of the earth-sciences information were individually elicited using a formal expert-elicitation process and then combined together by the LLNL team to produce the PSHA results. Separate elicitation processes were used for the seismic-source characterization and the ground-motion aspects.

(ii) The "EPRI" study (EPRI 1989) was sponsored by the Electric Power Research Institute. Its objective was to develop seismic hazard curves for most of the sites in the eastern U.S. (east of the Rocky Mountains) at which nuclear-power plants were then operating, although a few of the sites covered by the Livermore study were not covered. Like the Livermore study, the EPRI analysis was a broad regional study, which then calculated the site-specific seismic hazards from the regional PSHA information. For the seismic-source part of the analysis, the EPRI study utilized a large number of geoscientists who were grouped into several different seismic-source teams whose interpretations of the earth-sciences information were elicited, team-by-team, using a formal expert-elicitation process. The ground-motion part of the analysis was performed using a weighted combination of models developed by the analyst team. The seismic-source and ground-motion information was then combined together by the EPRI group to produce the PSHA results.

Although the Livermore and EPRI studies were similar in many ways, both technically and procedurally, they also differed significantly in a few areas. As mentioned, both broke important new ground, primarily with respect to the implementation processes used but in many substantive technical areas as well, and today both are key landmarks in the history of PSHA. However, for our historical purposes here, the most important fact about these two studies is that the *Livermore and EPRI mean seismic hazard curves for most sites in the eastern U. S. differed significantly*. This led, for several years after both studies were published, to considerable consternation and several efforts to understand what might underlie the differences. The reason for the consternation was that the differences between the Livermore and EPRI results had important implications for policy in a number of areas. However, no completely satisfactory explanation for these differences emerged, despite several important studies (both Livermore-EPRI comparison studies and new PSHA studies at various sites) that cast useful light on various technical and procedural issues.

Ultimately, although there was a strong feeling in the PSHA community that procedural issues rather than technical earth-sciences issues per se were an important reason for the differences, the inability to understand all of the differences between the Livermore and EPRI hazard results led directly to the formation of the SSHAC to perform this project. Originally, some of the sponsors and participants proposed that one key study objective should be to "resolve" the differences between the Livermore and EPRI studies. However, the Committee quickly realized that the new project would be most useful if it were forward-looking rather than backward-looking—specifically, if it could pull together what is known about PSHA in order to recommend an improved methodology, rather than specifically attempting to figure out which of the two studies was "correct," or which specific problems with either study were most important in affecting that study's specific results.

Therefore, although the Committee has carefully studied both the LLNL and EPRI studies (along with other past PSHAs) to obtain methodological insights, both positive and negative, we did not undertake a forensic-type examination to identify past "errors" or their implications. More broadly speaking, the Committee has attempted to draw

upon the entire body of PSHA literature and experience, which is of course much more extensive than the LLNL and EPRI projects, as important as they have been.

The above discussion is a natural introduction to a presentation of the SSHAC project charter and objectives, which are discussed next.

1.4 Objective of the Project

At the inception of the project, the three sponsors (DOE, NRC, and EPRI) provided an "objective" for the SSHAC effort, as follows:

> "The objective is to develop implementation guidelines, including a recommended methodology, suitable for the performance of PSHA for seismic regulation of nuclear power plants and other critical facilities."

Operationally, the SSHAC has taken its charter to be:

> To describe an up-to-date PSHA methodology, including guidelines and recommendations, that can guide the analyst both technically and procedurally.

Because PSHA results can be so important for both engineering design and public-policy decision-making, a goal of this project is that the PSHA methodology will ensure the stability of the numerical results for a reasonable period of time (five to ten years) or until significant new technical information presents itself.

This goal will be achieved by (i) ensuring that the assessment is based on unbiased interpretations of available data and information, and (ii) explicitly identifying and evaluating the uncertainties in the PSHA inputs, including both data and model inputs, and incorporating them in the composite measure of the uncertainty in the results.

1.5 Audience for the Report

This report has been written with four different audiences in mind:

- analysts who will implement the PSHA methodology (and for whom the specific guidance has been written);
- earth scientists whose expertise will be drawn upon by the analysts, and who will require an understanding of the entire PSHA process in order to participate most effectively in a PSHA project;
- technical reviewers who will be called upon to review a PSHA study, either to advise a study's sponsors of its validity or to provide support for a regulatory decision;
- PSHA project sponsors, meaning decision-makers in entities such as private firms or government agencies, who have a need for PSHA information and who are in a position to sponsor a PSHA study. Such sponsorship includes both financial and institutional sponsorship, and we have both in mind.

The first three audiences should be interested not only in the broad guidance but also in the specific technical details. The fourth audience, although perhaps not as interested in the detailed guidance about how to determine seismic sources or ground-motion attenuation, should be interested in how the committee envisions that a PSHA project must be put together, how the process is expected to work for different levels of effort, how to avoid the known pitfalls observed in past studies, and how to set realistic expectations as to the validity of the results.

1.6 Conditions and Limitations on the Guidance

In order to bound the scope, the Committee and its sponsors decided on several conditions and limitations that are important for any reader to understand. The principal ones are:

- Types of applications: In the past, probabilistic seismic-hazard analyses have been used in at least four quite different ways. These different types of applications, all of which are contemplated in the SSHAC guidance, are:

(i) to understand the seismic *hazard* at a specific site in order to establish site-specific safety regulations;

(ii) to guide the establishment of *specific criteria for the seismic design, evaluation, and/or retrofit* of a facility;

(iii) to provide the *hazard input* to a comprehensive probabilistic seismic-risk assessment for a facility, either existing or in the design stage; and

(iv) to support development of regional *seismic-hazard maps* used in broad applications such as building codes.

Of course, depending on the application, different levels of effort may be indicated.

- Breadth of application: Although the emphasis in the formal statement of objective is on "seismic regulation of nuclear power plants and other critical facilities," the SSHAC methodology can clearly be used more broadly. In fact, SSHAC has contemplated various broader applications from the start. Any attempt to apply the methodology to regions, sites, or facilities that are significantly different from "nuclear power plants and other critical facilities" should evaluate the methodology's applicability on a case-by-case basis, because SSHAC's preferred approach may not always apply directly to other facilities. However, the issues that are raised and discussed here, especially but not exclusively the procedural issues, apply in varying degrees to any PSHA project, and should be at least considered by sponsors and analysts before undertaking almost any PSHA. (See Section 5.1 for a description of four study levels that SSHAC has identified.)

- Site-specific vs. regional applications: PSHA can be applied not only to specific sites but also to broad regions. Both applications are contemplated in the SSHAC guidance.

- East and west: The SSHAC methodology is intended for application in both the eastern U.S. and western U.S. (more broadly, in both low-seismicity and high-seismicity regions of the country). Even though the specifics of implementation differ in detail in these two very different regions, the procedural aspects should be similar.

- Probabilistic vs. deterministic seismic-hazard analysis This project addresses the methodology for performing probabilistic seismic-hazard analysis. A number of non-probabilistic approaches to understanding seismic hazards are widely practiced and have considerable value in some applications. These non-probabilistic methods are usually called "deterministic" methods. Although it has been tempting to develop information about the similarities and differences between PSHA and some of the most widely-used deterministic methods, the Committee has explicitly not done so at the direction of the sponsors.

1.7 Philosophy of the Project

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 conducting the analysis. We believe that a well-defined philosophy establishes the foundation for developing the rules and guidance that are provided here.

We have identified five elements of our philosophy that merit discussion in this introductory chapter, and these will be discussed next:

- The level of prescription
- Stability
- The use of "experts" and the meaning of "consensus"
- Transparency
- Performing a PSHA using different levels of effort

The level of prescription: The SSHAC has attempted to provide explicit guidance, and,

where feasible, prescribed approaches for performing a PSHA. Sometimes, our guidance will require a certain methodological approach, while at other places we will recommend, or perhaps suggest, how the analyst should proceed, or in some cases we may merely allow a particular approach. This hierarchy in the guidance ("require," "recommend," "suggest," and "allow" and their opposites) is intended to be as explicit as the Committee believes can be supported by the information available. We realize that, because we have developed this guidance primarily with nuclear-power-plant and other nuclear-facility applications in mind, the hierarchical structure may not apply directly in other applications.

Sometimes, there will be several alternative approaches to a particular element in the PSHA methodology. Where the committee judges that these are equivalent, the guidance has attempted to identify one approach among the alternatives and to require or recommend it. This is not intended to denigrate the validity of the alternative approaches, but by narrowing the options we do intend to provide for a degree of uniformity, which will enhance the technical community's ability both to compare the work of different PSHA practitioners and to review it more easily. The SSHAC wishes to avoid the implication that using PSHA approaches *other than the one recommended here* would be invalid. Nevertheless, the Committee believes that the *issues* raised and dealt with herein should be considered by every PSHA sponsor and analysis team.

Stability: By following SSHAC's guidance, a PSHA practitioner should be able to provide reasonable assurance that the numerical hazard results will be stable for a reasonable period of time following the completion of the PSHA (unless significant new seismic information, which could occur at any time, calls for a major revisitation). This stability is important to both the technical and the policy-making community, and is achieved by ensuring the integrity of the process, involving two crucial elements: (i) completeness of the earth-science information used in the analysis, and (ii) a thorough evaluation of the uncertainties in professional interpretations of that information. As discussed below, SSHAC believes that without such an evaluation, the user—whether the technical user or the policy user—is not adequately served.

The use of "experts" and the meaning of "consensus": In writing the guidance in Chapter 3, the Committee has given careful attention to the role of "experts" in the PSHA process. As Chapter 3 describes, 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 a subject. A given PSHA project will utilize various types of experts in various different roles that SSHAC believes must be kept clearly separate, even if the same individual often changes roles in different phases of the same PSHA project.

In Chapter 3, we also dwell at length on the issue of "consensus," identifying four different types of consensus and describing how each plays its specific role in a PSHA project. The SSHAC believes it important to emph[illegible]here that, given the existence of differing inte[illegible]tions of the technical information input in [illegible]SHA, there is not likely to be a "consensus" [illegible]s that word is commonly understood by lay readers) among the various experts that a single interpretation of the earth-sciences information is the "correct" one. This is the case for both seismic sources and seismic ground-motion attenuation.

Rather, the following should be sought in a properly-executed PSHA project for a given technical issue: (i) a representation of the legitimate range of technically supportable interpretations among the entire informed technical community, and (ii) the relative importance or credibility—read "weight" even if not a numerical weight—that should be assigned to the differing hypotheses across that range. As SSHAC has framed the methodology (see the detailed discussion in Chapter 3), the PSHA practitioner is charged to seek out this information, whether by "sampling" a sub-set of

the community of experts or, if financially restricted, by drawing upon only the literature and his/her own judgment.

Transparency: The results of a PSHA serve a range of users with different needs. To assure that all of these needs are met, the information that is generated as part of the PSHA must be documented in a transparent way. Transparency of the PSHA, including not only the input data and models used but also the process employed and the results obtained, satisfies the needs of (i) the earth scientist who is interested in understanding the scientific issues, (ii) the engineer who must understand how the ground motion predicted at a given site has been derived, and has been related to the magnitudes and distances of the contributing earthquakes, (iii) the technical reviewer who must be satisfied with the completeness and scientific integrity of the earth-science interpretations and of the PSHA process, and (iv) the decision-maker concerned with the stability and integrity of the results as a whole.

Documenting the PSHA, including both the methodology and the results, in a transparent way allows all of these users to see how the constituent parts of the assessment fit together. This will reduce the apparent level of complexity generally associated with these assessments.

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. This is appropriate as a starting point, because some applications are so important that the sponsors can afford to devote upwards of a million dollars or more to the PSHA and the science upon which it is based.

However, the committee recognizes that some sponsors may not be able to devote such vast resources to a PSHA project, or may not even require a PSHA assessment of very large potential ground motions that would be associated with very rare events. In these cases, a scaled-down approach may be appropriate. To assist such sponsors, we have attempted to differentiate those elements of a PSHA that are essential to its success—that must be incorporated—from those elements where it may be feasible to compromise, accepting more uncertainty (and concomitantly, less confidence in the results) as the result of a smaller project scope. In any case, the basic constituent elements of a PSHA are the same in all applications, even if the process is different.

The committee emphasizes, however, that wherever we have indicated that certain types of compromises are acceptable, we nevertheless insist that there be no compromise in the rigor with which the PSHA is undertaken. Only the size of the residual uncertainties (which in any event will be large, even for the most expensive PSHAs) may be compromised; and even here, our committee requires that a careful characterization of both the source and size of the uncertainties be part of any PSHA.

1.8 Uncertainties in PSHA

In the introductory section 1.2 above, we mentioned that the results of PSHA, as defined for the purposes of this report, are expressed in terms of likelihoods—estimated probabilities in a given time period or estimated frequencies—that earthquakes producing various sizes of ground motion will occur at a given site or in a given region.

The SSHAC has adopted a probabilistic formulation for dealing with seismic hazards that embeds uncertainties in the core of the methodology. This has forced the Committee to try to deal directly with *all* of the various uncertainties that characterize our current state-of-knowledge.

Although the optimism of science in general leads some to believe that nearly all of the "uncertainties" in PSHA that we will deal with in this report are ultimately amenable to reduction, we recognize that for practical purposes many of them cannot be thought of or dealt with in this way. We define two different classes of uncertainties:

- Those that we will call *epistemic* are lack-of-knowledge uncertainties arising because our

scientific understanding is imperfect *for the present*, but are of a character that in principle are reducible through further research and gathering of more and better earthquake data.

- Those that we will call aleatory—"random" in character—are uncertainties that for all practical purposes cannot be known in detail or cannot be reduced (although they are susceptible to analysis concerning their origin, their magnitude, and their role in PSHA).

In the seismic case, it is helpful to consider a mental model in which some thousands of years of an earthquake catalog and site-specific ground-motion recordings were made available. In this case, the former epistemic uncertainty would be reduced to near zero, whereas the forecast of the maximum ground-motion at the site in the next year would remain subject to aleatory uncertainty.

The division between the two different types of uncertainty, epistemic and aleatory, is somewhat arbitrary, especially at the border between the two. This is because, conceptually, some of the processes and parameters whose uncertainties we will characterize here as aleatory ("random") may be partially reducible through more elaborate models and/or further study. However, for our purposes here, we will relinquish such a hope or expectation, and will treat some of the uncertainties in various processes and parameters as unknowably aleatory.

The conceptual difference between epistemic and aleatory uncertainty is an important element of SSHAC's approach to PSHA. In the chapters of this report that follow, we will provide methodological guidance that incorporates uncertainty analysis at the core of the approach, and that therefore cannot be implemented without an understanding of how uncertainties of both types are dealt with. Especially in light of the fact that our knowledge of earthquake phenomena is still so incomplete, which necessitates that the PSHA analyst must deal with diverse expert interpretations of the insufficient information that does exist, we wish to reinforce here, in the introductory chapter, that a PSHA that follows the rest of our guidance but that does not deal appropriately with both the epistemic and the aleatory uncertainties must be considered inadequate.[3]

This exhortation does not imply that every PSHA study must undertake a highly refined uncertainty analysis in order to be valid. Depending on the application, the uncertainty treatment may be adequate while relying largely on experience in similar situations and the judgments of the analysts for its support. However, the SSHAC approach emphasizes that unless the analysis team deals with the major uncertainties instead of "ducking" them, the PSHA results will not be complete, and the full description of the problem faced will not have been communicated to the users of the results.

1.9 Introductory Comments on a Few Other Issues

Regulatory applications: Another SSHAC objective is to develop a methodology that satisfies, when necessary, NRC requirements for nuclear-power-plant siting, including the ability to be reviewed by the NRC staff and adjudicated in an administrative hearing. Meeting this objective will allow the methodology to be used in other similar regulatory or quasi-regulatory settings, including those contemplated by DOE for its reactor and non-reactor facilities.

However, SSHAC has not given significant attention to the specific ways in which PSHA results have been used in the past, or may be used in the future, in the regulatory arena. Each sponsor of a particular PSHA study must work together with the project team to direct the project's efforts at those applications—regulatory or otherwise—that are that study's intended use. This includes such crucial issues as the scope of the project, any special documentation requirements, the relative emphasis on mean or median hazard results as the more important (if appropriate), and so on. In particular, it is not known to us whether the results of a given PSHA

[3] In certain applications, the objective is simply the mean annual hazard (that is, the expected value with respect to epistemic uncertainty). In this case, the result is not sensitive to the distinction between the two uncertainty types, but both must still be captured to obtain the correct value.

performed using the SSHAC guidance will be useful for nuclear-power-plant regulatory purposes.

Seismic hazard expressed in terms of ground motion: The SSHAC has thought about seismic hazards principally in terms of the ground motion that would occur at a given site. This ground motion can be expressed in many different ways (response spectral ordinates, peak acceleration, etc.) that are discussed in detail in the body of this report. Generally, the results of a PSHA are expressed in terms of the likelihood in one year that a certain level of ground motion may be exceeded, usually called the "probability of exceedance."

Seismic hazard at a rock outcrop vs. local site effects: Local site effects must be considered in any site-specific application to a facility, and some guidance on them is provided here. However, the SSHAC sponsors decided early in the effort that the principal emphasis should be on recommending a methodology to obtain the seismic hazard (ground motion) at a hypothetical (or actual) hard-rock outcrop at a given site.

Uncertainty: As discussed above, the Committee believes that a PSHA that does not deal with the various uncertainties properly is not useful for nearly all the contemplated applications. Therefore, the Committee has given special attention to guidance on this subject, which has turned out to be one of the major issues in the project.

1.10 Criterion for Success of the SSHAC Project

With PSHA, even cookbook-type methodological guidance allows flexibility in implementation. Of course, such flexibility means that different teams inevitably will interpret and apply the methodology differently. Early in the project, the Committee agreed that a key criterion for success of the SSHAC project would be that the recommended PSHA methodology, when independently applied by different groups, would yield "comparable" results, defined as results whose overlap is within the broad uncertainty bands that inevitably characterize PSHA results. For this to be true, we believe (as discussed above) that the uncertainties in the methodology must be confronted and dealt with head-on. No PSHA analyst should attenpt less, and no PSHA sponsor should accept less.

Furthermore, if the results of two such studies turn out not to be "comparable," following the guidance herein will provide a framework within which the differences can be identified and debated in a structured manner.

1.11 Road Map to the Report

The report is organized into several Chapters. Following this Introduction, Chapter 2 contains an overview of the PSHA methodology. Next, Chapter 3 provides the crucial guidance on structuring a PSHA project, including how experts are used and how the peer review process should be structured. The next two chapters present the methodology for characterizing both seismic sources (Chapter 4) and ground-motion attenuation (Chapter 5). This is followed by a discussion of the methodology for producing the PSHA results (Chapter 6) and guidance on obtaining insights from the results and on documenting the project (Chapter 7). A glossary and comprehensive list of references complete the report.

Material too detailed or outside the scope of the main report can be found in the Appendices.

2. OVERVIEW OF THE PSHA METHODOLOGY

2.1 The Basic Probabilistic Model

The model of the randomness (or aleatory uncertainty) of the behavior of the earth that underlies virtually all probabilistic hazard analysis is by now very familiar to scientists and engineers working in the field. The SSHAC endorses this model for all but certain uncommon cases where the information available may permit or require specific deviations. As with any effective representation of nature, the model below represents a compromise between complexity, availability of information, and sensitivity of the results.

The objective is to estimate the mean frequency per unit time or, alternatively, the probability in a given future time period that a specified level of some ground motion parameter will be exceeded at a site of interest. For example, the result might be the annual probability that the 1 hertz spectral acceleration at the site exceeds 0.3g. In general, one will seek this probability for a range of levels, i.e., as a function of the ground motion parameter value. Also, a suite of different ground motion descriptors will be studied, e.g., spectral accelerations for several different oscillator frequencies. We will focus here, however, on the simplest case.

The components of the aleatory model are those that (i) characterize the seismicity in the vicinity of the site, and (ii) represent the ground motion prediction of the effect at the site should an earthquake of given size (magnitude) occur a given distance from the site. These two general subjects will be dealt with in great detail in Chapters 4 and 5. For our purposes here we presume that the seismicity is represented by a set of **s** independent "sources" each with spatially homogeneous seismicity and that the ground motion prediction is characterized by a function g(m,r) that yields the mean value of the (natural) log of the ground motion parameter, lnA, given the magnitude, m, and the distance, r, of the event. To a first approximation, this function increases linearly in magnitude and decays logarithmically with distance. Further, the variability (event to event and site to site) observed in ground motion data is represented by a Gaussian distribution on lnA with standard deviation σ. σ is often assumed to depend on m.

For any given source of seismicity, the model assumes that earthquake events of "engineering interest," i.e., those above a magnitude threshold such as 5.0, occur with a mean annual rate ν. Further we assume that these events occur at relative frequencies, $f_M(m)$. This probability density function has a corresponding CCDF (complementary cumulative distribution function), $G_M(m)$, which is the fraction of events with magnitude m or greater. The common assumption is that the form of $f_M(m)$ is exponential:

$$f_M(m) \propto e^{-\beta m} \text{ for } m_o \le m \le m_u \qquad (2.1)$$

in which m_o is the lower threshold and m_u is the upper bound magnitude, the largest magnitude that this particular source is capable of producing, while β is the parameter that determines the relative frequency of larger to smaller events. This parameter is, within a constant, the traditional b value of the familiar Gutenberg-Richter relationships: $\beta = \ln(10)b \equiv 2.3b$. In certain applications this exponential magnitude frequency distribution may be supplemented by a "characteristic magnitude" distribution; this implies superposing a "spike" or narrow rectangular bar of frequency density at or about the value of this characteristic size; this size is usually m_u.

Finally the assumption that the source is spatially homogeneous implies that any point within that source is equally likely to be the hypocenter of the event. From this knowledge and the geometry of the source relative to the site, one can deduce a function $f_R(r)$ that describes the relative frequency of different site-to-earthquake distances. In plan view the geometry of the source may be a line (fault) or arbitrarily shaped area. The rupture may be considered to be simply a point or a line (a fault in plan view) of a length that depends on the magnitude. In the latter case the distribution on

distance (defined typically as the closest distance between site and fault) would also depend on magnitude, $f_R(r \mid m)$.

Then, of all those events of magnitude m at distance r from the site, the ground motion assumptions above imply that the fraction that causes ground motion greater than or equal to level **a** is

$$\Phi^c\left(\frac{\ln a - g(m,r)}{\sigma}\right)$$

in which $\Phi^c(\cdot)$ is the CCDF of the standard unit normal. Therefore the fraction of all events on the source that equal or exceed **a** is

$$\int\int \Phi^c\left(\frac{\ln a - g(m,r)}{\sigma}\right) f_R(r|m) f_M(m)\, dr\, dm$$

The mean annual frequency of such events is simply this fraction times ν, the mean rate of all events.

Then, to consider all s sources, we need simply sum the mean rates from each source leading to the following expression for $\lambda(\mathbf{a})$ the mean annual rate of events with site ground motion level **a** or more:

$$\lambda(a) = \sum_{i=1}^{S} \nu \int\int \Phi^c\left(\frac{\ln a - g(m,r)}{\sigma}\right) f_R(r|m) f_M(m)\, dr\, dm \qquad (2.2)$$

in which the subscripts on all the factors within the sum (ν, f_R, f_M, and even possibly g(m, r) and σ) are deleted for simplicity. This is the basic equation of probabilistic seismic hazard analysis. It is the simple algorithm by which the many important pieces of the total puzzle are finally integrated.

Under the additional assumption that the events in every source follow independent Poissonnian processes, the mean rate $\lambda(\mathbf{a})$ can be used to compute the probability of exceedance in any time interval of length t:

$$P[A > \mathbf{a} \text{ in time } t] = 1 - e^{-\lambda(\mathbf{a})t} \qquad (2.3)$$

in which P[·] is read "the probability of the event that." Note that for the small probabilities of usual interest in PSHA problems the value of $\lambda(\mathbf{a})t$ is small relative to unity, in which case the probability in Eq. 2.3 is approximately equal to simply $\lambda(\mathbf{a})t$. In different words the annual probability is approximately equal to the mean annual rate. Therefore the two phrases are used virtually interchangeably in common PSHA and in this report.

In certain problems it may be important to recognize that some of the events, such as the "characteristic" events, are not Poissonian in their temporal stochastic behavior. In this case it is usually sufficient (Cornell and Winterstein 1988) to replace the mean rate ν of such events by the time interval average of what is called the hazard function $h(\tau)$ which is a function of the time elapsed since the last such event on the source. It is in this case that one must distinguish carefully between the probability (Eq. 2.3) and the mean rate (Eq. 2.2). The *probability* is the appropriate item to calculate and report.

Further, it should be recognized that the spatially homogeneous areal source model used in many applications is not a physical characteristic of the earth but a simplified mathematical representation of a field of seismogenic structures that the earth scientist believes can be approximated adequately for hazard estimation purposes by such a model. Its application in practice will be discussed in Chapter 4.

2.2 Primer on Uncertainties

2.2.1 Introduction

The purpose of this section is to introduce several concepts and the associated terminology essential to the framework adopted for PSHA. We discuss both the nature of physical models ("models of the world")—recognizing that they can be either deterministic or probabilistic depending on the application—and our knowledge and ability to model the phenomena (the "world") of interest. We then acknowledge that models themselves, as well as the parameters appearing in them, may be uncertain and we introduce probabilities to express these uncertainties. The uncertainties that are part of the model of the world, if any, are called *aleatory* uncertainties (other names are

"stochastic" or "random" uncertainties). Even under "perfect information," i.e., when the model has been validated and the numerical values of its parameters are known, these aleatory uncertainties are still present (for a given model).

The uncertainties that stem from our lack of knowledge concerning the validity of the models and the numerical values of their parameters are referred to as *epistemic* uncertainties (in the literature, they have been referred to as simply "uncertainties"). As information is collected, the epistemic uncertainties are reduced. We prefer to use the terms aleatory and epistemic because they have a unique interpretation; alternatives (e.g., "uncertainty" for "epistemic") have multiple meanings.

We also discuss in this section the concept of model uncertainty in more detail, as well as the display and communication of the various types of uncertainty.

2.2.2 Deterministic and Aleatory Models of the World

The "model of the world" is the mathematical model that is constructed for the physical situation of interest, such as the occurrence and impact on a system of a physical phenomenon. The "world" is defined as "the object about which the person is concerned" (Savage 1972). Occasionally, we will refer to the model of the world as simply the model, or the mathematical model. Constructing and solving such models is what most physical scientists and engineers do. There are two types of models of the world, deterministic and probabilistic. A simple example of a deterministic model is the function g(m,r) that yields the mean value of the logarithm of the ground motion parameter at a specified site given the magnitude, m, and the distance, r, of the earthquake (see Section 2.1).

Many important phenomena cannot be modeled by deterministic models. For example, the actual ground motion parameter A can not be predicted precisely. We then construct models of the world that include this uncertainty. A simple example is the normal distribution (Section 2.1), i.e.,

$$\Pr(A > a) = \Phi^{c}\left(\frac{\ln a - g(m,r)}{\sigma}\right) \qquad (2.4)$$

where $\Phi^{c}(\cdot)$ is the CCDF of the standard unit normal.

The interpretation of this probability is the following: if we consider very many earthquakes all at a distance r from the site and of magnitude m, the fraction of events leading to a ground motion parameter A (the "random variable" of this problem) exceeding a given value a will be very close to this probability.

The uncertainty described by the model of the world is sometimes referred to as "randomness," or "stochastic" uncertainty. Stochastic models of the world are also called *aleatory* models [aleatory: of or depending on chance, luck, or contingency (Webster's 1988)].

In addition to the two examples from the model for ground motion cited above, we recognize that the representation of seismicity by a number of "sources" each with a specified mean rate of seismicity along with eq. (2.1) for the magnitude distribution of an earthquake occurring in a given source is an aleatory model.

2.2.3 The Epistemic Model

There are two additional types of uncertainties associated with a (deterministic or aleatory) model of the world. The model itself (or, the hypotheses behind it) may involve approximations, so that its predictions deviate by a fixed but unknown amount from observed values of the predicted quantity. The second type is associated with uncertainties about the numerical values of the parameters of a given model, e.g., the parameter β of eq. (2.1).

The *epistemic* probability model represents our knowledge regarding the numerical values of the parameters and the systematic over- or under-predictions of the model [epistemic: of or having to do with knowledge (Webster's 1988), see also Paté-Cornell and Fischbeck (1992); in risk assessment, this probability distribution function (pdf) is also referred to as a "state-of-knowledge" pdf (Kaplan and Garrick 1981)].

The epistemic model for the deterministic model of the world that consists of seismic sources allows alternate boundaries for each source. Each alternate map is associated with an epistemic probability. The upper-bound magnitude m_u of each source is assigned its own epistemic probability distribution. Similarly, for the ground motion model, eq. (2.4), the function g(m,r) itself is uncertain (epistemic model uncertainty); multiple alternatives are often considered. The value of the standard deviation σ of the aleatory model is also uncertain (epistemic parameter uncertainty).

We note that the probability distributions that reflect epistemic uncertainty are the ones that are "updated" as empirical evidence is gathered. For a given model, the updating of the epistemic distributions of its parameters is done using Bayes' theorem, as shown in numerous references, e.g., Lee (1989), Benjamin and Cornell (1970), Winkler and Hays (1975), and Apostolakis (1990). It can be shown that, when the empirical evidence is very strong, these epistemic distributions become delta functions about the exact numerical values of the parameters. At this point, no epistemic uncertainty about the numerical values of the parameters exists and the only uncertainty in the problem is the aleatory uncertainty in the model of the world. The latter can never be removed (unless, of course, we happen to change the model of the world).

The types of uncertainty that we have presented are defined from what can be called the "probabilist's" point of view. We have associated epistemic uncertainties with models and their parameters only, while aleatory variables appear in the model of the world. Unfortunately, this clear distinction cannot be maintained in PSHA, because the engineering use of the term "parameter" is not consistent with this formulation. For example, the stochastic ground motion model (Electric Power Research Institute 1993) includes what are called these parameters, such as the stress drop, $\Delta\sigma$, that are assumed to have both kinds of uncertainty.

In the formulation that we have presented, quantities such as $\Delta\sigma$ would be called aleatory variables and they, as well as their assumed probability distributions, would be part of the (aleatory) model of the world. The moments (or other parameters) of these aleatory distributions would, in turn, be assigned epistemic probability distributions. Having made this distinction clear, we must follow common engineering practice and call these quantities "parameters" of the models. Thus, a parameter in an engineering model may have related to it both aleatory and epistemic probability distributions. We still, however, make a distinction between, on one hand, quantities that deal with model uncertainty (see, for example, the discussion on D below), and parameters, such as $\Delta\sigma$, on the other. The term "model uncertainty" is used for the former, while "parametric uncertainty" is reserved for the latter. As an example, Table 5-2 in Chapter 5 shows this classification in the context of ground-motion models.

We point out that, even though we have discussed probabilities appearing in the model of the world and the epistemic model, and we have given them different names, leading philosophers of science and uncertainty (e.g., De Finetti 1974; De Groot 1988) believe that, conceptually, there is only one kind of uncertainty; namely, that which stems from lack of knowledge. Aleatory and epistemic uncertainties are a convenient way to distinguish between uncertainties that cannot be reduced (for a given model) and uncertainties that can be reduced as new knowledge is acquired.[1]

2.2.4 More on Model Uncertainty

Consider the case of a single deterministic model of the world which calculates the quantity y_c. Furthermore, we know that there are significant model uncertainties associated with this prediction. One way to describe this situation is to introduce a parameter D into the model of the world which may be multiplicative or additive. For example, we may assert that the actual value

[1] The distinction between aleatory and epistemic uncertainty may at first appear inconsistent with the Baysian view of probability, but, in fact, it is entirely consistent with this view. Aleatory uncertainties may be thought of as frequencies of a set of exchangeable events or as frequency distributions of an exchangeable set of continuous random variables. If the frequencies or frequency distributions are uncertain, it makes perfect sense to assess probability distributions over the unknown frequencies or parameters of the unknown frequency distributions.

may be obtained as $y_a = Dy_c$. In this case, the parameter D is multiplicative (this is, for example, how the EPRI stochastic ground motion model treats part of model uncertainty). This parameter may be interpreted as the ratio of the true value over the predicted value (by the model). Note that we still may assign both aleatory and epistemic uncertainty distributions to the parameters of the deterministic model that produces y_c.

The quantity D is a deterministic parameter of the model of the world. However, its numerical value may be uncertain, therefore, in the epistemic model we represent this uncertainty by a probability density function (pdf) g(d).

Let us consider again the implications of this formulation in terms of a thought experiment. After even one observation, the exact value of D will be known (since all epistemic uncertainty will have been eliminated) and it will represent the systematic bias of the model prediction y_c. This systematic bias is due to the incompleteness (or other shortcomings) of the deterministic model.

The EPRI (1993) engineering model of ground motions is an example of this formulation. The ground-motion amplitude (peak acceleration or spectral acceleration at a certain frequency) is given by

$$\ln A(m,r) = g(m,r) + \varepsilon_e + \varepsilon_a \qquad (2.5)$$

where g(m,r) represents a median attenuation equation, e_e is an epistemic variable with zero mean, and ε_a is a zero-mean aleatory variable. In this example, the additive quantity ε_e plays the role of D and it represents lack of knowledge about the difference between g(m,r) and the logarithm of the true median amplitude for this magnitude and distance. The aleatory variable e_a represents event-to-event and site-to-site scatter due to smaller-scale details of the earthquake source processes and of wave propagation through a heterogeneous crust. Note that in terms of actual observations, the observed scatter is due to both epistemic and aleatory uncertainty and one observation alone would not suffice to remove all epistemic uncertainty. This treatment in terms of model and parameter uncertainty was formalized by Abrahamson, Somerville, and Cornell (1990).

The issue of model uncertainty has been investigated in other applications of risk assessment also. An example from fire risk assessment is given in Siu and Apostolakis (1981 and 1985), in which the updating of the (epistemic) distribution of D is also investigated using Bayes' theorem.

Alternate formulations of model uncertainty have also been proposed. In one such formulation (Apostolakis 1990 and 1993), a number of models calculating $y_{c,j}$, j=1,...,n, are considered corresponding to a set of n hypotheses H_j, j=1,...,n. In the context of ground motion prediction, these alternative hypotheses may represent empirical attenuation equations derived by different investigators using somewhat different data sets and functional forms, attenuation equations derived using different methods (e.g., empirical vs. stochastic), or attenuation equations derived using a stochastic model and varying some key assumption or parameter in a discrete manner.

Another example of discrete hypotheses is in the specification of seismic sources (or seismic source maps). A certain tectonic feature may or may not be active (according to an activity probability, P_a). If it is active, uncertainty about its true geographic extent is typically represented by multiple alternative geometries with associated weights.

Discrete hypotheses with weights are used frequently in PSHA because they are flexible and they are intuitive. Under certain assumptions, a set of model weights can be developed that reflect the relative forecasting accuracy of approximate models. In Appendix J, we develop this view in more detail in the context of some simple opinion aggregation models that are useful for gaining insight.

An alternate interpretation of model uncertainty is offered in Winkler (1993) and Morris (1971). The outputs from a number of models are viewed as information that is to be processed using Bayes' Theorem. This formulation avoids the issue of

developing probabilities regarding the acceptability of models; this approach has not yet been used in risk assessment. The difficulty lies in its implementation, as Winkler states.

2.2.5 Communication of the Uncertainties

Having discussed the concepts of the model of the world and of the epistemic model, we must now address the question of how to communicate these results to others (possibly, decision makers).

To make the discussion concrete, let us consider the (aleatory) probability of no events in the interval (0, t) [see also eq. (2.3)], i.e.,

$$P[\text{no occurrences in t}] = \exp(-\lambda t) \qquad (2.6)$$

Furthermore, suppose that the rate λ has a discrete epistemic distribution given by the following set of doublets: $\{10^{-2}, 0.4\}$ and $\{10^{-3}, 0.6\}$, i.e., we judge that this rate can be either 10^{-2} per year with probability 0.4 or 10^{-3} per year with probability 0.6. The probability of zero occurrences in a period (0, t) can be calculated using eq. (2.6) and is given by a set of doublets: $\{e^{-0.01t}, 0.4\}$ and $\{e^{-0.001t}, 0.6\}$. The two terms in these doublets resulting from the use of eq. (2.7) represent aleatory uncertainty, while the probabilities 0.4 and 0.6 represent epistemic uncertainty.

A frequently used "point" estimate is the mean value of the probability of no occurrences, i.e., $(0.4e^{-0.01t} + 0.6e^{-0.001t})$. This is called the predictive probability of zero events in a period (0,t). Note that we have not considered any model uncertainty in this example.

We note that the use of the average (predictive) value for decision making can create problems, especially when the epistemic uncertainties are very large, as is frequently the case with model uncertainties [see eq. (2.6)]. This is because the average value can be greatly affected by high values of the variable, even though they may be very unlikely. The average value is only a summary measure of the full uncertainty, which is expressed by the set of doublets. In particular, when the epistemic uncertainty is very large, it would be incomplete to report only the average to the decision maker. In our example, reporting all the doublets to the decision maker communicates the full epistemic uncertainty.

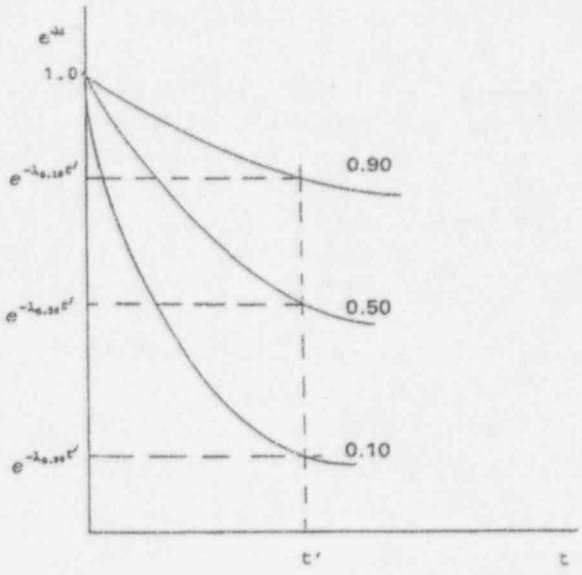

Figure 2-1 A family of aleatory curves displaying epistemic uncertainty

The simple example presented above captures the essence of the uncertainties that we have considered. Other, more realistic, cases can be easily constructed. For example, let us assume that the aleatory model is still given by eq. (2.6), but now the epistemic distribution of λ is a continuous pdf. It is then customary to display a family of curves for various values of λ (Kaplan and Garrick 1981). For example, Figure 2-1 shows three curves produced from eq. (2.6) with λ being equal to the 10th, 50th, and 90th percentiles of its pdf. Also shown are three values of the aleatory probability for a given t' (the analogy with the doublets discussed above can be seen). The interpretation of the curves is as follows: after a very long time, all epistemic uncertainty will have disappeared and the value of λ, and, therefore, the actual curve $e^{-\lambda t}$, will be known. At the present time, we judge with probability 0.90 that this "true" curve will be below the curve labeled "0.90" in the figure.

Figure 2-2 shows the results of a PSHA in this format. The interpretation is just as above. For example, we are 0.85 confident that the "true" seismic hazard curve will turn out to lie below the curve labeled "85th."

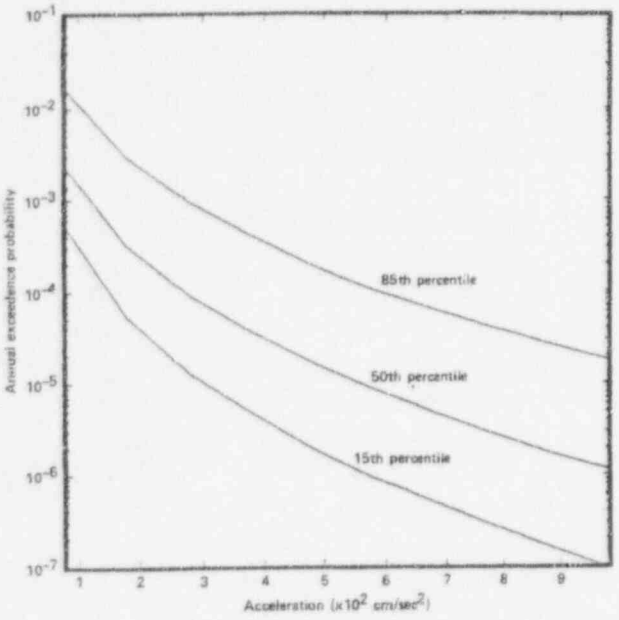

Figure 2-2 A typical family of seismic hazard curves

2.2.6 Further Comments on the Distinction between Aleatory and Epistemic Uncertainties

A recent paper by Veneziano (1994) questions the value of distinguishing between aleatory and epistemic uncertainties. The author claims that this distinction is ambiguous for geologic hazards and raises the issue of the time-dependence of this formalism. Veneziano argues further that what really matters to decision making is the total uncertainty, that is, what we have called the predictive distribution in the last section.

To address these concerns we note that there are both theoretical and practical reasons for this distinction. The theoretical foundation was provided by de Finetti in 1937 (de Finetti 1974; Press 1989). He first introduced the concept of *exchangeability*; an infinite sequence of random variables $X_1,...,X_j,...$ is said to be exchangeable, if the joint distribution of any finite subset of these variables is invariant under permutations of the subscripts. A special and familiar case is when these variables are independent and identically distributed. De Finetti's theorem states that, if the X_j are binary variables, the predictive distribution of r "1's" (for example, r "successes") in n exchangeable trials is given by

$$P[\text{r "1's" in n trials}] = \int_0^1 \binom{n}{r} \theta^r (1-\theta)^{n-r} \, dF(\theta) \qquad (2.7)$$

where $F(\theta)$ is some proper cumulative distribution function on (0, 1). Furthermore, the limit of the relative frequency (r/n), as n becomes large, is θ. This equation shows that the predictive distribution of r "successes" in n exchangeable trials may be obtained as if the trials were independent conditional on θ. In our previous terminology, we would say that the model of the world is the binomial distribution (in this case) and it is obviously aleatory. The epistemic model is the mixing distribution $F(\theta)$. It is this distribution that is updated as evidence becomes available. This theorem can be extended to general random variables and is the cornerstone of the subjectivistic (Bayesian) theory of probability.

While this theorem is fundamental, it does not tell us how to separate aleatory from epistemic uncertainties in actual applications. In the above example of binary variables, e.g., the familiar coin-tossing experiment, it is fairly obvious that the assumption of exchangeable trials is reasonable, therefore the natural candidate for the model of the world is the binomial distribution appearing in eq. (2.7). In a practical situation, the assumptions that may be used in the model of the world may not be obvious. It is useful to briefly discuss what really happens in practice.

In modeling a physical situation of interest, we do not decide *a priori* how to separate the uncertainties. In fact, the question does not even arise. What we do is build the best model that we can making assumptions that are defensible. If we decide that certain aspects of the problem require a probabilistic treatment, we introduce the appropriate models that capture our knowledge about these uncertainties. This is what physical scientists do.

After the model of the world is completed, the formalism that we have adopted requires that the

analysts answer several questions, such as. "Are the basic assumptions of the model valid?" "Are there alternate assumptions that one may adopt?" "Are the numerical values of the parameters of the model known?" In the simple example of eq. (2.7), the only question that is asked is the last one. It is at this point that a new set of uncertainties may be identified. Recognizing this reality and to facilitate *communication* we introduce the term epistemic uncertainties for this new set and we refer to the uncertainties in the model of the world as aleatory. From this perspective, we can now say that this distinction *occurs naturally* and it is not made in advance. Therefore, a clear prescription as to how to separate the two types of uncertainty (other than what we have already said) cannot be given *a priori*. Furthermore, the different terminology is not intended to imply that these uncertainties are of a fundamentally different nature.

The benefits from the formalism are far greater than just the facilitation of communication. By demanding that the analysts ask the preceding questions, this approach imposes a *discipline* on the analysis that has been found to be invaluable in practice. The interpretations of aleatory and epistemic uncertainties that we have discussed in earlier sections are new to most modelers and practitioners and using these concepts forces the analysts to really delve into the details of the models and consequences of various assumptions. An example from the NUREG-1150 studies will clarify this point.

In studying how large power reactors behave in postulated accidents, one important issue is the internal pressure that would cause the large containment building that surrounds the reactor to fail catastrophically. In the NUREG-1150 study of the Peach Bottom 2 reactor, a boiling water reactor, the approach to understanding at what pressure the containment would fail was to ask a group of experts to provide their judgments (Amos et al. 1987).

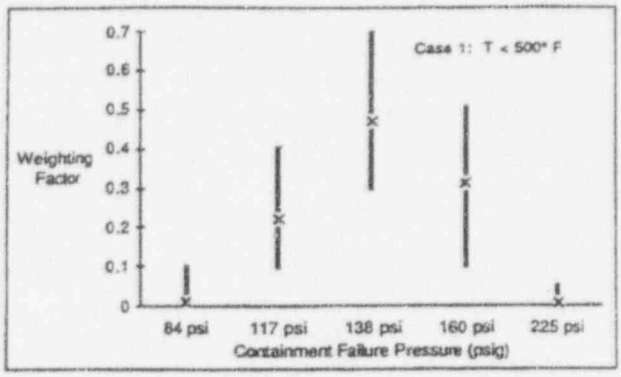

Figure 2-3 Weighting Factors for Containment Capacity for Low Temperature Conditions [NUREG/CR-4551/Volume 3, Figure 4-1]

Figure 2-3 shows the expert-supplied probabilities ("weighting factors") for the containment failure pressure (given certain conditions). Five pressure values are being considered. This model of the world implies that one of these pressures will actually turn out to be the true one and, at this time, we do not know which value it will be, i.e., the model of the world is deterministic. It is stated in Appendix A of the report that there might be some randomness about each value and that "there was a great deal of discussion concerning this issue due to the difficulties in defining the meaning of the failure pressure distributions derived for this issue. Each reviewer had a somewhat different interpretation of the input that was being required, as well as of the use of the input in the Limited Latin Hypercube sensitivity analysis." This means that the experts debated the validity of the assumption that the model of the world was deterministic (that only one value was the true failure pressure). It was finally decided that the aleatory variability was "generally small" and it was dropped from further analysis. We note that, if the group had decided to include aleatory uncertainty in the model, the question asked of the experts would have been "what is the fraction of times that failure occurs at each of these pressures and what is your uncertainty about this fraction?" The results of Fig. 2-3 are responses to a different question, i.e., "what are your probabilities that the true failure pressure will be one of the five values shown in the figure?"

The predictive distribution contains all the uncertainties and is the one that is used in formal decision theory to evaluate the expected utilities. What the presented formalism does is allow for the systematic assessment of this distribution. In practice, the epistemic uncertainties themselves may suggest possible actions, such as delaying the actual decision and doing more research to reduce the magnitude of the epistemic uncertainties.

As discussed in previous sections, new information is used to update the epistemic model using Bayes' theorem. This is based on the fundamental assumption that the models of the world with which we begin the analysis do not change. This is not the way engineering models, especially the ones employed in risk assessments, evolve in time. New evidence and advances in science very often lead to new models. In these cases, the old formulation does not apply anymore and one must start with new models of the world and ask, again, the above questions to assess the epistemic uncertainties associated with these new models. Thus, the new predictive distributions will be evaluated in the same manner as before. It is evident, therefore, that models used in present analyses may be used only for a limited time depending on how sound their assumptions are. For example, the models for PSHA that this report presents and the associated guidance for hazard assessment are expected t_ useful for the next several years. When new scientific advances necessitate a significant change in the models of the world, then the structure of the analysis will change and the formalism that we have discussed will apply to this new structure.

3. STRUCTURING AND IMPLEMENTING A PSHA

3.1 Introduction

The success of a PSHA project is principally determined by how it is structured and implemented to derive inputs; in particular, how this structure and implementation account for different technical interpretations of the available evidence and uncertainties. Despite the importance of these issues—that is, *how* the PSHA is conducted, rather than what goes into the analysis—very little written guidance has been developed. Such guidance should not only incorporate the evolving concepts related to uncertainty treatment and expert elicitation in general, but, perhaps more importantly, should draw on the experience base developed over the past decade or more in carrying out seismic hazard analyses.

PSHA inputs involve multiple issues, e.g., ground motion models, ground motion uncertainty, seismic source identification, seismicity parameters, etc. The complexity, importance and diversity of judgments within the appropriate scientific community regarding any one of these issues vary between study location (east vs. west U.S.), range of the study (site-specific vs. regional), and other factors.

3.1.1 Principle 1: The Basis for the Inputs

A basic principle defined by the Committee is that *the underlying basis for the inputs related to any of these issues must be the composite distribution of views represented in the appropriate scientific community*. Expert judgment is used to represent the informed scientific community's state of knowledge. Of course, it is impractical—and unnecessary—to engage an entire scientific community in any meaningful interactive process. Decision makers must always rely on a smaller, but representative, set of experts. Thus, we view an expert panel as a sample of the overall expert community and the individual Technical Integrator (defined later) as the expert "pollster" of that community, the one responsible for capturing efficiently and quantitatively the community's degree of consensus or diversity.

Regardless of the scale of the PSHA study, the goal remains the same: *to represent the center, the body, and the range of technical interpretations that the larger informed technical community would have if they were to conduct the study.*

3.1.2 Principle 2: A Clear Definition of Ownership

Another principle defined by the Committee with respect to deriving inputs is that *it is absolutely necessary that there be a clear definition of ownership of the inputs into the PSHA* (and hence ownership of the results of the PSHA). Therefore, this precludes the PSHA being performed by an analyst who simply accumulates inputs (either from the literature or eliciting the judgments of one or more experts) without establishing his/her responsibility for and ownership of aggregated results. That is, it is important that the analyst be an integrator in the sense of establishing his/her ownership of the results.

The number and type of PSHAs conducted over the past decade scales very much like an earthquake recurrence curve: hundreds of "small-magnitude" hazard studies are conducted annually to evaluate ground motions at the sites of new or existing conventional facilities (e.g., buildings, pipeline terminal facilities, hospitals); tens of "moderate-magnitude" studies are conducted for more critical facilities (e.g., high-rise buildings, nuclear production facilities, offshore platforms); and a few large-magnitude studies have been conducted over the past decade for highly critical and/or highly regulated facilities (e.g., nuclear power plants, high-level waste repositories). Given the nature of the "PSHA-experience recurrence curve," our collective experience base is decidedly skewed toward the smaller-scale studies. However, the sponsors of the SSHAC project are interested in guidance that is skewed toward application for more critical facilities.

Despite an emphasis on large-scale studies for critical facilities, SSHAC has devoted considerable time and effort in reviewing past PSHAs of all scales and applications in order to

learn which processes have worked, the pitfalls, and the processes that appear to hold the most promise in future application. One example of SSHAC's research effort is the Seismic Source Characterization (SSC) Workshop (see Appendix H), in which the focus of the discussions was on *process* issues: the manner in which SSC experts should be elicited, the degree of expert interaction desired, the value of workshops, methods for combining the interpretations of multiple SSC experts, etc. In SSHAC's ground motion workshops, concepts related to facilitation of workshops, roles and interactions of experts, and aggregation of expert interpretations were "tested" through a real application.

The two basic SSHAC principles discussed above—(1) inputs should represent the composite distribution of the informed technical community and (2) ownership of inputs established by an integrator—SSHAC recommends that the derivation of inputs be conducted by one of two approaches, either by a Technical Integrator (TI) approach or a Technical Facilitator/Integrator (TFI) approach. Appropriate definitions of these approaches and SSHAC recommendations of the structures of these approaches as a function of the importance, complexity, diversity of views and contentiousness of an issue is the subject of Section 3.1. Sections 3.2 and 3.3 outline the TI and TFI approaches, respectively. Because it has never been fully implemented, hence not documented, the TFI discussion in Section 3.3 (and Appendix J) is more comprehensive. Section 3.4 discusses peer review and summarizes SSHAC recommendations with regard to peer review of the TI and TFI approaches to deriving inputs on any issue related to PSHA. It should be noted that detailed guidance for implementing the TI and TFI approaches for seismic source characterization are found at the end of Chapter 4 and detailed guidance for implementing the TI and TFI approaches for ground motion assessment is given at the end of Chapter 5.

3.1.3 Definitions and Roles of Technical Integrator (TI) and Technical Integrator/Facilitator (TFI)

To outline clearly the Committee's recommended approaches to the PSHA input issues, it is necessary to define some important terms and concepts.

3.1.3.1 Project Sponsorship and Leadership

- The *Project Sponsor* is the entity that provides the financial support for the project, hires the study team (including the project leader), and "owns" the study's results in the sense of property ownership.
- The *Project Leader* (often one individual, but possibly a small team) is the entity that takes managerial and technical responsibility for organizing and executing the project, oversees all other project participants, and "owns" the study's results in the sense of assuming intellectual responsibility for the project's overall technical validity. The Project Leader makes decisions regarding the level of study of particular issues (discussed below).

3.1.3.2 integrators

Two types of integrators are considered:

- *Technical Integrator (TI)*: a single entity (individual, team, company, etc.) who is responsible for ultimately developing the composite representation of the informed technical community (herein called the community distribution) for the issues using the TI approach. As discussed later, this could involve deriving information relevant to an issue from the open literature or through discussions with experts.
- *Technical Facilitator/Integrator (TFI)*: a single entity (individual, team, company, etc.) who is responsible for aggregating the judgments and community distributions of a panel of experts to develop the composite distribution of the informed technical community for the issues using the TFI approach. The key differences between the TI and TFI approaches are (i) the facilitator role of the TFI in which he/she is responsible for

facilitating the discussions and interactions between experts and (ii) the use of "evaluator" experts, who act as individual integrators, in the development of the community distribution.

In the context of these discussions we use the term *interpretation* to denote a technical hypothesis (i.e., without epistemic uncertainty), and *evaluation* to denote a weighted set of hypotheses or interpretations. The evaluation process, then, is focused on epistemic uncertainties.

3.1.3.3 PSHA Issues

By reviewing past hazard studies and experimenting with "new" approaches, SSHAC has formulated a spectrum of approaches to structuring a PSHA. It is concluded that all approaches attempt to achieve several primary objectives. These objectives include: proper and full incorporation of uncertainties, inclusion of a range of diverse technical interpretations, consideration of site-specific knowledge and data sets, complete documentation of the process and results, clear responsibility for the conduct of the study, and proper peer review.

It is recognized that PSHA can, and should, be conducted for a wide variety of reasons and at various scales. There is nothing inherently "wrong" with the calculated results that come from a modest hazard analysis conducted by a single contractor; nor does the use of multiple experts in a large-scale project guarantee that the hazard results are more defendable (particularly if done poorly). They are, however, more likely to capture accurately the scientific community's information. The choice of the level of PSHA is often driven by the level of uncertainty and contention associated with a particular project, as well as the amount of resources available for the study. It is further recognized that particular components or *issues* of the PSHA (e.g., the slip rate on a particular fault, the maximum magnitude, or the amplitude of near-field ground motions) may have variable degrees of contention and/or uncertainty.

Table 3-1 Degrees of PSHA Issues and Levels of Study

ISSUE DEGREE	DECISION FACTORS	STUDY LEVEL
A Non-controversial; and/or insignificant to hazard	•Regulatory concern •Resources available •Public perception	1 **TI** evaluates/weights models based on literature review and experience; estimates community distribution
B Significant uncertainty and diversity; controversial; and complex		2 **TI** interacts with proponents & resource experts to identify issues and interpretations; estimates community distribution
C Highly contentious; significant to hazard; and highly complex		3 **TI** brings together proponents & resource experts for debate and interaction; TI focuses debate and evaluates alternative interpretations; estimates community distribution
		4 **TFI** organizes panel of experts to interpret and evaluate; focuses discussions; avoids inappropriate behavior on part of evaluators; draws picture of evaluators' estimate of the community's composite distribution; has ultimate responsibility for project

As used here, an issue may be one or a combination of input issues, e.g., an issue may only be ground motion models for median values, or an issue may include both the ground motion model for the median as well as the aleatory uncertainty in the ground motion. Table 3-1 illustrates the process of identifying the key technical PSHA issues and deciding the level of study that should be devoted to addressing the issues. It is assumed that individual issues or components of a PSHA can be evaluated separately, although, commonly in the past, the decision regarding the level of study has applied to the entire hazard analysis. In the left-hand column of the table, the degrees of issues are shown as A, B, and C. In deciding the degree of an issue, there are several considerations such as:

- the significance of the issue to the final results of the PSHA
- the issue's technical complexity and level of uncertainty
- the amount of technical contention about the issue in the technical community
- important non-technical considerations such as budgetary, regulatory, scheduling, or other concerns.

Degree A issues are non-controversial and/or have no significance to the seismic hazard results; Degree B issues are more controversial, complex, and significant to the hazard results; Degree C issues are often highly contentious, complex, and most significant to the hazard results. Obviously, there is a continuum of degree so that the three levels identified represent a coarse partition of the range of issue degrees. Some judgment must be made when classifying any particular issue for a given study.

For each issue (or for the PSHA as a whole) a decision must be made regarding the level of study that will be conducted to address the issue. The decision usually involves factors such as the regulatory framework, the resources (money and time) available to conduct the study, perceptions (including both the public and other stakeholders) of the importance of the project, and scheduling constraints.

3.1.3.4 Experts

Because of limited data, it is often necessary to interact with experts to derive necessary information regarding an issue. For purposes of PSHA issues, three types of experts, not necessarily mutually exclusive, are identified: proponents, evaluators, and resource experts. An important distinction is made here in the roles of experts as "proponents," as "evaluators," and as "resource experts." A *proponent* is an expert who advocates a particular hypothesis or technical position. The proponent role is common in science, whereby an individual evaluates data and develops a particular hypothesis to explain the data. The proponent's position is then challenged technically by his peers in professional debates and in the literature to see if it stands up to a variety of observations. The proponent of the hypothesis detaches himself professionally from the success or failure of the hypothesis; that is, although he argues for the viability of the hypothesis, he recognizes that it may ultimately be proven wrong. With time, the hypothesis will gain increasing support with additional data or will lose favor in the scientific community.

An *evaluator* is an expert who is capable of evaluating the relative credibility of multiple alternative hypotheses to explain the observations. The evaluators are expected to evaluate all potential hypotheses and bases of inputs from proponents and resource experts and provide 1) their own input and 2) their representation of the community distribution. The evaluator recognizes that the evaluation occurs at a particular point in time and, as a result, the viability of any particular hypothesis is uncertain and may not be proven until some time in the future. To evaluate the alternatives, the evaluator considers the available data, listens to proponents and other evaluators, questions the technical basis for their conclusions, and challenges the proponents' positions. In the end, the evaluator is able to assign relative credibilities to the alternative hypotheses. He recognizes, too, that no single hypothesis is likely to be the ultimate truth—it is only a current representation. Therefore, he finally may, for example, assign a smooth continuous (epistemic) uncertainty distribution over that parameter (e.g.,

the median peak acceleration of a magnitude 6 at 10 kilometer distance or the long-term slip rate on a fault) to which each hypothesis (model/interpretation/data set) assigns a unique value, and for which a finite set of weighted hypotheses would imply a simple discrete uncertainty mass function.

A *resource expert* is a technical expert with particular knowledge of a particular data set of importance to a PSHA. Commonly, a resource expert will have site-specific experience that will be of use to the evaluators. For example, a resource expert for a site-specific hazard analysis might be a geologist who has mapped and evaluated nearby faults. A resource expert might also have expertise in particular methodologies or procedures of use to the evaluators. For example, a resource expert may have developed new procedures for evaluating the completeness of earthquake catalogs or for processing catalogs to identify foreshocks and aftershocks.

3.1.3.5 Study Level

Table 3-1 summarizes four levels of study to address issues, which are shown roughly in order of increasing resources and sophistication. The TI and TFI roles are outlined for the various levels. Because the TFI, by definition, involves the "facilitation" of multiple experts, the TFI role does not appear until the Level 4 analysis. The TI, on the other hand, varies in his/her role from basing judgments on his own experience and literature to obtaining input from communication with other experts.

The roles and activities associated with the TI show increasing input from technical experts with increasing level. For example, at Level 2, the TI reviews the literature and contacts those individuals who have developed interpretations or who have particular site experience. At Level 3, however, the TI gains additional insight by bringing together the experts and focusing their interactions. In these sessions, the experts could have an opportunity to explain their hypotheses and data bases. Further, proponents or advocates of particular technical positions can defend their positions to other experts.

In the context described above, the Level 2–4 analyses involve the input from proponents who have developed technical interpretations regarding particular issues of importance to PSHA. Levels 2 through 4 differ in the degree to which these proponents are questioned directly and/or are given a forum for expressing their views. In Levels 1 to 3, the TI plays the role of the "evaluator." In Level 4, a group of expert "evaluators" is identified and their judgments are elicited. The TFI is responsible for identifying the roles of the proponents and evaluators and for ensuring that their interactions provide an opportunity for focused discussion and challenge.

It is important to note that in all four levels of hazard analysis, the responsibility for the success or failure of the analysis rests with the TI or the TFI. In the Level 1-3 analyses, the responsibility is clear inasmuch as the TI develops judgments and hazard inputs based on information gathered from others. In the Level 4 analysis, resources permit and the situation dictates multiple evaluators and hence a TFI to take responsibility for the aggregated product. The TFI must organize and manage interactions among the proponents and evaluators, must identify and mitigate problems that might develop during the course of the study (e.g., an expert who is unwilling or unable to play the evaluator role), and must ensure that the evaluators' judgments are properly represented and documented. In both the TI and TFI approaches, proper peer review must be conducted to review the process and substance of the study.

In the TI approach, it is clear that the *intellectual responsibility* for the study lies with the TI. Intellectual responsibility is defined as the responsibility not only for the accuracy and completeness of the results, but also for the process used to arrive at the results. In the TFI approach, *both* the TFI and the experts have intellectual responsibility for the results. The TFI has a further burden of ensuring that the process is properly implemented. In most cases, peer reviewers are expected to provide an *endorsement* of the process and results of the study. An endorsement is an affirmation that the particular

project meets his/her standards of quality, thoroughness, and validity.

Regardless of the level of the study, the goal in the various approaches is the same: to provide a representation of the informed scientific community's view of the important components and issues and, finally, the seismic hazard. ("Informed" in this sense assumes, hypothetically perhaps, that the community of experts were provided with the same data and level of interaction as that of the evaluators). This is done by the TI in the Level 1-3 studies, with various levels of input from representatives of the community and their literature. In the Level 4 analysis, multiple evaluators provide their review and synthesis of the available data and formulate interpretations that represent their assessments and uncertainties. As will be discussed in Section 3.3.4, the evaluators will be asked to represent both their own interpretations and uncertainties (Stage I elicitation) and their view of the informed community's composite interpretation (Stage II elicitation). In the latter sense, they are themselves each acting as integrators in evaluating the community's views.

Because there have been relatively few Level 4 studies (EPRI and LLNL are examples), there is not a large experience base on which to build guidance. Further, the adoption of the TFI process introduces certain new ground and processes for structuring expert interaction for which detailed guidance must be developed. For this reason a large part of this chapter is devoted to implementation advice for the TFI. In contrast, the TI process is much more common and founded in application. Therefore, the discussion devoted to the TI approach is more limited and based on numerous examples.

3.2 The Technical Integrator (TI) Approach

3.2.1 Introduction

PSHAs in which a single entity is responsible for specifying all inputs into a PSHA, as well as performing the necessary calculations, uncertainty analyses, and documenting the process and results have considerable precedence. In most of these applications, the specified inputs are developed and therefore owned by the single entity, that is, a TI or Technical Integrator. In most of these applications there has not been a formal elicitation of expert judgments. The single entity is the sole evaluator of the information available—either published, as espoused by proponent(s) or described by resource experts—and, hence, is sole developer of the representation of the community distribution. This feature applies to all three levels of study and is one of the distinguishing differences between the TI approach and the TFI approach in which a formal panel of evaluator experts is used jointly as (i) sources of their personal inputs, and (ii) source of representation of the community distribution. The distinction between the three levels of study using the TI approach is a matter of level of resources used by the TI to develop his/her representation of the community distribution.

In modest (Level 1) studies, the TI utilizes the interpretations found in the published literature, supplemented by informal discussions with other researchers. For example, consider a site-specific PSHA for a bridge site in southeastern Illinois whereby the engineers are interested in evaluating the integrity of the structure when subjected to 500-year ground motions. The TI should review the literature for previous hazard studies that have been conducted in the area (e.g., the EPRI and LLNL studies, the national hazard map by the USGS, studies by the Illinois Department of Transportation, recent studies of seismicity in the New Madrid seismic zone, recent paleoseismic studies in the Wabash Valley, etc.) The goal in these reviews is to understand and, in turn, represent in the hazard analysis the present level of knowledge and uncertainty in the seismic environment of the bridge site. Peer review for a Level 1 study can be quite modest (say 10% to 20% of the total effort), but still serves the valuable function of providing review of the process followed and review of the data and interpretations.

Assuming that the study was a Level 2 study (which implies additional resources), the TI would communicate with the authors of published studies and other local experts who have expertise

in the region or in regional ground motions. The goal in these interactions would be to hear and understand the technical positions taken by various proponents of particular hypotheses. For example, the TI might probe the basis of interpretations taken by a paleoseismologist who is advocating the view that his data set in the Wabash Valley suggests that large-magnitude earthquakes strike the region on the average every 250 to 500 years. What is the basis for identifying the paleoearthquakes? What evidence suggests that these events are large? What are the uncertainties in the age estimates? What do others think about these data and conclusions? In the course of these exchanges, the TI would evaluate the viability and credibility of the various hypotheses with an eye toward capturing the range of interpretations, their credibilities, and uncertainties. In effect, the TI is acting as an "integrator" of the various interpretations and is attempting to provide an overall assessment that would represent the *informed scientific community's view* of the subject, if the community were to make such an assessment. This goal should be common to all Levels of PSHA, whether a Level 1 TI approach or a Level 4 TFI approach.

To complete the example, assume that the bridge in Illinois is the largest suspension bridge in the Midwest and that the issue of large-magnitude paleoearthquakes has been deemed by the U.S. Department of Transportation as a critical issue that must be addressed in order for the state to qualify for federal retrofit funds. In this case, the project sponsor/project leader may conclude that this issue is a degree "C" issue, following the categorization discussed previously. Further, he may chose to conduct a Level 3 study focusing on the paleoseismic issue, feeding into a Level 1 PSHA for the remaining issues. To conduct the Level 3 analysis, the TI would bring together the technical experts and proponents of various hypotheses for debate and interaction in, perhaps, one or more workshops. The TI would focus the debates in a way that would highlight the issues of most significance to the PSHA (e.g., indicators of the magnitude, location, and recurrence of paleoearthquakes). The TI would probe the viability of the arguments for and against the hypotheses and would attempt to encourage active interaction of the advocates of various technical positions. The result of this process would be a representation developed by the TI of the diversity of interpretations and their uncertainties.

A key aspect of the TI approach is the use of peer review to assure that the process followed was adequate and to ensure that the results provide a reasonable representation of the diversity of views of the technical community. Peer review has long been a cornerstone of quality assurance procedures for PSHA. Usually peer review is conducted in the final *late-stages* of the project and involves the review of draft and final project documents. In recent years, through large projects such as the Diablo Canyon Power Plant Long Term Seismic Program, DOE's New Production Reactor studies, and the Caltrans Seismic Hazard Evaluations for the San Francisco Bay Area bridges, the process of a *participatory* peer review has been developed and implemented. In this approach, the peer reviewers are actively involved in reviewing the project throughout its implementation. In this way, the peer reviewers are able to provide advice regarding changes in the course of the study as it evolves. They are thus in a better position to evaluate the process of the study and not just the final results. Of course, this entails some loss of independence of the reviewers (see Section 3.4 for a more detailed discussion of peer review).

In application, the TI approach has been most commonly applied to site-specific seismic hazard analyses, and less commonly to regional seismic hazard analyses. Often, site-specific analyses include site-specific data that have been developed with the particular purpose of evaluating the seismic hazard. For example, the PSHA for the Diablo Canyon Long Term Seismic Program and the PSHA for Rocky Flats DOE site both included geologic or geotechnical data gathered with the specific purpose of evaluating the site ground motions. Because of this focus on site-specific information, the TI approach has been well-suited to directly incorporating this information into the hazard analysis through a thorough review by the responsible TI. The TFI approach, which includes the assessments by

multiple experts, requires the review of all pertinent data sets by the multiple experts in order for them to make an informed assessment. The additional resources required to do so are usually not available and/or the project sponsor decides that value gained for the additional resources are not required for the particular project.

3.2.2 The TI Process

This section summarizes the recommended process that should be followed according to the TI process. The guidance provided here is general and is not elaborate. This reflects the fact that the TI approach has common application and is well-tested. Detailed guidance is provided for seismic source characterization and ground motions applications in Sections 4.4 and 5.6, respectively. In the following discussion, the steps of the process are presented with an assumption of a moderate-scale (Level 2) analysis. The reader can infer that a Level 1 analysis would involve lesser activity and a Level 3 analysis would involve additional activities, particularly with regard to the communication with experts within the technical community.

Step 1 Identify and select peer reviewers

The Project Leader, perhaps in conjunction with the Project Sponsor, is responsible for identifying and selecting peer reviewers. Selection criteria for the peer reviewers includes such attributes as the following:

- Earth scientist having a good professional reputation and widely recognized competence based on academic training and relevant experience.
- Understanding of the general problem area through experience collecting and analyzing research data for the same or comparable environments.
- Availability and willingness to participate as a named peer review panel member, including a commitment to devoting the necessary time and effort to the project.
- Personal attributes that include strong communication and interpersonal skills, flexibility and impartiality, and the ability to simplify and generalize.

Peer reviewers, particularly those involved in a "participatory peer review" (see Section 3.4), should be prepared to question and provide meaningful guidance to the Project Leader and the TI (or TFI) on both the process being followed and the technical substance of the project. The project should be conducted such that the peer reviewers will endorse the process and the substance of the project at its completion.

Step 2. Identify available information and design analyses and information retrieval methods

The TI is responsible for assembling all relevant technical data bases and other information important to the hazard analysis. This includes any site-specific data that may have been gathered specifically for the hazard analysis (e.g., geologic maps, results of fault studies, geotechnical properties of soils, etc.). The TI also identifies technical researchers and proponents that he/she intends to contact during the course of the study to gain insight into their positions and interpretations (in a Level 3 analysis, the TI identifies those individuals that he intends to assemble for discussions and interactions). In addition, the TI defines the procedures and methods that will be followed in conducting the hazard analysis.

Step 3. Perform analyses, accumulate information relevant to issue and develop representation of community distribution

This Step is the heart of the TI work. Specifically, the TI is responsible for understanding the entire spectrum of technical information that can be brought to bear on the issue at hand. This includes the written literature, recent work by other experts, and other technical sources. In advanced technical work, it is always the responsibility of the investigator to learn about the most recent advances in the field, often by direct contact with other experts via personal correspondence, personal meetings, telephone conversations, and so on. In a Level 3 study, members of the technical community are also brought together and the TI orchestrates interactions and, possibly, workshops to focus the discussions on the technical issues of most significance to the hazard and to be sure that he is aware of the diversity in interpretations for these key issues. The TI uses

all of this information to develop a community distribution of the range of uncertainty for the particular issues being addressed.

Step 4. Perform data diagnostics and respond to peer reviews

Interactive peer review during the analysis is very important. The TI can use the peer review team as a sounding board to learn whether the full range of technical views has been identified and assimilated into the project. If key aspects are difficult to resolve because different technical views exist among respected experts, the peer reviewers are vitally important. Peer review of the process depends on the type of peer review used (see Section 3.4). If participatory peer review is applicable, on-going review would occur after steps 2, 3, and 4 with appropriate response to the reviews.

The fact that experts are not brought into the process in a formal sense, as in the TFI approach, means that the TFI guidance on "expert buy-in" does not apply directly. However, it does apply indirectly, and that aspect of the TFI guidance should be studied (subsection 3.3). Specifically, if the TI develops a controversial interpretation that represents an integration of diverse technical views of differing experts, it is very important that an attempt be made to obtain the views of the specific advocates of the various technical positions involved. The peer reviewers can verify that this contact has been fulfilled and that the various interpretations are properly represented. In SSHAC's opinion, if these experts can "buy into" the process that the TI has used to integrate the different views, the credibility of the ultimate result of the TI's effort will be significantly strengthened.

A variety of sensitivity analyses should be carried out and shared with the peer reviewers to understand the most significant issues, sources of uncertainty, and data sets used to address the issues.

Step 5. Document process and results

This step is vital to an understanding of the study by third parties. Although relatively straightforward, it is important to emphasize that the TI be attentive to the documentation guidance in Chapter 7.

3.3 The Technical Facilitator/Integrator (TFI) Approach

3.3.1 Introduction

In a significant enhancement to current practice, we introduce the concept of the Technical Facilitator/Integrator (TFI). The TFI is a single entity who has the responsibility and is empowered to represent the composite state of information regarding a technical issue of the scientific community. In the TFI process, the selected experts act, not as proponents of one specific viewpoint, but as informed *evaluators* of a range of hypotheses and models. 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.

The TFI process is centered on the precept of thorough and well-documented expert interaction as the principal mechanism for integration. Much of the "work" in the TFI process occurs in the context of face-to-face expert information and viewpoint exchanges that take place over a series of carefully structured meetings and workshops.

In contrast with the classical role of experts on a panel as individuals providing inputs to a separate aggregation process, the panel is viewed as a team, with the TFI as team leader, working together to arrive at, first, a composite representation of the knowledge of the group and, second, a composite representation of the knowledge of the community at large (these representations may or may not reflect panel consensus). The process is transparent to the experts at all stages in contrast with previous PSHA studies in which some experts have complained that the aggregation process was a "black box."

The TFI conducts individual elicitations and group interactions, and with the help of the experts themselves, integrates data, models, and interpretations to arrive at the final product—a

full probabilistic characterization of seismic hazard. Together with the experts, the TFI "owns" the study and defends it as appropriate.

The TFI is a special role that only comes into play in a Level 4 analysis in which an issue is complex and controversial enough to warrant the challenge and expense of a suite of multiple integrators. The advantages of bringing increased wisdom and experience to bear on a difficult problem come at the 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 past PSHA projects. On future projects that warrant Level 4 analyses, the TFI process described below offers some new and unique advantages over previous PSHA multiple-expert processes.

The distinction between the novel roles of experts as evaluators and integrators and the traditional role of experts as proponents of a particular scientific point of view is fundamental to the TFI process and is not well-defined in current multiple-expert use literature and applications. The TFI methodology does rely heavily, however, on relevant published decision science research, and incorporates our best understanding of state-of-the-art methods for eliciting and aggregating expert judgments. Moreover, the TFI concept is based on a detailed review of the problems and issues of past studies, particularly the large EPRI and LLNL PSHA studies of the late 1980's and early 1990's (discussed in Section 3.3.2.2 below).

The remainder of this section focuses on two unique TFI roles: 1) the TFI as a Technical Facilitator who structures and guides the interaction of a panel of experts, each of whom evaluates the full range of models and interpretations, supported by expert proponents who explain and defend specific models and interpretations, and 2) the TFI as a responsible integrator whose objective is to develop a composite characterization for the expert community based on the panel's inputs. 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 diversity in the set of panel estimates, each of which may itself represent a weighted combination of proponent models and positions. The TFI can be viewed as performing an integration assisted by a group of experts who provide integration advice.

TFI Responsibilities

In carrying out the two roles, the TFI conducts a systematic process, which entails a number of specific responsibilities:

- **Facilitator** —structures and documents full information, data and judgment exchange; stages effective, professional face-to-face debates and interactions in critical areas; ensures that the group identifies all strengths and weaknesses of key data and modeling approaches; elicits formal evaluations from each expert; creates conditions that enable a direct, non-controversial integration of the experts' judgments.
- **Integrator** —develops a final composite assessment (in explicit probability distributions that can be incorporated in the PSHA calculations); explains and defends this assessment before the panel; obtains feedback and concurrence (to the maximum degree possible); explains and defends the composite representation to the outside, i.e., to other experts, the peer reviewers, and all interested parties (e.g., policy makers and regulators).

It is clear that the TFI must have the stature and expertise to deal authoritatively with the multiplicity of disciplines and individuals. It is doubtful that one individual can be identified who will possess all of the qualities required of a TFI. It is more reasonable to anticipate that the TFI will consist of a small group of individuals, typically, two or three. At least one individual should have "substantive" knowledge of the subject matter, e.g., seismic source characterization or ground motion modeling; as a "specialist," he or she should be at least as qualified as the members of the panel on the technical issues. Another role (often another individual) will be that of a "PSHA expert" who knows how PSHA works and how the experts' inputs might affect the final results. One of the

substantive experts must be comfortable in the role of group facilitator (defined below). Finally, one member of the TFI team should be an "elicitation" expert (sometimes called a "normative" expert), i.e., an expert on individual- and multiple-expert elicitation processes, as well as in decision analysis and probability theory, especially on methods for processing evidence.

Goals of the TFI Process

In applications and presentations of the TFI process, observers have often asked the following questions:

> Does the TFI process always result in a consensus among experts?
>
> If not, are the expert judgments equally weighted?
>
> If, for some reason, the expert judgments aren't equally weighted, then what?
>
> Who chooses the weights and how?

These questions are natural because most if not all existing multiple-expert processes have 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 preset objective but rather proceeds through a pushdown list of objectives, attempting to achieve the simplest, least controversial end state possible. In designing the TFI process, we recognized that the answer to each question depends critically on the objectives of the exercise and on the specific issue being addressed. For example, while consensus and equal weights are highly desirable, they are only appropriate under certain conditions (described below). However, these conditions can be controlled and SSHAC believes that equal weights, at least, can usually be attained with sufficiently structured intensive expert interaction. Also, we shall describe different types of consensus, each of which has an *a priori* different likelihood of being achievable.

Notice carefully that each expert, as in Level 3, documents and takes technical responsibility for his or her personal interpretation. 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 expert-as-integrators estimates of the community distribution. The TFI is responsible for documenting and defending how the composite representation was developed, be it by equal weighting of the individual expert estimates of the community distribution or, if necessary, by means more appropriate to the particular circumstances.

Thus, rather than pre-specifying the outcome of the integration process, 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). In the rare case in which simple integration is not appropriate, additional guidance is provided.

Expert Roles

Title	Individual Evaluator	Integrator	Proponent	Resource Expert
Product	Individual Interpretation	Estimate of Community Distribution	Presentation of a Model	TFI Assistance

Reader's Guide

The TFI process includes four separate expert roles. As previewed above, the two primary roles for panel members are "Evaluator" and "Integrator," but both panel and non-panel member experts will, on occasion, be asked to play the role of "Proponent," and an important non-panel member expert role is that of "Resource Expert" who assists the TFI in a number of important activities (described below). To assist the reader, the table below summarizes these roles and their products:

The remainder of 3.3 and its companion Appendix J are written at two levels with some intentional redundancy. The first three subsections provide a basic understanding of the TFI process and its rationale. They are organized as follows:

3.3.2 Historical Context and Motivation for TFI Approach

3.3.2.1 General Approaches to Expert Use

3.3.2.2 Historical PSHA Approaches to Expert Use —Lessons Learned

3.3.3 Underlying Logic of TFI Process

3.3.4 TFI Elicitation Process

An important adjunct to this section, Appendix J, provides additional guidance, including background on the principles underlying the TFI process and presentation of a set of specific implementation tips and traps. It is organized as follows:

Appendix J. Guidance on TFI Principles and Procedures

Section 1. Guidance on Historical Approaches to Expert Use

Section 2. Guidance on Facilitation

Section 3. Guidance on Integration

Section 4. Guidance on the Two-Stage Elicitation Process

Appendix J's additional guidance is written at a more technical level and is "must reading" for potential TFI's who need detailed how-to-do-it instructions, and for expert-aggregation specialists who wish to delve more deeply into the expert-aggregation issues underlying the TFI approach.

3.3.2 Historical Context

To place the TFI approach in perspective, it is useful to review existing approaches to the use-of-experts problem. We start with a brief overview of general schemes and then focus on previous PSHA studies, highlighting some of the lessons learned that led to the TFI approach. This section is a condensed version of Section 1 of Appendix J, which goes into greater depth and includes specific references to related work.

3.3.2.1 General Expert Use Approaches

Historically, two basic types of expert use processes have been used in general practice, mostly focusing on the problem of aggregating the judgments of multiple experts:

- Mathematical Schemes, in which expert inputs are combined using a mathematical formula, and
- Behavioral Schemes, in which aggregation is accomplished through consensus building or some type of qualitative judgment by an individual or negotiated group decision

A great variety of mathematical schemes have been proposed and reviewed in the decision science literature, ranging from linear and logarithmic opinion pools, equal and non-equal weights on expert probability distributions, weights on the parameter values of underlying probability distributions, and Bayesian models (references are provided in Section 1 of Appendix J). Most behavioral schemes are centered around some type of group facilitation process in which the group, through either structured or unstructured interaction, is given the objective of reaching complete agreement on some technical issue.

Mathematical aggregation has several advantages. The logic is transparent and completely checkable. Combination formulas can isolate and separate specific assessments of dependence, expertise, and overlap, so that sensitivity studies

are straightforward. Unfortunately, given the current state of the art, there are several substantial disadvantages to mathematical aggregation. Mathematical models are not advanced enough to include all the factors that are important.

Classic consensus-building processes are usually designed to encourage a group to reach consensus on a technical issue, such as the best estimate of median ground motion for a region or the annual frequency of characteristic magnitudes for a fault. The major advantage of this type of scheme is that, if the information exchange is full and unbiased, and if the result truly reflects each expert's state of information, then the consensus result is appropriate, credible and non-controversial. Unfortunately, there are several problems with such methods. The overriding concern is whether the result is a true consensus that accurately reflects the diversity of education, experience and reasoning within a group, or whether it is more the result of negotiation and strong personalities. There is also the risk of understating the appropriate range of uncertainty by suppressing discussion of differences and focusing on points of agreement.

Should consensus on a technical issue be an objective? In theory, where there is substantial uncertainty, this type of consensus should rarely occur. In practice, technical consensus is better viewed as a convenient result, not as an objective. For example, in the SSHAC ground-motion workshops, the experts, even after thorough group interaction, had diverse judgments about which ground-motion model they would use if they had to use only one (for a given magnitude, distance and frequency). However, when asked to assign weights to the range of models, the weighting schemes were remarkably similar.

SSHAC believes that it is very important, whatever process is used, not to force unwarranted technical consensus that appears to be agreement but that does not reflect the state of information of any single reasonable individual. The SSHAC process is oriented towards potential consensus of a very different sort, that is, consensus on the best composite representation of the knowledge of the scientific community.

3.3.2.2 Historical PSHA Approaches to Expert Use —Lessons Learned

In seismic hazard analysis, both mathematical and behavioral schemes have been used. The analysts typically decide at which level aggregation will take place (e.g., at the ground motion prediction level and/or at the overall seismic hazard level) and they employ mathematical combination formulas, either explicitly (e.g., equal or unequal weights on expert probability distributions), or implicitly (e.g., Monte Carlo sampling, implying equal weights, perhaps after removing an outlier, implying a zero weight).

Motivation for TFI Approach

The previous PSHA exercises most relevant to the multi-expert situation were the large EPRI and LLNL studies (Chapter 5 on Ground Motion describes relevant aspects of these studies). SSHAC was fortunate to have the extensive cooperation of project leaders and participants in both studies. They openly and willingly discussed the strengths and weaknesses of those projects; indeed, many of the key EPRI and LLNL participants made substantial contributions to this report. The successful ideas from these projects and other sources, such as (Otway and von Winterfeldt 1992); (Meyer and Booker 1991), (Cooke 1991) and (DeWispelare, Herre, Miklas and Clemen 1993), provided much of the foundation for the SSHAC TFI approach. However, detailed analysis of the previous studies also pointed to some areas for potential improvement, which led directly to the TFI concept:

- Overly Diffused Responsibility Previous studies sometimes lacked a well-defined single entity, responsible for the composite results. Responsibility was typically diffused over a large group of experts, analysts and stakeholders in a nebulous way. In contrast, the TFI has explicit overall responsibility for the final PSHA product. In all cases, of course, the individual experts are responsible for their own interpretations and evaluations.
- Insufficient Face-to-Face Expert Interaction Previous PSHA studies have sometimes not involved sufficient, nor sufficiently

structured, expert interaction. While experts were queried and interviewed, there was insufficient time for intensive face-to-face technical interaction among experts, and little opportunity for structured technical discussions to clarify issues, challenge and defend positions, and resolve unintended differences. Past projects demonstrate that the experts' limited experience with probabilistic models and statistics requires a strong facilitator to probe into technical details to avoid unintentional characterizations of expert interpretations, and also to detect and correct an expert who is acting as a proponent rather than an objective evaluator.

In the EPRI and LLNL studies, some experts felt dissociated from the final results. SSHAC's workshops confirmed that emphasis on expert interaction as the principal mechanism for integration helps to ensure that both the panel and the TFI feel "ownership" of the composite results —all major agreements and disagreements are represented explicitly.

- Inflexible Aggregation Schemes Previous PSHA studies (as well as most other major studies employing multiple-expert elicitation) have generally not taken sufficient advantage of state-of-the-art concepts and principles from the fields of expert elicitation and aggregation. For example, most PSHA studies have made *a priori* decisions to apply equal weights to multiple-expert judgments. As we shall discuss below, equal weights are both desirable and often appropriate, but only if the expert-interaction process is carefully designed to ensure appropriate conditions for equal weights, and only if a careful check is made after the interaction to ensure that these conditions have been met (in some prior studies, the conditions for equal weights were clearly not met). Although there are no universal algorithms or recipes for aggregating judgments, SSHAC's recommended process incorporates key principles from the large body of helpful research and practice in expert aggregation into guidelines for the TFI.

- Imprecise or Overly Narrow Objectives Previous PSHA studies (and, again, virtually all previous multiple-expert studies) have generally not distinguished well between the ultimate objective of a composite representation of the panel itself and a composite representation of the expert community as a whole. This distinction can be crucial with respect to important issues, such as how to deal with panel experts with outlier opinions. Fortuitously, representing the overall community is not only a more desirable objective, but is actually more likely to be an achievable one.

- Outlier Experts Previous PSHA and multiple-expert studies have dealt awkwardly or not at all with the contentious issue of "outlier experts," experts who make interpretations that are significantly different than the those of the rest of the panel and that are not well supported by logic or data. Treatment of outlier experts can have a major impact on the final hazard distribution; indeed, this issue was a primary motivation for the TFI process.

- Insufficient Feedback Following the elicitation of expert judgments—but prior to finalization of the assessments—the experts should be presented with the results of their evaluations. This feedback is in the form of the calculated final results, interim results, and numerous sensitivity analyses. As an example of past problems, it has been shown in some studies that the assessment of earthquake recurrence parameters (a and b-values) without feedback can lead to problems. This is because, in addition to uncertainties in the values of the two parameters individually, the parameters are usually correlated with each other and this correlation needs to be specified (e.g., the probability that high values of a should occur with high values of b). Feedback could include the range of recurrence curves derived from the assessment of a and b-values and derived recurrence intervals for particular magnitudes.

3.3.3 Underlying Logic of TFI Process

We summarize here the basic structure of the basic TFI process. Figure 3.1 provides a "roadmap" for the process 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 TFI's Fundamental Objective

To understand the "tree" TFI process and its potential outcomes, it is essential to understand the unique objective of the SSHAC process: that is, to use the panel to represent the overall scientific community's state of knowledge. The underlying premise is that the primary objective for public policy making is not capturing the judgment of any individual expert (including the TFI), nor even capturing the composite judgment of any specific subset of experts (including the panel), but rather, capturing as best possible the composite judgment of the overall scientific community of informed experts. Characterizing the panel's own knowledge is an essential intermediate goal, but not the final product.

Of course, it is impractical to engage an entire scientific community (often hundreds or even thousands of scientists for a given issue) in any meaningful interactive process. Decision makers must always rely on a smaller, but representative, set of experts. Thus, the panel is viewed as a sample of the overall expert community.

Section 3.3.4 describes a two-stage elicitation process in which the panel members are asked in Stage I to represent their own positions as independent evaluators of data, models and interpretations (the traditional role of a scientist); and in Stage II, to play the role of integrators who attempt to represent the composite position of the community as a whole. This two-role distinction may appear subtle at first, but it has important practical implications for the process outcomes.

The following discussion is organized around the tree in Figure 3-1. In describing the possible process outcomes, we highlight some especially useful working principles for the TFI to apply at each stage; the reader should be aware, however, that in application, the TFI can "mix and match" the principles throughout the process. Also, these and other process guidelines are described in more detail in Section 2 of Appendix J.

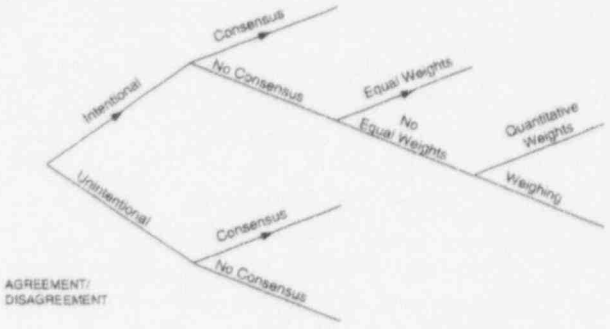

Figure 3-1 TFI Process Logic

Outcome 1: Consensus

The most desirable end state is consensus among the expert panel, but only if the experts truly agree after a full and intensive information exchange and interaction. There are two equally inappropriate outcomes the TFI must avoid: 1) the group achieving an artificial consensus that is not real (unintentional agreement) and 2) the group appearing to have substantial disagreements that are caused only by semantics and confusion rather than by substantive scientific differences (unintentional disagreement).

Types of Consensus

A key question we must address before proceeding is, "Consensus on what?" 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.

This could reflect agreement on a scalar parameter like the speed of light or density of the earth's crust, or agreement on a deterministic model and its parameters (e.g., ground motion attenuation as a function of distance), or sometimes just agreement on a functional form

(e.g., the attenuation curve is logarithmic). Importantly, this could reflect agreement within practical limits such that the final hazard distribution is insensitive to differences. This type of technical consensus represents the common use and meaning of the word, but is often an artificial objective and difficult to achieve.

Consensus Type 2:

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

This could reflect agreement about a probability judgment, the probability distribution resulting from a single model, or agreement on appropriate weights for a range of probabilistic models or positions. This type of technical consensus is also difficult to reach, but may be achievable for some issues after removal of unintentional differences by an appropriately facilitated TFI process (see below).

Consensus Type 3:

> All experts agree that a particular composite probability distribution represents them as a group.

Note that a group may agree on their composite representation, even if individuals have different positions. This type of consensus is generally easier to achieve than Types 1 and 2, especially if the experts recognize that substantial diversity among individual panel estimates tends to imply a wide range of overall uncertainty.

Consensus Type 4:

> All experts agree that a particular composite probability distribution represents the overall scientific community.

SSHAC seeks Type 4 consensus, which is potentially the easiest type of consensus to achieve. In the process of seeking Type 4 consensus, a useful intermediate step is to seek Type 3 consensus.

There is reason to be far more optimistic that the TFI process can achieve legitimate Type 3 or 4 representational consensus than one would be for an expert panel to achieve more traditional Type 1 or 2 technical consensus. In the TFI process, the issue is not consensus on scientific issues, which is almost impossible to achieve; acting as integrators, the experts only have to agree on the appropriate composite representation of the overall scientific community. As demonstrated in the SSHAC workshops, 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 seismicity or ground motion issue, than it is for the experts to agree on specific technical issues.

Here are some process principles especially useful in the early stages of the TFI process:

Experts as Evaluators, not Proponents Viewing the experts as evaluators (Stage I) who provide both interpretations of a range of data and models for the TFI 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, not for the purpose of creating a permanent advocate.

Emphasis on Expert Interaction The TFI must conduct structured, facilitated discussions among the panel 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 group interactions preceded and succeeded by TFI interaction with individual experts. Guidelines for how to conduct this interaction were developed and tested in the two SSHAC ground motion workshops (these are documented in Chapter 5 and Appendices A and B).

Isolate Sources of Disagreements Experts may disagree: about underlying scientific hypotheses and principles; about interpretations of different available data sets; about the values of model parameters; and, even with agreement on models, data and parameter values, about the ranges of the epistemic uncertainties that affect seismic hazard. Paradoxically, isolating and focusing discussion about the different potential types of disagreement

may actually move the group toward agreement on scientific issues. In the SSHAC workshops, the process of isolating sources of disagreement uncovered many common points of agreement and revealed a number of points of unintended disagreement. One participant remarked at the end, "It is astonishing how much everyone now agrees."

Active Listening A useful facilitation model is the concept of "active listening," in which a person's reasoning is not considered fully understood unless each listener, whether or not they agree with the reasoning, can explain it back to the person who made the point. It is extremely important for the TFI to summarize 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 as a facilitator to set the right tone. Two elements are critical: first, establish that the purpose is not to choose the best model or answer. The TFI concept is founded on the premise that there is no one correct model or answer, no single "winner" or "loser." Second, the purpose is not to achieve consensus (of any type, but especially Types 1 and 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.

Outcome 2: Equal Weights

When the panel members do not share the same composite representation of the community, the TFI must define the composite distribution. The TFI is neither constrained to use any fixed aggregation formula nor, in particular, to weight all expert inputs equally. Nevertheless, equal weighting has significant advantages and the TFI process is explicitly designed to create conditions under which equal weights will be appropriate. The attraction of equally weighting expert judgments is that it avoids at least two extremely difficult issues. First, one need not make what can be a very charged—and difficult to defend in the regulatory arena —judgment (Who is the best expert?), and second, one need not make what can be very difficult assessments (If not equal weights, what?)

It is essential for the TFI to understand clearly when equal weights are appropriate and when they are not. As we shall discuss, intensive interaction is perhaps the most effective way to create conditions under which equal weights are appropriate. In past seismic hazard and other public policy studies, equal weights were often used without this intensive interaction and without careful analysis of whether equal weights were appropriate. This can be dangerous in the seismic hazard arena: because of the logarithmic nature of key components of the seismic hazard calculation, equally weighting an indefensibly high probability given by one outlier expert can (as it has in some studies) swamp out the impact of all the other experts. The result is an answer that no one, not even the outlying expert, believes is representative of the overall community.

In the classic expert-use problem (see Appendix J for details), there are two fundamental conditions that must hold for equal weighting to be appropriate: first, the experts must either be completely independent —i.e., rely on independent data bases and models (this is virtually impossible), or be equally interdependent (expert dependence is more carefully defined in Appendix J). By exposing the expert panel to all models and data bases, the TFI process encourages equal interdependence. Second, the experts must be equally credible. In the TFI process, experts are methodically screened for their ability to be excellent scientific evaluators (see Section 3.3.4 for details on panel selection).

The Committee's methodological goal of representing the state of knowledge of the overall community of experts imposes another important condition that must be satisfied for equal weights to be appropriate. A set of n equally weighted experts, in order to represent the informed diversity in the whole community, must reflect an unbiased sample of the overall expert population. If, for example, an expert evaluator insists on giving weight to only one model, thereby acting as a proponent rather than an informed evaluator, giving that expert equal weight among the n experts overrepresents the strength of his or her position in the community. To understand this,

suppose that the group could be expanded to the size of the entire community by adding (n-1) new experts (i.e., the size of the community is (2n-1)). Then, the proponent would still be the only one holding his or her position, and weighting the experts equally would result in the appropriate weight for this position of 1/(2n-1). On the other hand, assuming that the (n-1) original representative expert positions are replicated by the new experts, equally weighting the non-proponent experts results in a weight of 2/(2n-1) for each of the unbiased positions, twice that of the proponent's weight. (Note that changing an individual expert's weight from 1/n may or may not change the composite representation significantly depending on how strongly hazard estimates based on his or her position deviate from the hazard estimate based on the composite distribution of the other n-1 experts.) Outlier experts are discussed further below.

Panel Selection and Removal For a Level 4 study, it is critically important to select a diverse group of experts, large enough to ensure that all credible points of view are represented, including all fundamental interpretations and modeling approaches. Using equal weights implicitly assumes that each expert is "standing in" for a much larger community of equally qualified experts. Thus, it is important that the set of experts be capable of representing the overall expert community as a whole.

Two of the most serious practical problems occur 1) when a expert behaves as a proponent, rather than an evaluator and 2) when an expert is not prepared and in some way does not live up to his or her professional time and work commitments. Careful panel selection using explicit selection criteria will greatly reduce the chance of encountering these problems. Nevertheless, SSHAC also recommends strongly that the TFI develop and discuss in advance with the panel formal criteria for dropping members from the panel (see Section 3.3.4 for more details). In the event of a problem, a determination is made by the TFI in close consultation and with the support of the overall panel.

Structure before numbers The focus in initial interactions should be on the logic of different basic approaches, rather than on variations of the same approach. There should be more dialogue at the level of structure than at the level of numbers. This avoids disagreement over small numerical issues local to the specific panel, and focuses on community—level issues that matter. As the interaction evolves, numbers become increasingly useful to the extent that they show how different modeling approaches work over ranges of applications and data and how much disagreements matter. A related lesson from the SSHAC workshops is that it is crucial to investigate early issues involving data underlying a model or its parameter values in order for the group to understand well the different model results and expert positions.

Sensitivity Analysis There is no reason to down-weight an expert's composite representation if the final answer is insensitive to the weight given to his or her position. If the expert's answer is not dramatically different than the average of the other positions, or if it results in a lower-than-average hazard probability (the hazard calculation is logarithmic), then it will likely not have an appreciable effect on the overall hazard calculation, especially the mean hazard curve. In this case, even if the TFI feels an expert's position is over-represented" by an equal weight, it is not worth the time, energy and possible controversy involved to down-weight that expert.

Outcome 3: Explicit Quantitative Weights

In any practical project, the number of experts (call it "n") is small relative to the larger population of equally qualified experts. If the TFI believes that if the panel were expanded to the size of the overall community, an expert's position would not be representative of 1/n of the community, then to give that expert's position weight 1/n would misrepresent the diversity in the overall community. In this case, unequal weighting may be appropriate. The situation need not be contentious and should be viewed as primarily a process issue. The relevant question is, "Is the expert's position, which is already a weighted combination of models, representative?" not the more personally threatening question, "Is the expert's scientific position correct?" The Committee believes that in the rare case in which

the representativeness issue arises, the expert should be given every opportunity to defend his or her position as being representative to the other experts and peer reviewers (especially participatory peer reviewers).

The issue of unequal weights is, of course, pertinent to the individual experts who will almost certainly want to give different weights to different models. In this case, the expert aggregation literature has some useful guidelines the TFI can pass on to the experts for how to determine these weights (Appendix J).

Outlier Experts 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: a) he or she makes an interpretation far different than the rest of the panel and b) the expert cannot support the interpretation with solid data or reasoning (from the points of view of the TFI and the other panel members). A past PSHA study provided an example of an expert who attached probability of unity to Modified Mercalli Intensity (MMI) XII earthquakes throughout the Northwestern 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" since the outlier expert does, in fact, represent 1/5 of the panel. Moreover, the outlier expert was selected carefully as being *a priori* as equally qualified as the other experts. Common sense says that the MMI XII expert should be down-weighted, but how can this be justified after the fact without superimposing the TFI's own scientific judgment on the process?

The perspective of developing a composite representation of the overall community of scientists affords a way out of the logical trap. When asked to identify other supportive experts, the outlier may even agree that he or she is the only one out of a hundred seismicity experts who would attach significant probability to a MMI XII earthquake. To represent the overall community, if we wish to treat the outlier's opinion as equally credible to the other panelists, we might properly assign a weight of 1/100 to his or her position, not 1/5.

Expert Aggregation Checklist Section 3 of Appendix J reviews a set of basic issues relevant to both expert aggregation (directly relevant to the TFI) and model aggregation (relevant to the TFI in guiding the experts as evaluators). The TFI should be aware of and carefully consider each aggregation issue at each stage of the process before final decisions are made concerning issues like equal or non-equal weights.

Outcome 4: "Weighing" rather than "Weighting"

Rarely, even after extensive interaction, will a situation call for some type of asymmetric treatment of expert-as-integrator representations. More commonly (but still relatively rare), the experts themselves, in their role as evaluators of models or proponent positions, may find simple fixed numerical weights to be inadequate. An example is in the ground motion arena in which many experts believe that the weights on different models should be a function of magnitude, frequency and distance (see the Ground Motion appendices). But there are even rarer situations in which explicit model weighting of any type is artificial, in which case an expert must "weigh" alternative models in a more general sense. A simple example will help to explain this concept. Two proponents have provided a TFI with their probability distributions on a scalar quantity y. These cumulative distribution functions (CDFs) are shown in Figure 3-2. The experts A and B have also supplied the reasoning (qualitative arguments) underlying their CDFs. If the TFI is constrained to use equal weights, he or she will do what the NUREG-1150 methodology required (Hora and Iman 1989) and will produce the curve labeled EW. For each value of y, the EW ordinate is one-half the sum of the ordinates of the curves A and B. The qualitative arguments that the experts have supplied play no role in this aggregation scheme, except, perhaps, to give legitimacy to the individual distributions.

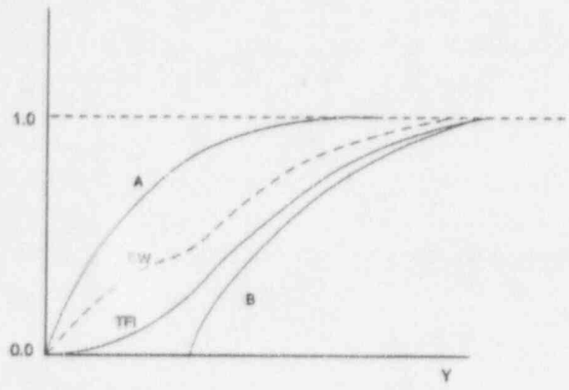

Figure 3-2 An Example of Behavioral Aggregation

Suppose now that the TFI studies these arguments carefully and finds that the reason why the two curves differ is the disputed applicability of a piece of evidence: Expert A believes that this evidence is convincing, while Expert B believes that it is not relevant. The experts are fully aware of this disagreement, and have discussed each other's rationales, but they are not willing to change their curves. Let us further assume that the TFI reaches the conclusion, based on the experts' interpretations, that the disputed evidence is most likely irrelevant at very low values of y, but cannot be completely dismissed for moderate values. The TFI, therefore, produces the curve labeled "TFI" to reflect this state of knowledge. This curve is presented to the experts and their subsequent arguments are evaluated by the TFI who may adjust the composite curve to reflect this feedback. Finally, the TFI reports the composite curve and the reasons that have led to its derivation (which, of course, includes reporting the individual curves and arguments, so that others may judge the validity of the whole exercise). This concept is consistent with Kaplan's idea of a "skillful user" (Kaplan 1992). It is easy to see why requiring the TFI to use explicit weights for this aggregation scheme would be artificial. Furthermore, this approach can mitigate contention based on different parties' complaints that their positions were not understood, because the explicit issues will have been explained and the TFI's reasoning documented, so that discussions on the merits can occur in an open context.

The Committee believes that while a "weighting" approach is not required of TFI's or TI's, explicit equal or unequal numerical weighting is highly desirable (if feasible) for several reasons: 1) Explicit weighting provides a decomposition in which different evaluations can be explicitly compared, 2) requiring explicit weights from experts tends to lower the possibility of eliciting extreme non-defensible opinions, and 3) there are probabilistic models (see Appendix J), albeit simplified, that provide theoretical underpinnings to the weighting process (as applied to either experts or models).

3.3.4 The TFI Process

We describe below a seven-step process for the TFI to follow to bring a multiple-expert project from problem definition to a successful conclusion. The seven steps are rather traditional, but some important novel aspects of the implementation are specific to the TFI process.

In particular, the goal of forming a composite representation of the scientific community suggests a natural two-stage elicitation procedure. We review this first because it provides useful context for not only the elicitation step, but for the expert selection, training, and aggregation steps as well.

Two-Stage Elicitation Procedure

A useful conceptual model of the expert panel is that it is an informed, independently-thinking sample of n evaluators who are representing a much larger community of N similarly informed evaluators (more precisely, representing the community's position if all in the community were equally informed, where "informed" includes a full understanding of relevant site - specific details). The TFI's problem is to collect information from the size n sample ($n < N$) in order to estimate the characteristics of the larger size N population. In many ways this is a classical problem in statistics, and many statistical insights apply directly. Section 4 of Appendix J presents a simplified mathematical version of this conceptual model.

The conceptual model suggests a two-stage process in which the expert panel members play two distinct roles. Here we highlight the elicitation process for each role. Appendix J provides for each stage a specific suggested list of the estimates and probability assessments required of the experts. Appendices A and B provide implementation details in the context of the two SSHAC ground-motion workshops.

Stage I Panelists as Independent, Informed Evaluators, Representing Themselves

Typically, the objects of a given elicitation are the parameters of an aleatory model, such as the mean rate or rupture velocity during an event or the median ground motion for a given distance and magnitude or even the (aleatory) standard deviation of the ground motion. The experts are asked to provide two types of assessments:

a) Each expert provides his or her best estimate (e.g., mean value). This is based on an evaluation of the full range of models, evidence, data and proponent positions in the community. The assessments are performed in the context of thorough facilitated interaction (including sharing of all relevant local or site-specific information) as described in Step 6 (analysis, aggregation, and resolution of disagreements).

b) Each expert assesses his or her epistemic uncertainty in the mean estimate. This is also based on thorough interaction; in particular, each expert is exposed to the full range of other panel-member estimates, which should often lead to appropriately wide distributions if there is substantial disagreement.

If the TFI's goal were to represent the panel's composite knowledge, the elicitation would stop here (after sufficient interaction, iteration, etc.). In fact, it is useful at this stage to construct an initial composite representation of the panel, but this is an intermediate product. A second stage builds additional information useful for extrapolating from the panel to the overall scientific community.

Stage II Panelists as Integrators, Representing the Overall Expert Community

In this stage, the panelists act as integrators (see Section 3.2), providing two types of assessments, based in large measure on what they learned from first-stage interactions with the other panel members:

a) Each expert provides an estimate of what the composite mean of the entire informed community would be; that is, assuming that an extensive elicitation were performed in which the community were provided the same information base and opportunity for interaction as the panel itself.

b) Each expert assesses an estimate of what the composite uncertainty in the community would be if an extensive elicitation were performed.

The Stage II assessments provide the TFI with information a) about each expert's judgment about how well his or her individual interpretation represents the overall community (it is entirely reasonable for a expert to say, "I recognize and can defend that my estimate is lower than the community average"), and b) about whether the panel believes its composite judgment is biased relative to the overall community.

The Stage II elicitation, since it is based largely on information generated in Stage I, should consume substantially less resources and time than the Stage I elicitation.

Seven-Step Process

The TFI must be involved in all aspects of a multiple-expert project in order to be able to take responsibility for the final product and to ensure that the involved experts take intellectual responsibility for the results. Based on their NUREG-1150 experience, Keeney & von Winterfeldt (1991) describe a seven-step process:

Step 1 Identification and selection of the technical questions

Step 2 Identification and selection of the experts

Step 3 Discussion and refinement of the issues

Step 4 Training for elicitation

Step 5 Group interaction and individual elicitation

Step 6 Analysis, aggregation, and resolution of disagreements

Step 7 Documentation and communication

Step 5 in Keeney and Winterfeldt's process was labeled "Elicitation." We have generalized the step to accommodate our special focus on group interaction.

Most of the discussion in the literature on multiple-expert applications, e.g., in Otway and von Winterfeldt (1992); Meyer and Booker (1991) and Cooke (1991), can be accommodated by this list of seven steps. In a project similar in spirit to the SSHAC project, DeWispelare and others (DeWispelare, Herre, Miklas and Clemen 1993) implemented an analogous formal expert elicitation process in their Yucca Mountain future-climate study.

We shall use the seven-step paradigm as a convenient way to structure our discussion of the TFI process; however, we pay special attention to the most unique SSHAC step, Step 6, where the TFI must act as both a facilitator for expert interactions (Step 6a) and as an integrator (Step 6b) responsible for producing a final composite representation of the expert panel.

Step 1. Identification and Selection of the Technical Issues

For our purpose here, a technical question is one that must be answered by the formal elicitation of expert judgments. Examples of questions from PSHA are the definition of the seismic source boundaries and the value of the maximum earthquake magnitude for each source in the seismicity portion of the study, and the median of the ground motion variable (PGA or spectral velocities) in the ground motion portion of the study. Clearly, such questions must have significant impacts on the results. Depending on the scope of the analysis and given the expense involved in a formal exercise, the TFI must develop criteria as to how the questions will be selected (relevant guidance is given in the seismicity and ground motion sections of this report). Some of the questions may be resolved by simply proposing a answer and soliciting comments from peers. The TFI should seek outside advice, e.g., from the study's sponsors and selected experts, when the questions are selected.

Step 2. Identification and Selection of the Experts

Attempting to define precisely who is an expert is not fruitful. In general, a candidate panelist must have a good professional reputation among his or her peers. In some recent studies (Trauth, Hora, and Guzowski 1993), a nomination process has been adopted, in which a long list of potential candidates is developed by consulting the archival literature and by asking technical societies, government organizations, as well as knowledgeable experts to submit names of researchers and practitioners. SSHAC strongly recommends this type of formal nomination process, and the development of a formal set of criteria for both selecting and potentially removing potential panel members.

For example, the following criteria were used to select the seismic source characterization experts for the ongoing Yucca Mountain seismic hazard analysis:

- Strong relevant expertise as demonstrated by professional reputation, academic training, relevant experience, and peer-reviewed publications and reports
- Willingness to forsake the role of proponent of any model, hypothesis or theory, and perform as an impartial expert who considers all hypotheses and theories and evaluates their relative credibility as determined by the data
- Availability and willingness to commit the time required to perform the evaluations needed to complete the study
- Specific knowledge of the Yucca Mountain area, the Basin and Range Province, or ground motion characterization
- Willingness to participate in a series of open workshops, diligently prepare required evaluations and interpretations, and openly explain and defend technical positions in

interactions with other experts participating in the project

- Personal attributes that include strong communication skills, interpersonal skills, flexibility and impartiality, and the ability to simplify and explain the basis for interpretations and technical positions.

In the same study, the following guidelines have been established for the removal of an expert from the panel:

"The need to consider removing an expert can only arise for failure to perform according to the commitments and demands of the project as stated in the expert selection criteria."

One or more of the following could prompt the need to consider removing an expert:

1) The person demonstrates unwillingness to perform as an expert evaluating credible models, hypotheses, or theories relative to the degree they are supported by data. This might be considered to be demonstrated if a person becomes a proponent of a single model, theory, or hypothesis to the exclusion of all others, or is unwilling to be guided by the data in making interpretations or expressing uncertainty.

2) The person is unwilling or finds it impossible to commit the time required to perform the evaluations needed to complete the study. This might be reflected in the person consistently being unprepared for workshops or interactive meetings with the Facilitation Team and/or consistently failing to meet established schedules for deliverables.

3) The person is unwilling to interact with other members of the project in an open and professional manner. This might be demonstrated by the person assuming a hostile and aggressive posture toward other members of the project or being uncooperative and disruptive in the workshops or interactions with the Facilitation Team."

A formal, well-documented selection and removal process can be extremely useful in highly charged political arenas in which the TFI must anticipate charges of bias. The TFI should play the principal role in creating nomination and removal criteria and in selecting the group, supported by the sponsors and possibly an advisory committee of experts.

It is important to ensure that the final group represents a broad spectrum of scientific expertise, technical points of view, and organizational representation. There are additional considerations as well. In the TFI process, evaluation ability and experience is especially important for the experts as informed evaluators. Also, the selection process should be influenced by the way the elicitation of the judgments will be handled. If the TFI plans to interact with the experts individually, it is important to select experts who are (or, are willing to become) somewhat familiar with the big picture, i.e., what PSHA is all about and how their input will be utilized. If, on the other hand, the TFI plans to form several teams of experts and interact with each team as a sub-group, then the concern should be making sure that each team includes all the necessary disciplines, e.g., for seismic source characterization issues, seismology, geophysics and geology. The need for each expert to have a broader perspective is not as pressing in the team case.

The advantage of forming teams is that, in highly multidisciplinary problems, each team can be chosen to have the necessary expertise to handle the problem. A drawback may be the presence of a strong personality who forces his or her judgment on the team, although an effective TFI will discern this and intervene to prevent it from happening. Furthermore, the presence of several teams provides additional assurance that a representative spectrum of scientific judgments will be obtained (i.e., assurance that the teams themselves can act as evaluators and integrators). In multidisciplinary problems, individual experts could have access to a supporting staff. Of course, the more elaborate the structure of the expert panels, the more costly the process. In the end, the

TFI will bear responsibility for both the selection process and the expert-panel structure.

Step 3. Discussion and Refinement of the Technical Issues

The TFI will hold a first meeting with the experts to discuss the technical questions that have been selected in Step 1 and to make sure that everyone understands them as intended (more meetings may be held, if necessary). The TFI needs to make sure that all experts have access to major sources of relevant data. An interaction of this kind is very important, because experience, for example, in the Ispra Benchmark Exercises on probabilistic assessments (Amendola 1986), has demonstrated that a major contributor to apparent disagreements is misinterpretation of the problem and its boundary conditions. Past experience in the LLNL/EPRI and other PSHA projects was similar. SSHAC workshops on seismic source characterization and ground motion confirmed that the participating experts felt strongly that detailed discussions and exchange of information prior to the actual elicitation were critical to the success of the exercise (see also the discussion on Step 6a below).

Through these interactions, the experts have an opportunity to provide input to the formulation of the technical questions and the precise formulation of the elicitation questions that will be asked. This formulation usually involves the decomposition of an issue into other issues that are judged to be easier to analyze. For example, one may wish to ask questions directly about a specific ground motion parameter or one may decide to consider several alternative models that estimate the parameter value, formulating the issues in the context of these models, i.e., asking questions about the numerical values of the parameters of these models, such as the expected stress drop.

The TFI's role in this step is primarily one of a technical facilitator (for more details on this subject, see Step 6a below). The TFI takes a proactive role by collecting and disseminating relevant information and by raising questions and encouraging all experts to learn the PSHA language and participate in the process. For example, this meeting offers a good opportunity for the TFI to discuss with the experts the concepts of aleatory and epistemic uncertainty (see Chapters 1 and 2 of this report). Such conceptual subtleties must be discussed so that the experts will have a clear understanding of the issues with which they are dealing (the Ground Motion appendices document such discussions).

Keeney and von Winterfeldt (Ref. 1991) recommend, and SSHAC agrees, that after the first meeting the experts should be given time to reflect on the issues and on the discussions that will have taken place at that meeting. They should then provide feedback to the TFI.

Besides the obvious benefits of eliminating misunderstandings, this step also influences the degree to which strong disagreements will surface during the processing of the judgments (Step 6, discussed below). We expect an informed group of experts that has debated the issues prior to the actual elicitation to be more likely to cooperate with the TFI in the formulation of the final composite judgment.

Step 4. Training for Elicitation

This step of the process is carried out by the elicitation experts of the TFI Team. The basic premise is that domain or substantive experts, i.e., experts on the relevant physical sciences, are not necessarily experienced at producing probability distributions that reflect their true state of knowledge. The language of probability may be foreign to them or they may be susceptible to various biases (Tversky and Kahneman 1974; Meyer and Booker 1991, Cooke 1991). Moreover, they should be familiarized with problem-structuring tools, such as influence diagrams (Shachter 1988; Oliver and Smith 1990; Call and Miller 1990) and logic trees (Coppersmith and Youngs 1986; National Research Council 1988).

The reluctance of some experts to speak in probabilistic terms may be overcome by explaining what probabilities are designed to do and by discussing some simple rules and exercises. The distinction between aleatory and epistemic uncertainty should be further explained in terms of concrete examples.

The possible biases may be characterized as being motivational or cognitive. Of course, the possibility of an expert having a motivation to distort his or her judgments deliberately should have been a factor in the selection of the experts (Step 2). Note that this does not necessarily mean that the TFI team should ignore candidates with motivational biases, just that these experts should properly play the role of proponents, not evaluators; in fact, the arguments that such proponents advance may be very useful to the panel's deliberations, even though the expert is known to be biased. The facilitation process described below is explicitly designed to expose and eliminate bias among panel members insofar as possible.

Cognitive biases, such as overconfidence and location bias, i.e., the reporting of narrower-than-justified probability distributions and the systematic over- or under-estimation of scalar quantities, have been discussed extensively in the cited literature. The TFI should explain to the experts the existence and nature of these biases in the hope that their impact will be minimized.

Step 5. Group Interaction and Individual Elicitation

An important aspect of the TFI process is the individual elicitation of probability judgments from individual experts. It is important to note, however, that the individual elicitations should be preceded by and followed by an important set of group interactions. We first address some individual elicitation issues and then discuss the relationship of individual elicitation to the group interactions.

We will not devote much space to individual elicitation here only because it is dealt with extensively elsewhere (including the references cited above). However, we do not want to minimize the importance of obtaining an accurate probability statement from each individual expert on all uncertainties of interest. Such a statement is useful, not only for characterizing each expert's position in a form usable for seismic hazard analysis, but also for ensuring full and unambiguous communication among the expert panel.

The actual elicitation process should be conducted with in-depth, face-to-face individual interviews, possibly supplemented by (but not replaced by) the use of preliminary questionnaires. When expert teams are employed, it is important to elicit the team as a group, possibly supplemented by preliminary individual interviews. The structure of the questions to be asked depends on the subject and will be developed by the TFI by taking into account the relevant literature.

A relevant point here is that the decision analysis literature advises that the experts should be asked to express opinions only on observable (at least in principle) quantities. In particular, this advice says that questions on event rates and moments of distributions should be avoided , because they are not "observable." Such a requirement would not allow the TFI to ask questions about the rate of occurrence of earthquakes in a seismic source, nor about the logarithmic standard deviation of the ground motion variable. This would be a mistake in the PSHA context, because the experts are very comfortable with these quantities. Asking the experts questions on "observable" quantities is based on the assumption that this would help them work with quantities that are easier to visualize and understand. In the earthquake community, long experience with data and analyses have made the experts very comfortable with the quantities cited above, so that related questions are meaningful to them.

An important element of the process, regardless of whether or not expert teams are formed, is the extensive use of consistency checks and providing feedback to the experts regarding the possible implications of their judgments. The idea is to challenge the experts and to invite self-scrutiny as much as possible. This is a key function of the TFI both as an informational resource to the expert group and as a facilitator of the group interactions and is discussed in detail in Step 6a below.

Before and after the individual elicitations, a number of types of group interactions need to take place. Chapters 4 and 5 present specific examples of types of workshops and meetings that enable these interactions. Here, we review briefly some generic interactions that are essential to success:

Information Meetings

There need to be informational meetings of at least three types (although not necessarily separated in time):

1. Background on objectives of study and overview of TFI process

   The experts need to understand the TFI process and their different roles in it. The experts must also understand clearly the distinction between the Stage I elicitation objectives and the Stage II elicitation objectives (as described above and expanded in Section 4 of Appendix J). In particular, assessing the possible scientific positions of the overall expert community will require a new way of thinking for most experts, so special care must be taken to ensure that the questions are well-defined, meaningful and thoroughly explained.

2. Background on the specific problem

   Depending on the scope of the study, the panel needs to be briefed by site or regional specialists who provide local or problem-specific knowledge that the panel members will not generally have. Also useful are presentations by local proponents and, possibly, site visits to give the panel first-hand familiarity with the study area. The experts should be encouraged to interact and exchange ideas and interpretations with the specialists.

3. Background on Hazard Analysis

   To be maximally effective, the experts must understand how their judgments will be used. They should be provided with a review of basic hazard methodology, the role of probabilistic judgments and the importance of sensitivity analysis.

Issue Interaction and Data Needs Review

The experts should work together to define and discuss the important issues on which uncertainty needs to be quantified—i.e., those variables that will require individual elicitation. Using the process described in Appendix J, the TFI structures interaction among panel members, specialists and proponents, facilitates debate and keeps the group focused on the sensitive parameters and issues.

It is also important to provide the experts with a detailed review of existing data and literature. The experts should be permitted to request additional data summaries and additional reports and papers.

Post-elicitation Feedback and Interaction

The TFI should summarize the result of the individual elicitations and provide this information as feedback to the entire panel. Panelists should be encouraged to amend their estimates, if they wish, after observing the other experts' judgments. Finally, it is often quite beneficial to conduct a post-elicitation group interaction to enable the experts to ask questions or address important differences or new issues arising out of the individual elicitation. Also, it is useful to structure group interaction to exchange viewpoints in preparation for individual expert-as-integrator assessments of the community distribution (Stage II) which must logically follow after the Stage I expert-as-evaluator assessments.

Step 6. Analysis, Aggregation, and Resolution of Disagreements

This step is where the SSHAC process deviates most from prior PSHA studies and the multiple-expert-use literature. Recall that the TFI has two fundamental roles: that of a Facilitator whose job it is to ensure that the knowledge, data and models of the expert community are fully and accurately elicited, and that of an Integrator whose job it is to ensure that the diverse information is integrated into a form useful for decision making that is a consistent and accurate representation of the state of information of the expert community. Because aggregation, if necessary, must follow the analysis of disagreements, it is natural to divide Step 6 into two successive steps: Step 6a, "The Role of TFI as a Facilitator," and Step 6b, "The Role of TFI as an Integrator."

Step 6a. The Role of TFI as a Facilitator

The TFI facilitation process is designed to encourage both the TFI and the experts to

understand explicitly the data bases and reasoning upon which different model estimates and expert interpretations are predicated. Moreover, it also demands explicit understanding concerning the rationale underlying each expert's uncertainty assessments.

SSHAC believes that successful integration is best achieved through proper facilitation of intensive interaction; hence, in the TFI process, the facilitation role of the TFI is paramount. A number of facilitation tips were provided in the previous sections. A longer list with more comprehensive discussion of facilitation principles and guidelines for potential TFIs is provided in Section 2 of Appendix J.

Step 6b. The Role of TFI as an Integrator

There are no cookbook formulas for integration (see Section 3 of Appendix J), but there are many useful concepts and models that can be used by the TFI. Even in the facilitation role, it is critical for the TFI be aware of certain key expert aggregation issues. Appendix J summarizes a set of fundamental expert-aggregation issues, including:

- Different Degrees of Expertise
- Outliers
- Non-Independent Experts
- Equal Weights
- Non-Equal Weights
- Level of Aggregation

The SSHAC process requires the TFI to be familiar with these issues and models, and to review them at each stage of the process (hence the need for an elicitation expert as part of the TFI team). There are three basic reasons for this:

1. The TFI must have a basic understanding of expert-aggregation issues in order to steer the expert interaction process to result in the simplest possible (e.g., equal weights) integration procedure. Moreover, the issues provide a checklist for the TFI to use in determining when it is appropriate to halt the process.
2. If it is determined that non-equal weights or "weighing" of the experts-as-integrators composite representations is the appropriate integration procedure, the aggregation issues and models provide useful information for how to do the non-equal weighting or weighing.
3. For experts acting as individual evaluators who must weight scientific models and interpretations, the aggregation issues and associated aggregation models can be directly useful. Since the experts are unlikely to be familiar with aggregation concepts, the TFI will need to use the aggregation issues and models to guide the experts in defining and assessing the weights.

We emphasize that the TFI does not need to use any prescribed, rigid combination formula, such as a fixed weighting scheme. Nevertheless, mathematical expert aggregation models have an important supporting role in the TFI process. A number of simplified expert-aggregation models are presented in Appendix J, Section 3. Also included is a new mathematical model specifically relevant to the TFI process. The TFI utilizes these models to check the implications of various assumptions, so that the ultimate aggregation (even if purely behavioral) will be sound and defensible. For example, the TFI may choose to process some disputed evidence using a number of aggregation models to illuminate the numerical impact of specific assumptions. This approach was used in Chibber, Apostolakis, and Okrent (1994) to estimate the pressure increment in the Sequoyah nuclear power plants containment vessel breach. The inputs from three experts, as reported in NUREG-1150 (Hora and Iman 1989), were processed using Bayesian methods under a number of assumptions regarding the degree of dependence among the experts, as well as the amount of their systematic biases.

Step 7. Documentation and Communication

The primary incentive for the formal elicitation of expert judgments is to supply credibility to the study. It is evident, therefore, that an essential element in accomplishing this is carefully and

thoroughly documenting every step of the process, as well as the results. It is important that each expert panel member document not only his or her own scientific position, but also his or her estimate of the community position. These detailed records will also prove invaluable when the TFI presents and defends the study to third parties, including regulatory agencies. Documentation is discussed in more detail in Chapter 7.

3.4 Peer Review

SSHAC recommends that peer review be conducted in both the TI and TFI processes. The purpose of the peer review is to provide assurance that a proper process has been followed, that the study incorporates the diversity of views prevailing within the technical community, that uncertainties have been properly considered and incorporated into the analysis, and the documentation of the study is clear and complete. Peer review has a long history of application in quality assurance for scientific endeavors including seismic hazard analysis. Classically, peer review is conducted by 1) one or more technical peers of the study participants who are "independent" of the study, and 2) at the end of the project. In recent years, experience on several large projects has shown that the active "participation" by peer reviewers throughout the course of the study can provide valuable input to the process being followed and can serve to define mid-course corrections that can improve the quality of the final product. This experience and these concepts are described in the guidance provided below.

3.4.1 Structuring the Peer-Review Process

If a PSHA project is to be successful, the crucial need for a strong peer review process cannot be overemphasized. What this means, in practice, is that the peer reviewers must be "peers" in the true sense: recognized experts on the subject matter under review. In the discussion below, we will assume that the Project Sponsor has assembled a peer-review panel, headed by a chairman who is responsible for writing the panel's reports (with the provision for the expression of minority views if appropriate). However, the Sponsor may in some cases use individual peer reviewers not assembled into a panel. For example, in a Level 1 analysis a review by a single peer reviewer may be sufficient to assure reasonable quality.

We will also assume that the peer-review panel reports are addressed to the Project Sponsor or the Project Leader, depending on the sponsor's desires, provided that the peer reviewers can act, and feel that they can act, to provide independent comments.

3.4.1.1 Participatory vs. Late-Stage Peer Review

In order to lay the foundation for our recommendations, we differentiate between two different types of peer review:

- A participatory peer review is an ongoing review that provides the peer reviewers with full and frequent access throughout the entire project. The process is structured to seek peer-review comments at numerous stages, and includes peer-review interaction with both the study team and, if appropriate, with the consultants and/or experts whose input is important to the final product. The principal benefit of a participatory peer review is that, if problems are discovered, the opportunity exists for a mid-course correction without the need for work to be substantially redone at the end. One limitation: peer reviewers might lose their objectivity as they interact with the project over time.
- A late-stage peer review is a review that occurs only after the project has been almost completed. Usually, such a review takes place when a draft of the final report has been prepared, or when the project's bottom-line results are close to being in final form. Sometimes, a late-stage peer review can examine an intermediate-stage result when it has been almost completed. The principal characteristic of a late-stage peer review is that, if major problems are discovered, the work may need to be substantially redone, without the mid-course-correction benefits of a participatory peer review. The use of a late-stage review is, therefore, a "gamble"—usually an informed gamble, of course—on the part of the sponsors that major problems

will not be discovered. A late-stage review has the benefit of a perception of complete independence.

Although these types of peer review are discussed separately here, it is possible for any given PSHA to include *both* a participatory and late-stage peer review.

3.4.1.2 Technical Peer Review vs. Process Peer Review

In the context of a PSHA project, we also need to distinguish between two different PSHA areas that require peer review:

- Technical peer review is the review of the earth-sciences aspects of a PSHA study: seismic-source characterizations, ground-motion models, the completeness and quality of the data set used to derive these inputs, etc. It also includes review of the PSHA calculation methods, the final seismic-hazard results and the sensitivity studies analyses. Reviewing this aspect requires expertise in the relevant earth sciences and calculational methodologies.
- Process peer review is the review of how the PSHA study is structured and executed. Because a PSHA must rely so heavily on expert interpretations of the admittedly inadequate earth-sciences information, the process peer review must concentrate on assuring that consideration of the uncertainties and the elicitation and incorporation of expert judgments is done well. Reviewing this aspect requires expertise in expert elicitation, statistical analysis, and related disciplines, as well as adequate familiarity with the technical issues and methods involved in a PSHA project.

3.4.2 Recommendations Concerning Peer Review

We have described two different methods for peer review, and two different subjects that require peer review:

- peer-review methods:
  - participatory peer review
  - late-stage peer review
- subject matter:
  - technical peer review
  - process peer review

We also have described two different approaches to address the complex technical issues involved in a PSHA project, the TFI and TI approaches. There are 4 different combinations of peer review structures to discuss for each of the two approaches. Table 3-2 contains a summary of our guidance concerning peer review.

Rationale: SSHAC's rationale for the peer-review guidance in Table 3-2 is as follows:

When structuring a peer review for the TFI approach, SSHAC recommends a participatory peer review over a late-stage peer review. When structuring a process peer review, SSHAC strongly cautions that a late-stage review can be very risky because accomplishing the process correctly is vital, and there are many process pitfalls that could benefit from a mid-course correction. In a technical peer review, SSHAC recommends a participatory review; however, this is not a strong recommendation—we believe that a late-stage technical peer review can be sufficiently effective, because the interactions among the various experts during the elicitation process, if done correctly, can provide many of the benefits of a participatory technical review.

When structuring a peer review for the TI approach, SSHAC believes that a participatory peer review is strongly recommended, if not essential. This recommendation holds for both the technical peer review and the process peer review. Although the process aspects using the TI approach may often be uncontroversial, SSHAC's reasoning is that, because the TI is conducting the entire analysis "in-house," there are significant opportunities for problems with both the technical and process aspects, and a late-stage review can be risky. For the technical aspects, the risk can sometimes be smaller (and more manageable) than for the process aspects, provided that the technical issues are not too contentious. For the process aspects, SSHAC believes that the risks associated with a late-stage review are likely to be great.

Table 3-2 SSHAC Recommendations on How to Structure the Peer Review Process

APPROACH	SUBJECT MATTER	METHOD	SSHAC RECOMMENDATION
TFI	Technical	Participatory	Recommended
		Late-stage	Can be acceptable
	Process	Participatory	Strongly recommended
		Late-stage	Risky: unlikely to be successful
TI	Technical	Participatory	Strongly recommended
		Late-stage	Risky but can be acceptable
	Process	Participatory	Strongly recommended
		Late-stage	Risky but can be acceptable

4. METHODOLOGY FOR CHARACTERIZING SEISMIC SOURCES

4.1 Introduction

This chapter summarizes important considerations in characterizing seismic sources for PSHA. Seismic Source Characterization (SSC) refers to the component of PSHA in which the locations, size, and frequency of future earthquakes are estimated. Because it is not yet possible to predict the location, size, and timing of the *next* earthquake, analysts attempt to determine the average *rate* of earthquake occurrence and use this rate as an indication of the likelihood or probability of future earthquake occurrence. The indication of rate, then, is a distinguishing feature of PSHA and a key parameter to be assessed for earthquake sources.

SSC is a multi-disciplinary activity that entails various aspects of the earth sciences including seismology, geology, and geophysics. The multi-disciplinary nature of source characterization means that a variety of expertise is required. Further, because of the limited knowledge of earthquake processes, the judgments of earth sciences experts (either formally or informally elicited) are required.

The three key elements of seismic source characterization are

- **Seismic source locations/geometries** Seismic sources are depicted in map form and represent locations within the earth's crust that have relatively uniform seismicity characteristics. Variations in the estimates of the geometries of sources reflect uncertainties in the spatial distribution of future seismicity. The probability of activity is assessed for each seismic source. Seismicity parameters (recurrence and Mmax) are specific to each seismic source.
- **Maximum earthquake magnitude** Maximum magnitudes (Mmax) are the largest magnitudes that a seismic source is capable of generating. Mmax is the upper-bound magnitude to the earthquake recurrence (frequency–magnitude) curve.
- **Earthquake recurrence** Earthquake recurrence is the frequency of occurrence of earthquakes having various magnitudes. Recurrence relationships or curves are developed for each seismic source and reflect the frequency of occurrence (usually expressed on an annual basis) of magnitudes up to the maximum.

The methods that are used to assess these three elements are different and, as a result, the three-part subdivision above will be used in the subsequent discussions of methodology.

The purpose of this chapter of the report is twofold: (1) to summarize the seismic source characteristics that are required for PSHA, and (2) to review approaches that can be used to characterize the epistemic uncertainties in SSC. These two sections of the chapter are not intended to be discussions of the "how-to" of seismic source characterization. The published literature provides reasonably complete discussions of the methods and scientific bases for characterizing sources for PSHA (e.g., Schwartz, 1988; Reiter, 1991; Coppersmith, 1991). These methods will be briefly summarized here. Likewise, various methods have been used to quantify the epistemic uncertainties in the elements of SSC and require only summary mention. Effort will be made, however, to distinguish among alternative methods for characterizing uncertainties, to recommend preferred approaches, and to note the pitfalls of these methods.

Section 4.4 of this report contains recommended methods for implementing SSC that incorporate expert judgment in quantifying uncertainties. The section is a principal focus of the SSC discussion because very little documentation of such methodologies exists in the literature. Further, it is the responsibility of SSHAC to review the methodologies and to make recommendations that are particularly appropriate to PSHA and its various components, including SSC.

A challenge in developing guidance for SSC is the requirement that the SSHAC-recommended

methods be appropriate for all parts of the United States. The approaches to source characterization, perhaps more than any other aspect of PSHA, depend upon the earthquake environment being considered. (Note that this is not strictly a function of "eastern" versus "western" U.S.; most of the western U.S. is characterized by low rates of seismicity, and some areas of the eastern United States are seismically active). In highly active areas of the western United States, the locations and geometries of seismic sources (in this case faults) are usually less uncertain than the recurrence rates appropriate for the sources; in turn, the recurrence rates are almost exclusively based on geologic data. Seismicity data play an important role in identifying sources and specifying the recurrence of small-magnitude events. In the low-activity eastern United States, geometries of seismic sources (typically area sources) are often highly uncertain and recurrence rates are derived almost exclusively from observed seismicity data, which are mostly small-magnitude earthquakes. Detailed analyses and procedures required for characterizing source geometries and recurrence, eastern United States versus western United States will not be enumerated; rather, this chapter will focus on methods for addressing the uncertainties associated with each and, in this way, find some common ground. The discussion of seismic sources is divided along the lines of various source types, as opposed to tectonic environments, which should assist in the application of the methods.

Section 4.2 summarizes the seismic source characteristics that are required for PSHA, and Section 4.3 discusses methods for characterizing epistemic uncertainties in SSC. Section 4.4 presents recommended methods for incorporating expert judgment in source characterization.

4.2 Seismic Source Characteristics Required for PSHA

The seismic source characteristics that must be assessed for probabilistic seismic hazard analysis are described below. The types of sources and the means of characterizing their earthquake behavior varies with the seismotectonic environment. Therefore it is useful to consider first the types of seismic sources that might be defined and then center the discussion on methods for these particular types of sources. Seismic sources can be categorized into four basic source types, shown in Figure 4-1:

Type 1 Faults, represented as lines or planes

Type 2 Area sources enclosing concentrated zones of seismicity

Type 3 Regional area sources

Type 4 Background area sources.

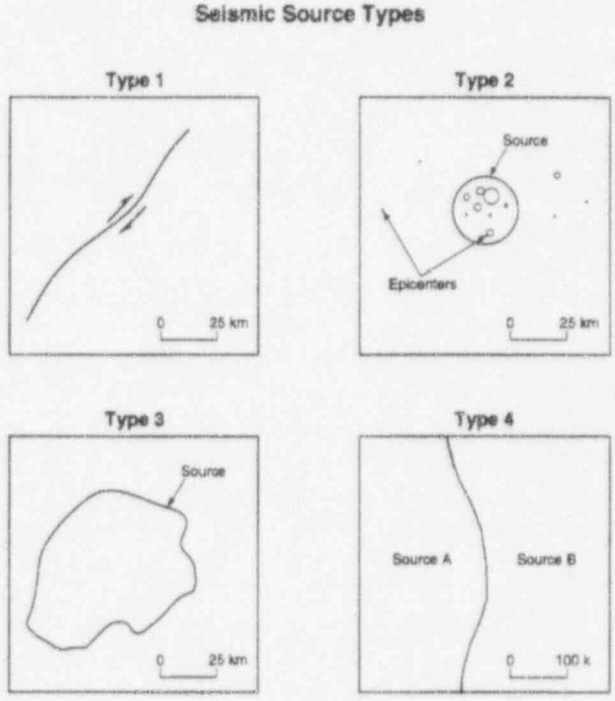

Figure 4-1 Diagrammatic representation of the four general types of seismic sources discussed in the text. Type 1 is a fault source and Types 2 - 4 are area sources. Type 2 is a source whose boundary encloses a zone of concentrated seismicity; Type 3 is a source defined by regional seismotectonic characteristics; and Type 4 is a regional background source (note scale).

Although these categories are arbitrary, they are useful in discussing the various data and methods used to characterize them. The basic source characteristics for all source types are the same (i.e., location, maximum magnitude, and recurrence); however the particular parameters and data sets that are used to define these

characteristics may be quite different. For example, slip rate is an important parameter for a fault source, but it is not applicable for a regional area source.

Although this section presents the source characteristics required for PSHA, it does not present a detailed description of the manner in which these characteristics can be assessed. For comprehensive descriptions of methods and the scientific basis for characterizing earthquake sources, refer to the published literature (e.g., Schwartz, 1988; Reiter, 1990; Coppersmith, 1991).

The following discussion is divided into the three principal components of seismic source characterization: source location and geometry, maximum earthquake magnitude, and earthquake recurrence. It should be recognized that, because of limitations in data, it will not be possible to assess all of the characteristics described below as part of any given seismic hazard analysis. For example, paleoseismic data may not be available to evaluate recurrence rates for a particular seismic source. However, the discussion here is given in terms of a reasonably complete set of alternative approaches. It is recognized that other characteristics besides those discussed are likely important to earthquake ground motions (for example, dynamic stress drop and the coseismic distribution of slip on a fault). However, these characteristics are not yet commonly included (at least explicitly) in probabilistic seismic hazard analysis.

4.2.1 Seismic Source Locations and Geometries

A seismic source is a construct developed for seismic hazard analysis as a means of approximating the locations of earthquake occurrences. A seismic source is defined as a region of the earth's crust that have relatively uniform seismicity characteristics, and is distinct from those of neighboring sources. It is possible to allow for some variation of seismicity parameters (a- and b-values) within a given seismic source. Typically, however, the distribution of Mmax and the probability of activity (defined below) are assumed to be uniform within a seismic source.

Each seismic source must be defined by its location in order for the distance distribution to a site of interest to be calculated in the hazard analysis. In theory, the level of detail necessary to describe the location and geometry of sources can be uniform for large regions. In practice, however, the level of detail in specifying the locations and geometries of seismic sources can vary as a function of distance from the site. Because the amplitude of ground motions attenuates with distance from the source, at large distances even large-magnitude earthquakes will not result in significant ground motions at the site. From the standpoint of seismic hazard analysis, this means that the inclusion of these distant sources in the analysis is not required because they do not contribute to site ground motions.

This means that there are distances beyond which detailed source characterization is not necessary. For example, for a site in the western United States (with its attendant attenuation), it is likely that sources more than about 300 km from a site of interest do not need to be considered; "detailed" source characterization need only be carried out within, say, 100 km from the site. "Detailed" source characterization would include specifying the mapped location and three-dimensional geometries of faults. At greater distances, the effect on hazard from faults and area sources is similar. Thus faults and small area sources at larger distances can usually be generalized as large area sources. Further, if fault sources (or Type 2 localized area sources) are nearby they will likely be most important to the hazard results and will, therefore, preclude the need to characterize sources in detail out to large distances.

To provide guidance on this issue, the following source-to-site distances are suggested for detailed source characterization and source identification, as a function of whether or not nearby faults are present:

Western U.S.	Source-to-Site Distance (km)
Maximum distance for source identification	300
Distance for detailed source characterization	
Faults within 50 km of site	100
No faults within 50 km of site	150
Eastern U.S.	**Source-to-Site Distance (km)**
Maximum distance for source identification	500*
Distance for detailed source characterization	
Faults within 50 km of site	200
No faults within 50 km of site	300

* In certain cases, where a highly active distant source is present, capable of generating large-magnitude earthquakes (e.g., New Madrid), distances up to 1,000 km may need to be considered.

For example, for a site in the eastern United States that has faults (or localized sources) within 50 km of the site, seismic sources should be characterized out to distances of about 200 km of the site. The difference between the western and eastern U.S. is related to differences in the ground motion attenuation between the two regions.

The "western U.S." is defined roughly as the region of Mesozoic–Cenozoic deformation of the earth's crust lying west of the Rocky Mountain front. The definition of locations and geometries varies with source type between faults (type 1) and area sources (types 2–4), as discussed below.

Fault Locations and Geometries (Source Type 1)

At a minimum, the location of fault sources must be identified in map view. Usually a fault map depicts the line of intersection of faults with the ground surface. In the case of blind faults that do not intersect the surface, the location of the shallowest extent of the fault should be indicated on the fault maps. With the occurrence of the 1983 Coalinga earthquake and the 1994 Northridge earthquake has come an increasing recognition of the important contribution that blind or buried faults can make to seismic hazard, particularly within regions of compressional tectonics.

Faults may be represented as "line" sources using the fault maps or, if sufficient information is available, by three-dimensional fault planes. The need to characterize the three-dimensional geometry of a source is greatest where the source-to-site distance is small. For example, if a fault is less than 10 km from a site, the direction and amount of dip away from or toward the site can have a large impact on the source-to-site distance.

A primary geometric characteristic is the dip angle, expressed by convention as 90 degrees for vertical faults and decreasing as the fault approaches the horizontal. The direction of dip must also be specified.

The updip and downdip extent of the fault within the seismogenic crust must also be specified for three-dimensional faults. Because seismic hazard analysis attempts to portray the earthquake generation process, a three-dimensional rupture is assumed to occur during earthquake generation. The area of this rupture, as measured on the fault surface in square kilometers, is directly proportional to earthquake magnitude. Empirical relationships, such as that given by Wyss (1979) and Wells and Coppersmith (1994) describe the area of rupture for given magnitudes. In order to model the occurrence of earthquake ruptures for hazard analysis, an estimate must be made of the downdip extent of the fault within the

seismogenic part of the crust. Such an estimate is commonly developed by considering the maximum focal depths of seismicity in the vicinity of the fault or in the region.

Another characteristic of faults that must be assessed is the style of faulting, generally defined as strike-slip, normal, and reverse faulting. This assessment can come from geologic studies of the fault, focal mechanisms from associated seismicity, or tectonic considerations (Coppersmith, 1991).

Area Source Locations and Geometries (Source Types 2–4)

It is universally true that earthquakes are the result of differential slip on faults. However, in many areas, such as most of the eastern U.S., the identification of the causative faults giving rise to seismicity is problematic. To accommodate this uncertainty in fault location, area sources were invented and have common application in PSHA. It is recognized that a homogeneous area source used in PSHA is not a physical characteristic of the earth's crust but is a simplified representation of one or more seismogenic structures whose location is unknown. The area-source boundaries enclose regions that earth scientists believe are relatively uniform with respect to the PSHA application.

Although the data used in their identification can be significantly different, the depiction of area sources is essentially the same for all source types discussed. Seismic sources are defined by their boundaries shown on maps. Although these boundaries may be considered "fuzzy" boundaries (Bender 1986), most commonly they are assumed to be sharp and to define differences in the maximum magnitude and recurrence rate between one zone and another. (An exception is variation in recurrence parameters within an area source). As discussed previously, area source boundaries can be defined by a variety of characteristics including concentrations of seismicity, changes in tectonics, and geologic boundaries.

Although the source map is the only required product, an assessment be made of the depth distribution of seismicity (which defines a seismogenic volume) is also recommended, particularly if the depth is anomalous relative to other regions. Also, the expected style of faulting should be evaluated. Uncertainties in source boundaries are incorporated into the hazard analysis through the identification of alternative source configurations, each with its own relative weight or credibility.

Data Used to Define Source Locations and Geometries

The identification of seismic sources is a critical part of seismic hazard analysis and involves a wide range of data types and scientific interpretations. The purpose of this section is to identify the types of data that can be used to develop source interpretations and to provide an indication of the relative usefulness that various types of data may have in making source assessments. No requirement is being made that all data discussed be developed for all hazard analyses—some hazard studies may require more data than others depending on the scope of the analysis. It is a requirement, however, that all *available* data of the type indicated be considered in characterizing seismic sources. Gathering additional data is a function of their importance to the analysis, potential benefits to be gained from further reducing uncertainties, and the like.

Table 4-1 summarizes the types of data used to define each of the four types of seismic sources and the relative usefulness of each data type. Relative usefulness in this context means the degree to which that particular type of data provides a strong technical basis for the source definition. For example, if fault sources are being identified, a map of young (Quaternary) faults is judged to provide a strong basis for defining fault sources in hazard assessment, whereas a map of older (pre-Quaternary) faults is judged to provide a relatively weak basis for defining fault sources. Likewise, the nature and spatial patterns of instrumental seismicity are most important in defining Type 2 and 3 area sources, while various types of geological structural data play a lesser role. Note that, in real application, the quality of various data can vary significantly. This variation can have an important impact on its usefulness in source definition.

Table 4-1 Data Used to Assess Seismic Source Locations and Geometries and Their Relative Usefulness

TYPE OF SOURCE	DATA/BASIS FOR SOURCE	RELATIVE USEFULNESS/ CREDIBILITY (1: high, 3: low)
Type 1: Faults	Mapped fault with historical rupture	1
	Mapped Quaternary fault at surface	1
	Mapped localized Quaternary deformation, inferred fault at depth	2
	Borehole evidence for fault, especially in young units	2
	Geophysical evidence (e.g. seismic reflection) of fault at depth	2
	Map of pre-Quaternary faults	3
Type 2: Concentrated Zone	Concentrated zone of well-located instrumental seismicity	1
	Mapped fault(s) at surface or subsurface in proximity to seismicity	1
	Zone of historical/poorly located seismicity	2
	Structural features/trends parallel to seismicity zone	2
	Focal mechanisms/stress orientation	3
	Rapid lateral changes in structures/tectonic features	3
Type 3: Regional Zone	Changes in spatial distribution/concentration/density of seismicity	1
	Regions of genetically-related tectonic history	1
	Regions of similar structural styles	2
	Changes in crustal thickness or crustal composition	2
	Regions of different geophysical signature	3
	Changes in regional stresses	3
	Changes in regional physiography	3
Type 4: Background Zones	Regional differences in structural styles/tectonic history	1
	Major physiographic/geologic provinces	1
	Changes in character of seismicity	3

4.2.2 Maximum Earthquake Magnitudes

The maximum earthquake magnitude that a seismic source is capable of generating defines the upper bound to the earthquake recurrence relationship. Because the assessment of the maximum magnitude often includes approaches different from those used to evaluate the remainder of the recurrence relationship, maximum magnitudes and earthquake recurrence assessments are discussed separately.

Faults (Source Type 1)

There are two basic approaches to assessing maximum magnitudes for fault sources: constraints provided by historical seismicity and provided by estimates of maximum dimensions of rupture. In most cases, the historical record for individual faults is short relative to recurrence intervals for the largest earthquakes; thus the probability that the historical record includes the maximum event is usually small. However, if the historical record includes a significant earthquake that can be associated with the fault (say, a surface-rupturing event such as the 1857 earthquake on the San Andreas fault), it may provide an estimate of the maximum magnitude. In cases where the historical event was associated with coseismic rupture, the extent of that rupture can be evaluated geologically relative to other constraints on the maximum rupture dimensions.

Earthquake magnitude is well-correlated with rupture dimensions. It follows that if rupture dimensions associated with a maximum earthquake on a fault can be estimated, the maximum magnitude can be assessed. Fault rupture parameters that have been shown empirically to be correlated with earthquake magnitude include rupture length, rupture area, maximum surface displacement, and average surface displacement (Slemmons, 1977; Bonilla and others, 1984; Wells and Coppersmith, 1994). The evaluation of these parameters for an individual fault includes paleoseismic investigations of the extent of past ruptures and other geologic constraints (see discussions in Schwartz and Coppersmith, 1988; Schwartz, 1989; Coppersmith, 1991). Commonly, a number of potential rupture dimensions can be estimated (e.g., rupture length, rupture area, displacement per event) and a magnitude estimated for each. Paleoseismic data regarding the number of events and rupture dimensions are usually associated with considerable uncertainty. The final maximum magnitude estimate for a fault source should be a distribution of magnitude values. The distribution should reflect the uncertainties in the estimates of rupture dimensions and their relative credibilities. Any constraints provided by the historical seismicity record can also be included in the maximum magnitude distribution.

Area Sources (Source Types 2–4)

The assessment of maximum earthquake magnitudes for area sources is particularly difficult because the physical constraint most important to the assessment—the dimensions of fault rupture—is not known. As a result, the primary methods for assessing maximum earthquakes for area sources usually include a consideration of the historical seismicity record and analogies to other sources.

In assessing the maximum earthquake, the historical seismicity record takes on great importance—particularly in terms of the locations and sizes of older earthquakes. Extensive studies of the distribution of intensities, and relationships between isoseismal distributions and magnitude, have been initiated with the ultimate goal of using them in evaluating the size and location of older events.

Studies of the sizes of historical earthquakes associated with the area source of interest should be made. It is possible that, after the historical record has been examined, it will be concluded that the record provides no particular constraint on the estimate of maximum earthquake for the source. Alternatively, the maximum historical earthquake for the zone may be assessed as a lower bound or best estimate of the maximum magnitude for the source. In cases where the maximum historical earthquake has not been assessed to be equivalent to the maximum possible earthquake, past practice has included adding an increment of one-half magnitude unit or one intensity unit to the maximum historical earthquake. This practice implies that, because the

historical record does not include the maximum event, the recurrence interval for the maximum possible event is longer than the historical period. Thus, the addition of a magnitude unit is equivalent to a shift to longer recurrence intervals on the recurrence relationship for the source (an approximate recurrence interval of 10 times the historical record, for typical b-values).

Other considerations in assessing maximum earthquakes for area sources are analogies to other sources. The source of interest may be tectonically similar to another source such that their maximum earthquakes are also deemed to be similar. For example, in past practice in the eastern U.S., the tectonic association of certain large-magnitude historical earthquakes, such as the 1886 Charleston earthquake, was evaluated relative to the possibility that such an earthquake could occur in other sources having similar tectonic characteristics. At present, the tectonic characteristics that are most important to controlling maximum earthquakes are not well-known, but could include whether or not the source is characterized by past rifting or extension (Johnston and Kanter, 1990). Recently completed studies (EPRI, 1993) have examined the possible tectonic constraints on maximum earthquakes within sources in stable continental regions.

In some cases, it may be possible to incorporate considerations of possible rupture dimensions into assessments of maximum magnitudes for area sources. The lengths of zones of concentrated seismicity (source type 2) may be assessed to represent maximum lengths of rupture. The dimensions of tectonic elements within a source may also provide physical constraints on maximum earthquakes. For example, a source that is defined as including a region of crustal deformation may include a consideration of the dimensions of faults within the deformation zone.

The uncertainties associated with the assessment of maximum earthquake magnitudes for area sources must be incorporated into a probability distribution for each source. The technical basis for the assessment and the associated data must be fully documented.

4.2.3 Earthquake Recurrence

Earthquake recurrence relationships express the annual frequency (which is usually assumed to be constant in time) of earthquakes having various magnitudes up to the maximum magnitude and they must be developed for each seismic source. The methods for developing these relationships are usually different for fault sources than for area sources.

Faults (Source Type 1)

The development of recurrence relationships for fault sources can include information from both the historical seismicity record and the geologic record. Typically, observed seismicity provides constraints on the frequency of small-magnitude events and the slope of the recurrence curve; the geologic record provides the frequency of larger-magnitude events.

To use observed seismicity to estimate earthquake recurrence first requires that an assessment be made of which events can be associated with the fault of interest. For instrumentally recorded earthquakes, a corridor around the fault should be specified that accounts for the dip of the fault and the epicentral location uncertainties. Associations with older historical earthquakes must consider uncertainties in epicentral locations.

The use of observed seismicity for recurrence assessment, for either faults or areal source zones, must account for incompleteness in the catalog as a function of magnitude, location, and time. The recurrence rate that is needed for seismic hazard analysis is the rate of independent main shocks, which are typically assumed to be distributed randomly in time. Therefore, dependent events (foreshocks, aftershocks, clusters) must be removed for use in the hazard analysis.

In plotting recurrence from observed seismicity (for example, Figure 4-2), it is helpful to indicate the average or mean frequency at particular magnitudes as well as to indicate the statistical variability of the frequency estimate for that magnitude (e.g., Weichert, 1980). Such a plot, expressed for example with 5- and 95-percent confidence limits, typically shows the progressively larger errors with increasing

magnitude. This is directly due to the occurrance of progressively fewer events as the magnitude increases.

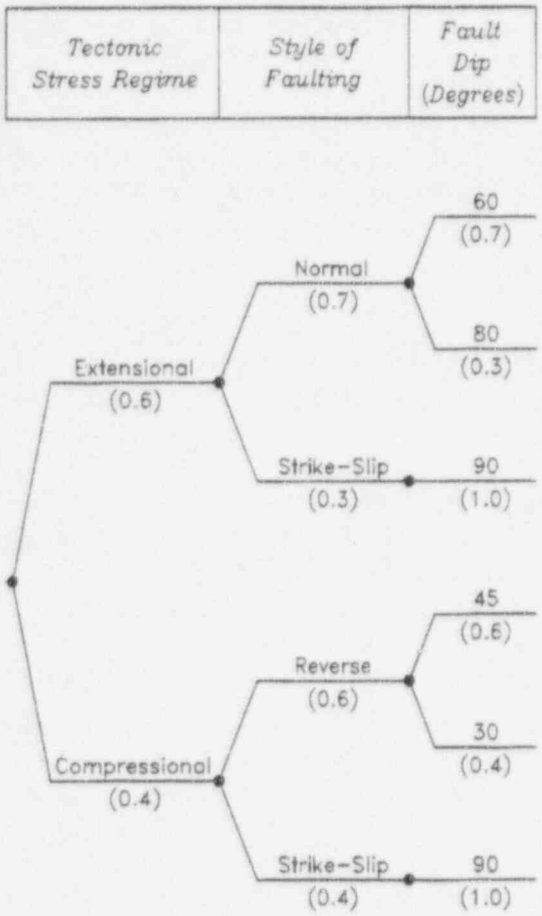

Figure 4-2 Example logic tree illustrating the manner in which assessments of the tectonic model can affect assessments of seismic source characteristics such as source geometry. In the example, the assessment of the tectonic model (in this case the nature of the regional stress regime) affects the assessments of the expected style of faulting and, in turn, the dip of faults.

Geologic data often provide valuable information regarding the recurrence of larger-magnitude earthquakes. Paleoseismic data can provide assessments of the recurrence intervals associated with earthquakes that have ruptured the surface. In using paleoseismic data, the uncertainties the recurrence intervals and the magnitudes of the paleoseismic events should be included. Another type of geologic constraint on earthquake recurrence is provided by the fault slip rate. The slip rate can provide an estimate of the average rate of release of seismic strain energy. (See Coppersmith, 1991 for discussion of the use of slip rate for recurrence estimation.) To use the slip rate, it must be partitioned into various earthquake magnitudes according to a magnitude-distribution model. Two alternative models are the truncated exponential model and the characteristic earthquake model (Youngs and Coppersmith, 1985). The characteristic earthquake model appears to be more valid for describing the recurrence behavior of individual faults (e.g., Schwartz and Coppersmith, 1984; Youngs and others, 1993). Both the exponential and the characteristic earthquake model require an estimate of the b-value in the exponentially distributed part of the recurrence curve. This estimate is commonly derived from the average b-value in the region based on observed seismicity. The uncertainties in slip rates and magnitude-distribution models should be incorporated and documented.

A suggested representation of earthquake recurrence relationships for individual faults is to indicate the frequency of observed earthquakes, with associated statistical error bars, the recurrence intervals from paleoseismic data, and the mean recurrence curves derived from the slip rate and magnitude-distribution model (Figure 4-2).

Area Sources (Source Types 2–4)

The assessment of earthquake recurrence for area sources commonly relies heavily on catalogs of observed seismicity. To maximize their utility, seismicity catalogs should be reviewed for uniformity in designation of magnitudes and for completeness as a function of magnitude, location, and time. The association of older historical events with particular seismic sources should be assessed bearing in mind the location uncertainties. For example, whether a large-magnitude historical earthquake, such as the 1886 Charleston earthquake, occurred in one source or another may be important to estimates of recurrence within those sources.

The observed seismicity rates can be plotted as mean frequencies for each magnitude, along with the statistical uncertainties due to the number of

events within each magnitude bin (e.g., Figure 4-2). Using these observed data, with the maximum magnitude estimate, a recurrence curve is fit. A reasonable method for curve-fitting is the maximum likelihood method because it accounts for the decreasing number of points as magnitude increases. The result is a recurrence curve that expresses the recurrence rate for various-magnitude earthquakes up to the maximum. Various methods for expressing the uncertainty in recurrence curves are discussed in Section 4.3. An appropriate magnitude-distribution model for area sources is a truncated exponential distribution.

The degree of variation, or "smoothing," of the a- and b-values within an area source can be specified. Uniform a and b throughout the source represents maximum smoothing, and different levels of smoothing can be identified. Guidance on the use of spatially varying recurrence parameters within a seismic source is given in Section 4.3.5.

4.3 Characterizing Epistemic Uncertainties in Seismic Source Characterization

Section 4.2 presented the basic elements of seismic source characterization that are required for PSHA. All of the elements discussed are uncertain and this epistemic uncertainty can be addressed in a variety of ways. In this section, approaches to characterizing the uncertainties in SSC are discussed.

4.3.1 Seismic Source Location and Geometry

Two basic approaches have been commonly applied in characterizing the uncertainties in source location and in specifying the activity of sources: alternative maps of seismic sources each associated with a relative weight, and alternative configurations of a seismic source each associated with a relative weight or probability of activity. Both of these approaches are acceptable and the preference for one or the other depends upon the SSC expert.

Probability of activity is an expression of the likelihood that a particular seismic source is seismogenic or capable of generating significant earthquakes. This assessment is most commonly made for individual faults, but has also found application (for example, in the EPRI eastern United States study) in assessing particular area sources interpreted on the basis of various tectonic features. In many cases, there may be uncertainty regarding whether or not seismic sources shown on source maps are active. Hence an assessment of the probability of activity must be made. An equivalent assessment is the probability of existence of a particular source zone.

The activity of fault sources is commonly assessed using the criteria developed from regulatory experience. For example, the concept of fault "capability," which is given in NRC's geologic siting criteria for nuclear power plants (10 CFR Part 100, Appendix A), is a common basis for assessing activity of faults. Fault activity assessment usually involves criteria that are believed to provide an indication of the potential for future earthquake occurrence. Such criteria include spatial association with past earthquakes, evidence for geologically recent displacement, structural association with other active faults, and the like. The relative usefulness of these various criteria is often quite different and should be identified.

The probability of activity of source zones has been evaluated in two alternative, equally credible, ways in the EPRI and LLNL studies for the eastern U.S. In the EPRI approach, tectonic features that might be seismogenic were identified and their probability of activity assessed. The criteria for assessing the activity of a feature are first identified and defined. Criteria include such attributes as spatial association with large- or small-magnitude earthquakes, evidence of geologically recent slip, orientation relative to the regional stress regime, and the like. The relative weight or relative value of each criterion in assessing the probability of activity is evaluated generically in a "tectonic feature matrix." Then these criteria are applied to each feature to assess its probability of activity. The seismic sources interpreted from the tectonic features (i.e., "feature-specific source zones") are then assigned

a probability of activity equivalent to that of the feature.

In the LLNL study, the probability of activity is defined as the probability of "existence" of a particular source zone. In practice, rather than making the assessment on a source-by-source basis, alternative source maps are developed—each map having its own probability of existence or credibility. The hazard calculations include this probability in combining the alternative maps. The probability of activity/existence expresses the uncertainties in the locations and geometries of seismic sources for the PSHA. In all applications of the probability of activity or existence, the criteria for making the assessment must be documented, the relative value of the criteria must be evaluated, and the basis for the assessments must be documented.

In expressing the probability of activity it is important to specify clearly the criteria that are being used to evaluate the activity and the relative weight that the criteria have in the evaluation. In the EPRI procedure, the criteria and their relative weight were specified using a "tectonic feature matrix" and were used to evaluate a large number of features. In addition, dependencies among sources may need to be indicated. In some cases, for example, one interpretation of the configuration of a seismic source may be judged to be mutually exclusive with another configuration, or one interpretation may be judged to depend on other interpretations. In these cases, additional assessments that describe these dependencies need to be made in order to properly combine all of the sources in the seismic hazard analysis (EPRI, 1989).

Another way in which tectonic interpretations are linked with seismic source geometries is through considerations of tectonic models. Alternative tectonic models for a region may imply different source geometries. For example, alternative tectonic models for a region may imply that mapped faults are either high-angle strike-slip faults or low-angle thrust faults. The uncertainty in tectonic models should be treated first in terms of alternative models, each with its relative weight. Then the alternative source geometries that are implied by these models can be developed as a function of each particular tectonic model. The resulting alternative tectonic models can be summarized in a logic tree format (see example in Figure 4-2).

Uncertainties in all of the parameters defining the geometry of individual sources can be characterized using weighted alternative parameter values or estimated continuous distributions. These parameters include maximum depth of seismogenic crust, focal depth distribution, fault dip angle and direction, total fault length, and updip and downdip extent (for blind faults).

4.3.2 Maximum Earthquake Magnitude

Maximum earthquake magnitude is a parameter for each seismic source. As such, its uncertainty can be defined by discrete alternative values with relative weights or using a continuous probability distribution. In addition to direct assessments of Mmax, it is also common to display and incorporate the uncertainties in the parameters and models that were used to derive the maximum earthquake as well. For example, maximum magnitudes for fault sources are typically estimated based on estimated maximum dimensions of rupture, including maximum surface rupture length, subsurface rupture length, maximum displacement, and average displacement. These rupture dimensions are, in turn, empirically related to earthquake magnitude. For a given fault having data related to each of these dimensions, it may be useful to express the relative weight to be given to each of the them in assessing the maximum earthquake. In addition, if multiple segmentation models are used to estimate rupture length, these models should each be associated with a relative weight. Clearly, a logic structure is a convenient way to express the relative weights applied to various approaches and parameters used to assess the maximum magnitude. Using a logic tree format, the maximum magnitude distribution for the source is simply a probability distribution of the type shown in Figure 4-3. The discrete Mmax distribution can be used directly in the seismic hazard analysis.

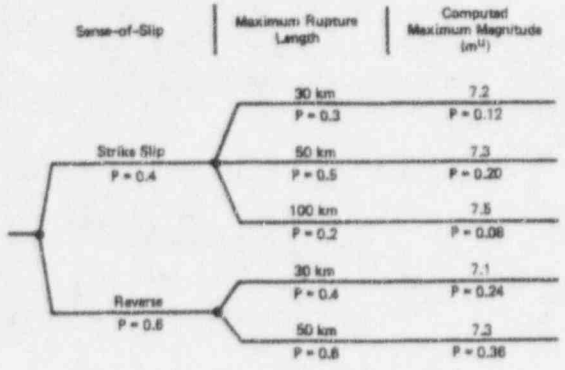

a) Logic Tree of Evaluating Maximum Magnitude

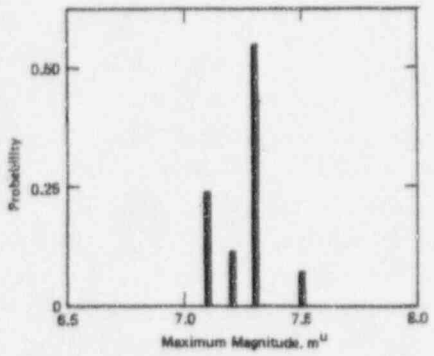

b) Discrete Distribution for Maximum Magnitude

Figure 4-3 Example logic tree showing the manner in which assessments of fault rupture dimensions, and associated uncertainties, leads to a probabilistic distribution of maximum magnitude. In the example, the sense of slip on the fault is uncertain and the expected maximum displacement per event is assessed conditional on the sense of slip. Each displacement value is related empirically to earthquake magnitude. The probability associated with each magnitude on the end branches is the product of the probabilities on the branches of the logic tree leading to the end branch. The result of the analysis is a discrete distribution of maximum magnitude, which can be used directly in the PSHA.

For assessing the Mmax of area sources, the procedures discussed in Section 4.2 are used, and a distribution of Mmax is usually assessed directly. An approach for assessing Mmax for sources in the eastern United States has been proposed by EPRI (1993), which is based on tectonic analogies between the eastern United States and other stable continental regions worldwide. In the procedure, a *prior* distribution of Mmax is assessed based on a statistical analysis of the global data base, and this distribution is updated based on source-specific information.

4.3.3 Earthquake Recurrence

Earthquake recurrence for individual seismic sources is defined by the a-value (also called the activity rate), b-value (slope of the recurrence curve expressing relative number of exponentially distributed small- and large-magnitude earthquakes), and Mmax. As discussed in Section 4.2, alternative magnitude distribution models are often important for describing the recurrence behavior of individual faults. The goal of uncertainty characterization for recurrence is to define the range of variation of the frequency-magnitude distribution. There are several ways to do this, depending on the type of seismic source.

Consider first area sources, for which the basis for recurrence estimation is observed seismicity. The first source of uncertainty is the magnitude of earthquakes contained within any catalog. As discussed in Chapter 5, the preferred magnitude for PSHA is moment magnitude and, until the eastern United States catalog can be translated to moment magnitude, Nuttli magnitude (m_{bLg}). In the eastern United States, most of the catalog of instrumental earthquakes is given in terms of Nuttli magnitude, although the M>4.5 historical earthquakes have been converted to moment magnitude using isoseismal areas (Johnston and others, in EPRI, 1993). Johnston and others provide uncertainty estimates in the moment magnitudes for each of the historical earthquakes in Johnston's catalog. Likewise, EPRI (1989) considered the uncertainty in the m_{bLg} estimates in the catalog and propagated that uncertainty into the recurrence analysis. Commonly, recurrence curves for sources are fit to observed data using a maximum likelihood procedure, to account for variations in the number of earthquakes in each magnitude bin. The statistical variability in the mean recurrence within each magnitude bin can be defined using Weichert's method (1980). Based on the observed earthquake counts (accounting for catalog incompleteness) and based on the

assumptions above, a plot can be developed showing the observed counts by magnitude, the variability in mean rate at each magnitude bin, and a maximum likelihood fit to the observed data. An example is shown in Figure 4-4. If the seismic source is very active and has generated a large number of earthquakes throughout a range of magnitudes, then the recurrence relationship derived directly from observed data may be sufficient to describe the uncertainties in recurrence for the source. Unfortunately, this is rarely the case. Typically, the observed earthquakes are few in number and small in magnitude. Hence, additional effort is required to assess the uncertainty in recurrence parameters.

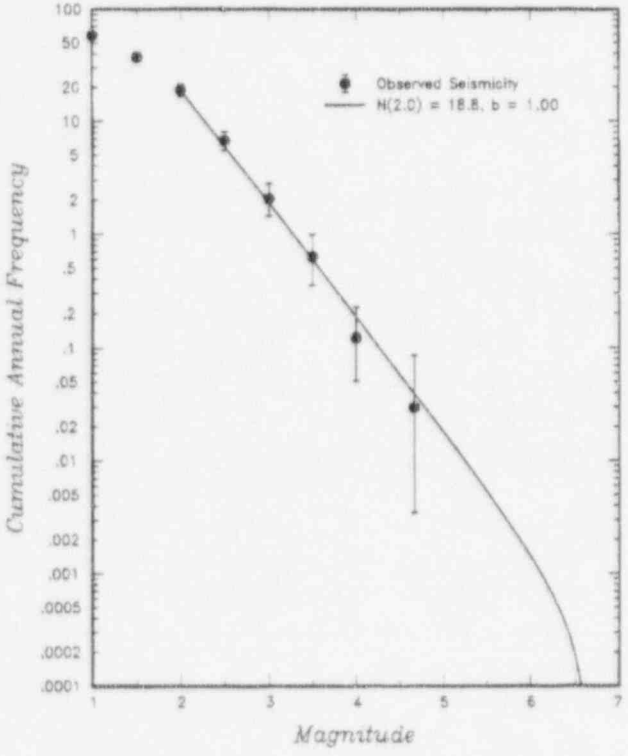

Figure 4-4 Example recurrence curve and observed seismicity for an area seismic source. The curve is the maximum likelihood truncated exponential recurrence relationship. The dots denote the mean annual frequency of observed earthquakes and the vertical error bars denote the 90% confidence interval on the cumulative rate of observed earthquakes (corrected for completeness). The parameters of the truncated exponential recurrence relationship are the cumulative annual frequency of events larger than magnitude 0 (the a-value), and the slope of the $\log_{10}$ frequency-magnitude recurrence curve, b.

Two alternative approaches have been used to describe the uncertainties in recurrence relationships (e.g., Savy et al. 1993). (Again, we are discussing an area seismic source and assume an exponential magnitude distribution). In the first approach, uncertainties in a-values and b-values are defined, including the correlation between the two parameters. Experience in the 1989 LLNL study (and corrected in the 1992 study) has shown that unintentional combinations of a- and b-values can result if the correlations between a and b are not defined. For example, suppose that one expresses the mean and uncertainty in a-value and the mean and uncertainty in b-value for a particular source. Unless the correlation between the two variables is specified, there may be combinations of a- and b-values that lead to unintended recurrence rates (e.g., a high a-value may be combined with a low b-value, resulting in high rates for large-magnitude earthquakes).

In the second approach, frequencies or recurrence intervals are assessed at particular magnitude levels. In the LLNL (1992) approach, these frequencies were elicited at two levels: at lower magnitudes where observed data are present and at larger magnitudes close to the maximum. The uncertainty in the frequency estimate can be described by a best estimate and a range of values. The net effect of this approach is also to eliminate unintentional extreme recurrence distributions that can result from assessing a-values and b-values independently.

The choice of the magnitude distribution model is usually based on the type of seismic source being considered: the exponential magnitude distribution is commonly considered appropriate for area sources (which presumably contain multiple faults), and the characteristic earthquake model is commonly considered appropriate for individual faults. There may be cases where the choice of the magnitude distribution model is uncertain. For example, a relatively small area source that includes a highly active zone of seismicity (e.g., the New Madrid seismic zone) may be characterized by either an exponential distribution (because of its areal extent) or a characteristic distribution (because the seismicity may be dominated by a single fault).

For fault sources, the uncertainties in earthquake recurrence are generally related to uncertainties in the models and parameters that are used to make the assessment. For example, a common approach to assessing earthquake recurrence is the use of fault slip rate, whereby the fault slip rate (which is uncertain) is multiplied by the area of the fault (also uncertain) and the rigidity of crustal rocks to arrive at a total average seismic moment rate. This seismic moment rate is then partitioned into earthquakes of various sizes according to a magnitude distribution model such as the characteristic earthquake model. Alternative approaches to estimating fault-specific recurrence are the use of paleoseismic recurrence intervals (having uncertainties in both the intervals and the sizes of the paleo-events) and geodetic strain data (uncertainties much like slip rate data). In all of these cases, a logic tree procedure is an effective way to sequence the models and parameters leading to the recurrence estimates and to propagate the uncertainties into the recurrence distributions. For fault sources, the observed seismicity is usually too sparse to provide a strong constraint on the recurrence rate, but, for more active faults, could control the recurrence rates in the lower magnitude part of the distribution.

4.3.4 What SSC Information is Elicited and What is Calculated?

The purpose of this section is to summarize information that must be elicited from SSC experts and describe which information can be calculated by the hazard analyst. The goal here is to provide an idea of the types of tasks that SSC experts should be prepared to accomplish.

At a minimum, the SSC experts should be prepared to provide the following:

- Seismic source map and alternative maps or alternative source configurations and the probability of activity for each.
- Any source activity dependencies (i.e., the assessment that one source is active if another is active).
- Focal depth distribution for all sources.
- Three-dimensional geometries for faults and associated uncertainties.
- Maximum magnitude distribution for all sources.
- Designation of an earthquake catalog for each source. Time periods over which the catalog is complete (either zero or fractional).
- Choice or approval of a magnitude distribution for each source. Where appropriate, multiple models should be specified with weights, or distributions of parameter values should be given (if one model is used). If the exponential model

  (Log N = a - bM) is used, an a-value and b-value must be specified, and the expert should use either:

  — pairs of values with weights

  — joint distributions of a and b with the correlation specified

  The magnitude distribution may vary in space within a source area.
- For faults, the expert may specify the distribution as above, or may use slip rate, b-value and magnitude distribution to specify the recurrence rate. These parameters can be readily transformed to magnitude recurrence information by the analyst. As is the case for areal sources, either discrete or continuous distributions may be used, but correlations must be specified.

4.3.5 Considerations on the Spatial Variation of Seismicity Within a Seismic Source Philosophical Basis and Implications of the Assumption of Homogeneous Seismicity

It has been assumed in many seismic-hazard studies that seismic sources of types 2, 3, and 4 have homogeneous seismicity; i.e., that the a-value and the b-value are the same for all points within the seismic source. According to this assumption, if the seismicity catalog were extensive enough over time, one would observe the same density of earthquakes (events per unit

area) in any small area within a given source. This assumption has two very important implications on the calculated seismic hazard, as follows:

1. On the mean hazard. All sites located within a homogeneous seismic source (and sufficiently far from the source boundary) will have the same mean hazard due to this source, regardless of the spatial distribution of historical earthquakes within the seismic source.

2. On the statistical uncertainty in hazard. The activity rate and b value for this seismic source are calculated using all the historical earthquakes in the source. The statistical uncertainty in the rate and b value are lower than they would be if this source was sub-divided into two or more smaller sources.[1]

These two effects are particularly important for large seismic sources of regional or tectonic-providence dimensions (i.e., source types 3 and 4).

The assumption of homogeneity is almost always made for the sake of simplicity (i.e., fewer parameters are required) and is driven more by ignorance than by a firm belief in homogeneity (e.g., the expert does not sub-divide this large source because he/she does not know how to sub-divide it, not because he/she thinks it has homogeneous seismicity). If a seismic source (particularly sources of types 3 and 4) is defined on a basis other than patterns of seismicity (see Table 4-1), there is no reason for the assumption of homogeneity to be valid.

In any seismic hazard analysis, the assumption of homogeneous seismicity must be justified and alternative assumptions may have to be included in the model of seismic sources. As a minimum, one must confirm that the assumption of homogeneous seismicity is not inconsistent with the spatial distribution of historical seismicity, using the statistical tests to be described below or other appropriate statistical techniques. If the assumption of homogeneity is not consistent with the data, and the source is a significant contributor (>30%) to the hazard at the site, the source should be sub-divided into more homogeneous sub-sources or the assumption of constant rate and b throughout the source must be relaxed by using the EPRI approach (see EPRI 1986; VanDyck 1986) or a similar approach.

Statistical Tests for Homogeneity

The following simple statistical test indicates whether the assumption of homogeneous seismicity is consistent with the spatial distribution of historical seismicity within a seismic source. The test consists of the following five steps:

1. Sub-divide the seismic source into smaller sub-sources using, for example, the 1-degree or 0.5-degree grid used by EPRI (1986).

2. Calculate the observed historical earthquake rate in each sub-source.

3. Calculate the expected number of earthquakes using the homogeneous model in each sub-source, considering the sub-source area, the length of the catalog, and the catalog-completeness assumptions.

4. Compare the expected and observed numbers of earthquakes in each sub-source and flag those sub-sources with statistically significant differences. Figures 4-5 and 4-6 illustrate this test. Figure 4-5 shows the source and its historical seismicity; Figure 4-6 shows the flags associated with the 10% and 2% significance levels.

5. Examine the number and pattern of flags to determine if the assumption of homogeneous seismicity is consistent with the catalog. If α is the significance level used in step 4, one would expect approximately a fraction α of the sub-sources to have the associated flags. Too many flags indicate that the assumption of homogeneous seismicity is inconsistent with the catalog; too few flags indicate that the catalog is too limited to provide any indication about spatial patterns of seismicity. Even if the number of flags is not unexpected,

[1] As a rule of thumb, the coefficient of variation in the activity rate is approximately $n^{-1/2}$ where n is the number of earthquakes in the seismic source. The standard deviation in the b value is also proportional to $n^{-1/2}$.

the concentration of positive or negative flags in a certain portion of the source is an indication that the assumption of homogeneous seismicity is inconsistent with the catalog (at a spatial scale larger than the grid size)[2]. In Figure 4-6, the number of flags clearly indicates that the assumption of homogeneous seismicity is not consistent with the spatial pattern of seismicity in the catalog.

This test is implemented in the EQPARAM code developed by EPRI (EPRI 1986). The code estimates seismicity parameters under assumptions of homogeneous or spatially varying parameters, but it may be easily used to perform these tests only. Also, this test is relatively easy to implement as a stand-alone code.

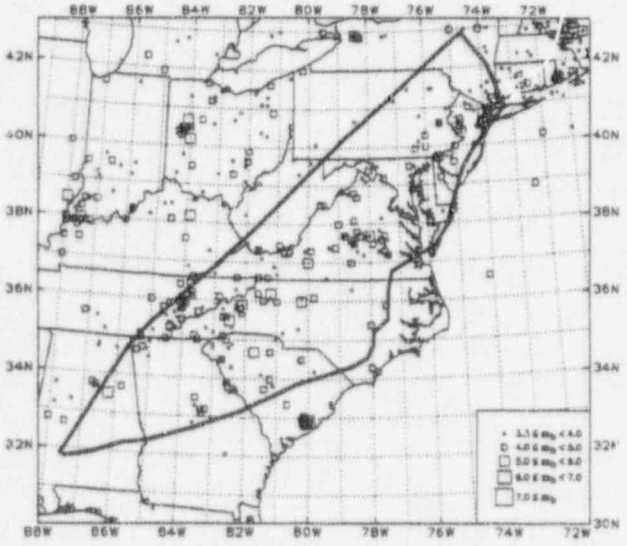

Figure 4-5 Map showing a background source for the southern Appalachians and the historical seismicity in the EPRI catalog.

[2]This test is not a really a spatial-homogeneity test. Rather, it is a series of univariate significance tests. Thus, the test requires some interpretation from the expert or analyst in Step 5. On the other hand, the test is easy to implement, intuitive, and very informative. Other tests for spatial homogeneity are available in the literature (e.g., Ripley 1981).

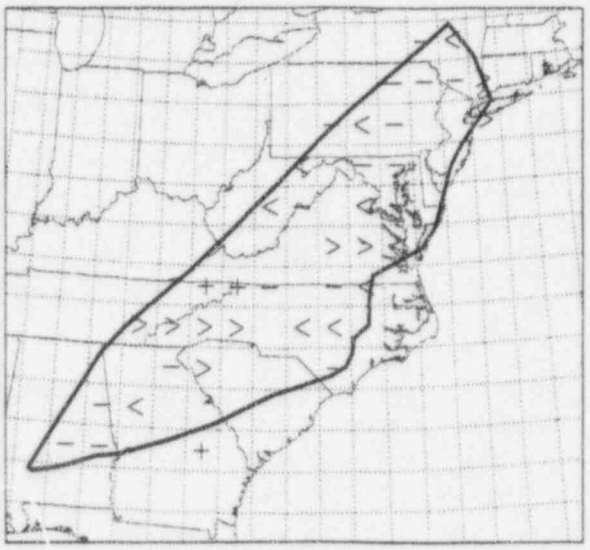

Figure 4-6 Diagnostic flags from the statistical test of homogenaity of seismicity for the source shown in Figure 4-5. "+" ("-") indicate that the observed count in a sub-source is significantly higher (lower) than predicted at the 10% significance level; ">" ("<") indicate that the observed count is significantly higher (lower) at the 2% level. Approximately 8% (i.e., 10-2) of the sources should have ± flags; 2% of the sub-sources should have ">" or "<" flags.

Special Circumstances Requiring Models with Spatial Variability

Even if the above statistical tests do not reject the assumption of homogeneous seismicity, there may be situations where this assumption alone may not be sufficient for the characterization of seismic hazard and its uncertainty at a site. The following two criteria are proposed in this regard.

Spatial variability should be considered for a seismic source, even if the assumption of homogeneous seismicity is not rejected, if the following two conditions apply:

1. Earthquake count in the source. The earthquake count is very small, so that it provides little indication about the spatial distribution of seismicity in the source (e.g., some sub-sources contain one to three events, others contain none).

and

2. Percent contribution to seismic hazard at the site (based on preliminary seismic-hazard results). The source contributes more than 30% of the seismic hazard at the site (for any exceedance probability or ground-motion measure of interest).

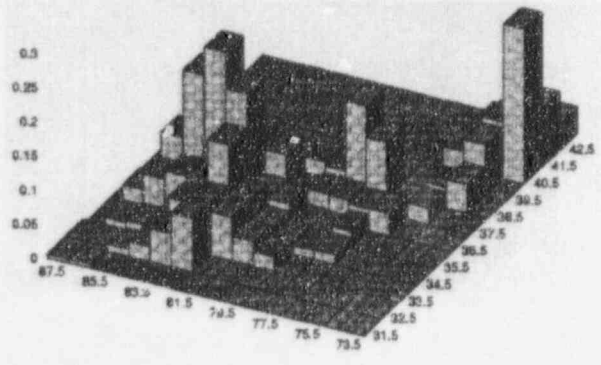

Figure 4-7 Spatial distribution of rate (m_b>3.3) per unit area (events/yr/deg^2) for the source of Figure 4-5, obtained with the assumption of low smoothing on *a* and high smoothing on *b*.

The motivation for requiring the use of spatial variability, even though it is not required by the catalog, is that the objective of a seismic-hazard study is to quantify both the mean hazard and its uncertainty.

Spatially Varying Seismicity Parameters: EPRI Model

The EPRI model for spatially varying seismicity parameters is presented here as one possible model for relaxing the standard assumption of homogeneous seismicity parameters throughout the seismic source. Another possibility is simply to sub-divide the seismic source into a few sub-sources, so that the seismicity becomes more homogeneous[3].

The EPRI model sub-divides the seismic source along a one-degree latitude-longitude grid, resulting in sub-sources of dimensions of one square degree or less. Values of a and b are estimated for each sub-source. This model can accommodate observed spatial variations of historical activity within a given source. This approach constitutes a moderate departure from traditional seismic hazard analysis, in the sense that it assumes (or can assume) a relatively smooth spatial variation of the activity rate. Also, each sub-source retains all the properties of the seismic sources of traditional seismic-hazard analysis. This model may be considered as intermediate between historical and conventional seismic hazard analysis.

In order to avoid problems with sub-sources that have low or no earthquake counts, and to reduce the uncertainty in the estimates of a and b, smoothing assumptions are introduced, which impose dependence between the seismicity parameters in adjacent sub-sources. Thus, the seismicity parameter in one sub-source depends, to some extent, on the earthquake counts in adjacent sub-sources within the same source. Conceptually, the smoothing assumptions may be interpreted as prior distributions on the degree of spatial roughness of a and b within the seismic source. Because this is not an easy concept, experts typically specify multiple values of the smoothing parameters, with associated weights, as an indication of their subjective uncertainty about the appropriate prior distribution.

Smoothing is specified separately for a and b. The smoothing assumptions range from full smoothing to no smoothing. Full smoothing on both a and b is the same as assuming that seismicity in the source is homogeneous; no smoothing on both a and b is the same as treating each sub-source as a separate source. Typically, b is assumed to be smoother than a, because b has been observed to be more geographically stable. The statistical test described earlier provides guidance for the selection of smoothing assumptions.

The seismicity parameters a and b for each sub-source are estimated using maximum penalized likelihood, where the penalty terms represent the smoothing assumptions. The result is a pair of "maps" for a and b within the source. As an example of the type of results obtained, Figure 4-7 shows the activity rates for the source in Figure 4-5, calculated under the assumption of low smoothing on a and high smoothing on b. Because the equivalent number of parameters

[3]One potential problem with the approach of sub-dividing the source into a few sub-sources is the choice of where to subdivide. Unless there are sharp contrasts in seismicity or obvious boundaries suggested by the the geology or geophysics, the choice of boundaries may lead to biases.

being estimated is larger than when homogeneity is assumed, the associated statistical uncertainty is higher. This uncertainty depends on the earthquake counts in the various sub-sources and on the smoothing assumptions; lower smoothing leads to more uncertainty in seismicity parameters. This higher uncertainty is not necessarily undesirable, because it may be more realistic.

Analytical calculation of the statistical uncertainty in the seismicity parameters is difficult because the smoothing introduces correlation. One simple way to quantify this uncertainty is by using a technique known as "bootstrapping," where artificial catalogs are generated (using the actual catalog or the estimated seismicity model) and maps of a and b are calculated for each artificial catalog. One then propagates this uncertainty into hazard space by calculating the seismic hazard associated with each alternative map of a and b.

The statistical uncertainty in the hazard—due to statistical uncertainty in spatially varying models of seismicity—must be quantified as part of the hazard calculations as described above. In fact, this statistical uncertainty may often be more important than the uncertainty about the proper level of smoothing.

The EQPARAM software package (EPRI 1986) performs all the calculational steps described above, including bootstrapping. Some further enhancements to these techniques have been proposed and tested (Veneziano and Luna Pais 1986; Veneziano and Chouinard 1987), the most significant enhancement is the optimal selection of smoothing parameters. These enhancements are not currently implemented in EQPARAM.

Appendix I contains detailed examples showing the a and b maps obtained under different smoothing options, the associated statistical uncertainties, and the effect of these uncertainties on the calculated seismic hazard.

4.3.6 Significant Changes in Hazard due to Seismic Source Characteristics

A significant amount of effort must go into seismic hazard analysis to obtain meaningful results, and this effort should be used in the most efficient way possible. To this end, it is important to examine which SSC parameters contribute significantly to seismic hazard, and to determine when changes in those parameters make significant differences to the computed hazard. The identification of important parameters can then be made on an informed basis so that maximum effort can be guided toward evaluating those SSC parameters that make the most difference to the hazard.

The real benefit in considering which parameters contribute to significant changes in hazard comes from being able to concentrate on the evaluation of important parameters (both in the sense of the best estimate and of the uncertainty) while neglecting, or treating in an approximate fashion, other parameters that are not significant or are only marginally significant. Thus, consideration of significant parameters involves both an evaluation of what drives the seismic hazard in the sense of the best-estimate hazard, and what contributes significantly to uncertainties in hazard.

To these ends, the Committee has formulated a procedure to guide the evaluation of which SSC parameters deserve the most scrutiny. The procedure is presented in Appendix G. Also given in Appendix G are a series of hazard analyses conducted by Risk Engineering and LLNL for a set of conditions related to source-to-site distances, focal depth distribution, maximum magnitudes, earthquake magnitude distributions, b-values, and a-values (activity rate). The goal of these analyses was to determine the relative importance of certain SSC parameters and combinations of parameters relative to the best-estimate hazard and the contributions to the uncertainty in the hazard (as a function of the structural period of the ground motion estimate and probability level). Some of the important conclusions of these analyses (which are given fully in Appendix G) are the following:

- Uncertainty in fault location causes a moderate sensitivity for most sites for high-frequency ground motions, and less sensitivity at low frequencies. For source zones, this applies to sites located outside the

source, but especially near the source boundary.

- Sensitivity to depth distribution is negligible except at small source-to-site distances (less than 50 km).
- Sensitivity to maximum magnitude is largest at large source to-site distances. It increases with ground motion amplitude, and is largest when the mean Mmax values are lower. (The sensitivity is greater when the mean Mmax is 6.0 rather than 7.5 for fixed a- and b-values).
- Sensitivity to the b-value is moderate, except at small source-to-site distances (less than 25 km).
- Sensitivity to whether an exponential or characteristic magnitude distribution is used depends on whether a slip rate constraint or a seismicity constraint is used to fix the rate of activity (a-value). If a slip-rate constraint is used, the maximum sensitivity occurs for very close or very distant sites. If a seismicity constraint is used, calculations at all distances are sensitive to the choice of the model.

4.4 Specific Expert-Elicitation Guidance on Seismic Source Characterization

4.4.1 Introduction

Seismic source characterization involves assessment of the location, rates, and maximum size of future earthquakes, which are variable, i.e., have aleatory uncertainties. Also, there is sparse historical evidence in most areas as well as limited understanding of the mechanisms associated with earthquake occurrences. Thus, our ability to model earthquake occurrences is subject to epistemic uncertainty. Because of the limited experience and understanding there is a diversity of interpretations of seismic source characteristics within the informed technical community and, for purposes of PSHA, it is necessary to capture the community distribution of source characteristics. Thus, as discussed in Chapter 3, SSHAC recommends using either the TI or TFI approach to derive the SSC inputs for a PSHA. The study level (1-4) will depend on the scope of the study and the expected complexity and contentiousness of the SSC.

The approaches discussed below are consistent with the general guidelines and concepts regarding the TFI and TI approaches discussed previously in Chapter 3. However, the procedures and methods discussed in this section are specific to SSC and differ somewhat from those procedures outlined for incorporating judgments related to ground motions (Section 5.6). For example, an essential and first step in seismic source characterization is the identification of seismic sources. In most regions of the U.S., the interpreted geometry of seismic sources will vary with the source characterizer and, therefore, each expert's map of sources will be different. The subsequent characterization of these sources (e.g., by recurrence parameters) will be specific to the particular interpretation of the expert. Because of this, there is no easy way to compare the results of the characterization from one expert to the next directly. More importantly, there is no easy way to integrate the results of the analysis at an intermediate step (say, the seismic source maps), nor can the final results of the seismic source characterization be readily combined, other than at the final step of the seismic hazard analysis. An exception might be in highly active tectonic environments in which the seismic source maps among various experts (reflecting active faults) might be very similar. Also, it may be possible to arrive at a consensus source map that a group of experts can all endorse. In this case, the uncertainties in scalar quantities (e.g., the slip rate on a particular fault) may be amenable to integration across multiple experts. In the future it may be desirable to move SSC in a direction that allows for more integration at intermediate levels of the analysis; for example, through the development of consensus seismic source maps for regions of the U.S.

A SSHAC-sponsored workshop designed to examine the pros and cons of SSC expert elicitation methodologies (see Appendix H) is the resource for the following discussion. The participants at the workshop were SSC experts who themselves have been elicited as part of

several seismic hazard analyses. As such, their experience represents a unique data base from which to draw conclusions about which SSC elicitation approaches "work" and which "don't work." Many of the elements of the recommended SSC expert elicitation methodology find support in the conclusions drawn by the experts at the workshop, as well as reviews by SSHAC of several PSHA projects conducted for both regional and site-specific applications.

In discussing the recommended methodologies for incorporating SSC expert judgment, the section begins with recommended approaches to quantifying SSC expert judgments using either the TFI approach or the TI approach, then considers how the approaches may vary as a function of the resources available for the project (resource-intensive versus modest resources) and the application (site-specific versus regional hazard assessment).

From the standpoint of seismic source characterization for PSHA, the Committee concludes that either the TFI or TI approaches can be used to quantify SSC characteristics and uncertainties, depending on the expected contentiousness of SSC in the region of interest. Because of SSHAC's emphasis on capture of the diversity of interpretations within the informed technical community, we will emphasize Study Levels 3 and 4, discussed in Section 4.4.3 and 4.4.2 respectively, based on the use of multiple experts as the primary sources of inputs. Modifications, assuming only limited resources and site-specific versus regional studies, are discussed in Section 4.4.4.

4.4.2 The TFI Approach

The TFI approach to deriving SSC inputs for a PSHA is to be used for those studies in which there is considerable diversity of interpretations of the seismic sources and/or the seismicity in the region of interest. Use of the TFI approach is based on the premise that representation of the community distribution of SSC's is best derived by eliciting inputs from a panel of experts, acting as evaluators and individual integrators. The products of the elicitation are:

- Alternative seismic source maps and distributions of seismic source characteristics from each expert representing his/her SSC with uncertainty

and

- Alternative seismic source maps and distributions of seismic source characteristics describing each expert's view of the informed technical community's distribution of seismic sources and seismicity.

An important TFI function is to facilitate during the workshops prior to the elicitation and involve proponents, resource experts, and evaluators. These workshops must include discussions of the historical data bases of earthquakes, geologic and tectonic models regarding the localization of seismicity, models of seismic source interpretations, frequencies and distributions of magnitudes of earthquake, as well as methods and procedures for analyzing and summarizing the historical data for use in developing SSCs.

Another important part of the TFI process is the elicitation of inputs from the evaluator experts. Because the experts need to provide descriptions of aleatory uncertainty and to describe their epistemic uncertainties in providing these descriptions, it is essential that the elicitation involve individual interviews. It is also important that experts be educated and trained in the concepts of aleatory and epistemic uncertainties as well as in ways of formulating and quantifying their epistemic uncertainties. The basic steps in the recommended methodology for SSC are given below in terms of the specific application to SSC.

1. Conduct careful expert selection The process of expert selection should be based on a clear set of criteria aimed at capturing a full range of diversity of expert interpretations.

2. TFI role The technical facilitator/integrator should play a strong role, running workshops and expert interactions, monitoring the behavior and participation of the experts, conducting calculations and sensitivity analyses, documenting the final results, and taking intellectual responsibility for the results of the project.

3. Provide a uniform data base to all experts SSC-related data sets, as defined by the experts themselves, should be provided to all of the experts in formats most useful to the experts.

4. Conduct multiple expert interactions Interaction among SSC experts is strongly recommended, through such vehicles as workshops, small working meetings, etc.

5. Elicit SSC judgments from experts Individual expert elicitations should be conducted through person-to-person interviews. Elicitations of expert teams is also acceptable.

6. Conduct sensitivity analyses and submit feedback to experts Following the elicitations, extensive sensitivity analyses should be conducted by the TFI and provided to the experts. They then should interact again as a group to review their interpretations.

7. Finalize SSC interpretations and combine at hazard level Integration/aggregation of SSC interpretations usually occurs at the hazard level. The TFI should create the proper conditions, through the application of 1 through 6 above, to combine the expert judgments using equal weights. Allowance should be made for cases where unequal weights are appropriate (see Section 5.3).

8. Peer review An active or "participatory" peer review should be conducted throughout the study with the particular focus of the process that was followed in conducting the SSC assessment.

Each of these components of the methodology is discussed below.

Expert Selection Because the TFI approach relies on the direct judgments of SSC experts as basic input to the PSHA, the selection of the experts is very important. Further, a desirable outcome of the SSC expert elicitation procedure is to develop a strong, defendable basis for equally weighting the interpretations of the SSC experts when combining the assessments for the hazard analysis. A key part of that basis is that the experts were selected according to a set of criteria that ensure high-caliber, equally-qualified experts.

Two equally acceptable alternatives for expert selection are that it be carried out by the TFI or by the project peer panel. In either case, the entire expert selection process must be thoroughly documented, such that an independent third party could review and understand the procedure followed based on the documentation.

The criteria for selecting the SSC experts must be established and documented. Important criteria should include: geologist or seismologist with a strong professional reputation and widely recognized competence based on academic training and relevant experience, tangible evidence (e.g., peer-reviewed publications and reports) of relevant studies and experience, and availability and willingness to commit the full time and effort needed for the study. In addition, the individual must be willing to forsake the role of a "proponent" espousing a particular hypothesis for that of an "evaluator" who considers all viewpoints and evaluates their relative credibility. In addition to selection criteria for individuals, the project leader should ensure that the experts as a group represent diversity in technical interpretations, areas of technical expertise, and institutional and organizational backgrounds.

If the SSC elicitation will be conducted with teams rather than individuals, the experts should be selected such that a diversity of views and expertise is represented across the teams. In addition, the experts should be informed that they will be working in a team environment.

Following development of the selection criteria, a large pool of potential experts should be identified. This can be done by identifying a few potential experts first and then asking each of them for their nominations. Alternatively, an independent panel may nominate potential experts. The large pool is then narrowed down to a smaller number (about 7 to 15 experts for regional studies, perhaps fewer for site-specific studies) depending on the range and diversity of views. These individuals should then be contacted and informed fully of the purpose of the study, the

specific assessments that they will be asked to make, the manner in which their assessments will be used in the PSHA, and the time and effort that they will need to devote to the project. Either lack of availability or insufficient motivation to commit to the required level of effort are sufficient grounds to exclude that expert from consideration.

At the start of the project, the TFI should establish and discuss with the experts the criteria that would be used to remove an expert from the panel. As discussed in Section 3.3, these criteria would include such problems as the lack of commitment on the part of one of the experts to devote sufficient time to the project, an expert who refuses to forsake the role of a proponent for that of an evaluator, an expert who lacks the interpersonal skills to interact with the other experts in a professional manner. The responsibility for monitoring the performance of the experts lies with the TFI and he, along with the Project Leader and Sponsor, are responsible for the removal of an expert from the panel.

Role of the TFI The TFI role implies proactive participation in dealing with the SSC experts and their elicitation. The TFI team includes a technical peer of the experts and can take a leadership role in selecting the experts, organizing and directing the workshops and other expert interactions, facilitating expert interactions, and monitoring the participation of the experts. Behavioral approaches to achieving integration include active participation in workshops, challenging the interpretations of the experts, looking for areas of consensus among the experts, weeding out differences of opinion due to misunderstandings or definitions, etc. As an example, the TFI should be responsible for developing the sensitivity analyses and feedbacks that are provided to the experts following their elicitations but prior to finalization. This is an opportunity for the TFI to focus the discussion on the implications of the assessments to the hazard results, and the technical basis for the diversity of interpretations.

The TFI approach as applied to SSC may differ somewhat from the TFI approach outlined for ground motions (Section 5.6). In particular, the role of the TFI in integrating the interpretations of the SSC experts occurs primarily through the process of selecting, training, and interacting among the experts (i.e., the seven key elements of the TFI approach discussed earlier). Because there is, as yet, no apparent way to combine SSC expert's seismic source maps (except, perhaps, in highly active areas where the fault locations might be agreed upon, or other cases where a set of sources can be chosen), there is no way to, in turn, develop a distribution of recurrence parameters that properly expresses the range of interpretations across the experts. In contrast, it may be possible to develop these types of distributions on ground motion values for a given set of magnitude and distance combinations. Thus, the role of the TFI in the ground motion case can include an appropriate 'weighing' of the alternative interpretations to arrive at a distribution for certain parameters. In SSC, each expert develops a set of seismic sources and associated parameters (and their uncertainties). These assessments can only then be combined at the end or hazard level, and such combination will require a "weighting" of the interpretations from the various experts. The TFI is responsible for selecting and implementing the scheme for integrating the SSC expert's interpretations. Recommended methods for this integration are discussed further below.

Despite the differences in detail of the TFI approach for ground motions and SSC assessment, they are founded on the same premise: a fundamentally important component to the integration of multiple expert judgments is interaction among the experts. Interaction is a mechanism for resolving unintentional disagreements, presenting interpretations, challenging the interpretations of others, reviewing data bases, and, ultimately, for assisting the experts in their evaluations. Expert interaction is discussed further below.

Data Bases A major responsibility of the TFI is to provide a comprehensive and uniform data base to the experts. Early in the study, the experts and the TFI should identify a comprehensive set of technical issues that will need to be addressed in the source characterization. The data needed to

address those issues should then be identified by the experts in a workshop forum. Formats should be specified, responsibilities assigned for data retrieval, and a realistic schedule established for data compilation. Depending on the wishes of the experts, much of the data base may be amenable to compilation and transfer in digital form to the experts. Care should be taken to clearly define the data needs in terms of "raw" data and any processing that is required by the experts. Any data processed according to the request of a single expert should be made available to any other expert who may desire it. The objective of this effort, which for large regional studies can be very resource-intensive, is to provide all of the experts equal access to the data that are most pertinent to the assessments that they will need to make. Although most of the data base effort will occur early in the project, provision should be made to allow for additional data requests or data processing later in the project.

Expert Interactions Interaction among the SSC experts is a fundamentally important aspect of the SSHAC methodology. Time for expert interaction should be allowed in workshops, small group meetings, and informal communication. Source characterization experts are few in number, often rely on the same data bases, and interact frequently as part of their professional activities. Therefore, the interactions recommended here are a natural scientific extension of the way that earth scientists formulate their ideas about seismic sources.

The purposes of SSC workshops are the following:

- Identify technical issues of greatest importance to PSHA
- Specify the data needed to address these issues
- Educate the experts on the available data and seismotectonic interpretations (resource experts, who themselves are not elicited, may make presentations at the workshop)
- Educate the experts on the methods and procedures that are available to characterize seismic sources and clearly specifying how their assessments will be used in the seismic hazard analysis
- Train the experts on expert elicitation procedures
- Review, technically challenge, and defend SSC hypotheses and interpretations
- Complete behavioral integration of expert interpretations
- Include observation by those not directly involved in the elicitations (e.g., sponsors, regulators, etc.)

The discussion and challenge of interpretations in SSC varies with the topic being considered. For example, seismic source maps developed by experts in the eastern United States are commonly quite different. In a workshop setting, experts can discuss the technical basis for the configuration of their source zones. However, because the sources drawn by any two experts are different, it is difficult to directly compare and challenge the interpretations. Likewise, because recurrence parameters are associated with particular seismic sources, it is often difficult to challenge the parameter values for a given source zone. In contrast, the seismic source maps in active areas are often quite similar among multiple experts. For example, the major faults are usually depicted and characterized. In these cases, the SSC experts are able to discuss and challenge the SSC interpretations made for particular faults, and to directly compare alternative interpretations. For example, the slip rate and recurrence intervals on the San Andreas and Hayward faults could be discussed and debated in a workshop environment.

The number of workshops and their content will vary with the particular study but must cover, at a minimum, the following areas: 1) identification of technical issues and data bases to address them, 2) available data and seismotectonic interpretations relevant to the study region, 3) available procedures and methods for defining seismic sources, specifying earthquake recurrence, estimating maximum magnitudes, and characterizing the uncertainties in these assessments, 4) procedures that will be followed

in eliciting expert judgments and for monitoring the performance of the experts, and 5) feedback of preliminary interpretations and sensitivity analyses.

The schedule, content, and conduct of SSC workshops are the responsibility of the TFI. Ample time must be provided to prepare for each workshop. Responsibilities for presentations should be given well in advance, and written materials should be provided prior to the workshop. During the workshop, the TFI should lead the technical discussions and ensure that all topics are given adequate time for consideration. It may be useful to have "proponents" advocating particular hypotheses or viewpoints present their technical arguments. The TFI must maintain a balance between rigid control and free-wheeling discussion. Less-vocal participants should be encouraged to voice their opinions. The goal of the workshop is communication, education, and reduction of unintentional disagreement.

In addition to workshops, other small meetings may be held to ensure progress and assist the experts. Individual experts may wish to work as small teams to develop their interpretations and to get feedback from their peers. The TFI may need to provide guidance regarding SSC-related procedures. If a team approach is being followed, these small meetings will likely be conducted within the team.

Workshops are effective mechanisms for the free exchange of data and interpretations. They are not recommended as a vehicle for performing the actual elicitations. However, it may be useful to conduct "example elicitations" in a workshop or meeting format to illustrate the manner for eliciting expert judgments, the procedures for quantifying uncertainties, and the methods for documenting the assessments.

Elicitations of SSC Experts The elicitation of either individual experts or of teams of experts is acceptable, although the procedures for doing so differ somewhat.

Assuming that the experts have received training in elicitation procedures and, as appropriate, have undergone example elicitations to provide a degree of comfort in the process, it is recommended that individual experts be elicited in small interview sessions. Present at the sessions should be the expert, the TFI team consisting of the technical expert and an elicitation expert with experience in subjective probability assessment. It is also acceptable to have other "resource specialists" (say, with specialized knowledge of statistics, tectonics, etc.) available to provide information to the expert. Every effort should be made to put the expert at ease in the elicitation and to maintain flexibility in the questioning to allow the expert to express his/her interpretations and uncertainties in his/her own way. The elicitation is not a final examination; the expert may bring any resource material that he/she feels will assist in making the assessments.

It is important for the individual being elicited to be discouraged from playing the role of a "proponent" who advocates a single hypothesis or viewpoint. The role of an "evaluator" is to evaluate the technical merits of all hypotheses and to assess the relative credibility of each. In doing so it is expected that multiple hypotheses will have some level of credibility to the expert, and the credibilities can be readily quantified using subjective probabilities. As discussed in Section 3.3, it is useful to have two stages of elicitation: a Stage I elicitation in which the expert plays the role of an evaluator who represents his own range of knowledge and uncertainty, and then a Stage II elicitation in which the expert is asked to represent his/her assessment of the diversity of views that result from questioning the larger informed technical community. A purpose of the second elicitation is to identify those cases where an expert may recognize that his interpretation is significantly different from his perceptions of the interpretations that the scientific community would have if they were similarly informed.

Elicitation of SSC expert judgments using written questionnaires is not recommended. Experience has shown that they are often subject to different interpretations and can be confusing. It is acceptable and even desirable, however, to provide a written questionnaire to the expert prior to an interview session for information purposes

only as a means of focusing attention on the pertinent SSC issues.

The elicitation must be carefully documented, including both the interpretations and uncertainties expressed by the expert and the technical bases for the interpretations. A recommended procedure is for the TFI to record in writing the elicitation and, subsequent to the session, provide the written documentation to the expert for review and revision. Acceptable, but less preferred, alternatives are to record the session electronically and provide a summary of the transcription to the expert, or to require that the expert furnish the written documentation following the session. Experience has shown that these alternatives are not as efficient as the preferred approach.

It is possible that the interview session, which is usually an exhausting experience for the expert, will not cover the entire source characterization required for the analysis. Alternative, equally acceptable, remedies for this problem are the following: conduct a follow-up elicitation session to complete the assessment, or allow the expert to complete the assessment privately and provide the results at a later time. In order for the latter approach to work, the elicitation session must have covered the entire spectrum of source characteristics for several sources in the interview session (presumably the sources of greatest significance to the hazard). The expert can then complete similar assessments for the remaining sources.

Another expected circumstance is the case where an expert may feel uncomfortable with answering a particular element of the assessment, usually because of a lack of experience or expertise in a particular area. It is recommended that the SSC analysis allow for the expert to decline answering questions in these areas. To do so, the SSC analysis should, prior to the elicitation, have made every effort to select the proper experts and to educate the experts. Failing this, procedures should be established for dealing with the problem and these should be communicated to the experts prior to the elicitation. For example, if the expert declines to define seismic sources in geographical regions remote from his/her past experience, the expert could use sources defined by the other experts. If the expert declines to provide certain parameters for his/her sources, an agreed-upon 'default' methodology for estimating the parameter should be invoked. Obviously, this is a difficult problem in the SSC area, but because of the multi-disciplinary nature of the problem, is one that should be anticipated. The use of multi-disciplinary expert teams is one way to mitigate the problem.

In some cases, the expert may prefer to specify a methodology or procedure for calculating a particular parameter, rather than provide a direct assessment of the parameter value itself. This is acceptable and provision should be made for the TFI to conduct the calculation in a timely manner (if possible, at the interview session itself) then ask the expert to review the calculation.

Despite the small experience base on which to make the recommendation, it is advised that teams of experts be elicited in the interview session. Questions will be posed to the team as a whole and the team as a unit will be responsible for developing a consensus interpretation and uncertainty distribution that captures the diversity of their individual views. Clearly, depending on the questions being asked, some members of the team will defer to others at different points in the assessment and this is acceptable. It is the responsibility of the TFI to lead the team through the assessment and to emphasize that the team's responsibility is not to reach agreement on each issue, but to develop a range of assessments (e.g., alternative seismic source configurations, distributions of recurrence parameters) that effectively captures the thinking of the team as a whole. It is expected that the team elicitation will be more time consuming than individual interviews and may require multiple sessions to complete. Documentation must be conducted in the same manner as for individual interview sessions.

Feedback and Sensitivity Analysis Following the elicitations and prior to the final seismic hazard calculations, the preliminary results of the elicitations and a variety of sensitivity analyses should be prepared by the TFI, provided to the experts, and discussed in a workshop format. The

purpose of this exercise is to (1) allow each expert to see the preliminary interpretations made by the other experts, (2) understand the implications that the assessment has to the seismic hazard calculations, (3) identify the key SSC assessments that are most important to the hazard results, and (4) compare the "predicted" seismic source characteristics to "observed" data. The following examples show the types of results that might be provided to each expert in advance of the feedback workshop:

- Seismic source maps and seismicity parameters for all experts
- Mean and fractile seismic hazard curves based on the assessments of each expert
- Plots showing the relative contributions of each seismic source, magnitudes, and distances to the mean hazard
- Plots showing the contributions of various SSC uncertainties to the total variance in the seismic hazard results
- Comparisons of the predicted recurrence rates for each source with the observed seismicity for various completeness periods
- Comparisons of the total predicted regional recurrence rates with the observed seismicity rates
- Sensitivity analyses showing the variability in mean hazard as a function of assessed ranges of seismicity parameters.

The TFI should review the results with each expert and identify possible problems or inconsistencies. For example, an expert's interpretations might predict that the rate of seismicity (say, the number of M>5 earthquakes per year) is significantly larger than or smaller than the observed rate from historical seismicity. The TFI should review the technical basis for the expert's assessment, ensure that the expert is aware of the difference between his estimate and the historical record, and provide an opportunity for the expert to revise his assessment if desired. There are certain key assessments in SSC that are subject to possible misinterpretation or "error" unless the expert is educated by the TFI. These assessments are discussed in Chapter 3 on Seismic Source Characterization. One example is the assessment of the recurrence parameters a and b. Assignment of values to these parameters (and their uncertainties) without consideration of the correlation between the parameter values can lead to some unintended combinations of a and b-values and their associated recurrence rates. It is the responsibility of the TFI to identify possible problem areas and, through questioning the expert, ensure that there are no unintended results.

At the feedback workshop, each expert should be provided the opportunity to discuss his/her interpretations and evaluations with the other experts focusing on the technical basis for the assessment. The TFI is required to lead a constructive scientific debate, to look for areas of consensus, to resolve misunderstandings or different assumptions, and to assist the experts in making explicit assessments. The experts are expected to challenge and defend different interpretations. The TFI should focus the discussions on those aspects of the interpretations that are most important to the seismic hazard results (for example, the recurrence rate per-square-kilometer for the "host zone" containing the site). As in previous workshops, the feedback workshop is vital to ensuring interaction among the experts, which in turn is a key mechanism for the behavioral integration of the experts' assessment.

At the workshop, it should be emphasized that the assessments are preliminary and, following the workshop, the experts will be encouraged to make any changes that they feel are appropriate in light of the discussions and feedback that they have received. It should be emphasized by the TFI that there is no need nor any desire for the experts to agree with each other after having seen where their assessments stand relative to the other experts' assessments. Rather, it is important for each expert to understand the *technical reasons* why his evaluations fall where they do relative to the others and to be sure that these reasons make sense to him.

Finalize SSC Assessments and Integrate

Following the workshop, the experts should

finalize their interpretations and documentation. It is unlikely that another elicitation session will be required to do so, but it may. The TFI should take the responsibility for ensuring that the finalization is done properly and in a timely manner.

Two mechanisms are used to integrate the interpretations of the SSC experts: (1) behavioral integration related to interactions of the experts throughout the project, and (2) weighting of the expert interpretations. The first mechanism should be the primary mechanism for integrating SSC interpretations. Further, it is desirable for the TFI to be in a position to combine the SSC interpretations using equal weights, unless compelling reasons exist for unequal weighting. To integrate SSC expert interpretations properly, the TFI should have accomplished the following during the process:

- Selected highly qualified experts
- Established the commitment of each expert to the project and worked to motivate the expert throughout the project
- Disseminated a comprehensive and uniform data base to all of the experts
- Educated the experts in all aspects of seismic source characterization, including areas of limited expertise, and trained the experts in elicitation methodologies
- Facilitated interaction of the experts such that a free exchange of data and interpretations occurred as did scientific debate of all hypotheses
- Allowed for experts to decline answering certain elements of the assessment for which they did not feel qualified
- Provided feedback and sensitivity analyses to the experts, checked for unintentional errors, and facilitated discussion and challenge of preliminary interpretations
- Provided an opportunity for each expert to modify his assessments in light of feedback from the TFI and interactions with the other experts
- Obtained explicit agreement from each expert (or team) that the other experts' (teams') interpretations are understood and are valid alternative interpretations/representations

If the above recommended criteria have been met, the TFI should have created the proper conditions to apply equal weights to the SSC expert evaluations. However, it is recognized that certain circumstances may arise that would signal the need for unequal weights (see discussion in Section 3.3). If these or similar conditions arise, it is the responsibility of the TFI to communicate to the expert the concern and provide an opportunity for improved performance.

The goal in the procedures described in Section 3.3 and applicable to SSC is not to arbitrarily assign unequal weights to SSC evaluations. In fact, it is anticipated that in most cases the assignment of equal weights will be highly defensible. However, the procedures are designed to provide the TFI a mechanism to deal with particular "problem" circumstances where unequal weighting is more appropriate. As discussed, the procedure should only be applied in cases where unequal weights leads to significant differences in the hazard results. The calculated seismic hazard results for both the equal weights and unequal weights must be documented as part of the SSC report.

Peer Review The advantage of the TFI approach is that the active interaction of SSC experts provides for a de facto "peer review" of the technical substance of the SSC assessment. The TFI approach is designed specifically to encourage the presentation and technical challenge of hypotheses. In addition to this informal technical peer review, it is also recommended that an explicit process peer review be conducted. The focus of the peer review would be the implementation of all approaches and methods for the SSC assessment, including selecting the experts, compiling and distributing data, eliciting experts, etc. The recommended peer review procedure is a "participatory" peer review, whereby the peer reviewers interact with the TFI throughout the study and gain first-hand knowledge of the assessment through this interaction. A "late-stage" peer review, although

acceptable, may not be sufficient to identify any necessary mid-course corrections in the project. The results of the peer review should be documented and included with the final report.

4.4.3 Technical Integrator (TI) Approach

The TI approaches (Study Levels 1 through 3) for deriving SSC inputs have been applied, in the past, primarily to site-specific studies. The particular level used depends on the amount of published information available and/or the diversity of interpretations within the informed technical community about the seismic sources and seismicity relevant to the region of interest.

For parallelism with the TFI discussion (Section 4.4.2), the discussion below describes a Level 3 study. The basic elements of the TI approach to SSC, assuming a participatory peer review, are the following:

1. Select peer reviewers A panel of peer reviewers is selected by the sponsor in the same way that the group of experts was selected following the "expert" approach.

2. TI assembles all data bases and defines SSC procedures SSC-related data bases are compiled by the TI (for site-specific studies this may involve collecting new data). Methods and procedures are selected for the SSC analysis and reviewed by the peer reviewers.

3. TI conducts SSC analysis The TI conducts the source characterization with the requirement that a wide range of technical interpretations be represented and included.

4. Peer reviewers interact with TI The peer reviewers meet frequently with the TI to review and criticize the compiled data base, the procedures to be followed, and the interpretations being made on source characteristics.

5. Peer reviewers submit written review of preliminary analysis A draft SSC analysis, which fully documents all of the assumptions, methods, and assessments, is submitted to the peer review panel. The peer reviewers provide written comments on the draft report keeping in mind their charge to verify a full range of alternative interpretations.

6. Review comments are addressed by the TI Written responses to peer reviewer comments are prepared by the TI and changes are made to the analysis. The TI report is finalized. The peer review panel submits its endorsement of the final report.

Each of the above steps is discussed below.

1. Peer Reviewer Selection Because they will actively participate in reviewing the elements of the SSC analysis and will be asked to endorse the process and results, peer reviewers should have the same high qualifications that are required of the experts in the TFI approach. Likewise, in large complex studies, the peer reviewers should agree to a significant commitment of time and effort. Therefore, the peer reviewers should be selected by the sponsors in the same manner as discussed previously and, as a group, it is desirable that they be similarly balanced.

2. Data Bases and SSC Procedures SSC-related data bases should be compiled and formatted in the same manner as the TFI approach. The peer review panel can be instrumental in helping to define data needs and availability. In site-specific studies, various types of site data may exist and may need to reviewed and evaluated for their accuracy and pertinence to the SSC analysis. For example, geologic mapping and boreholes may exist in the site vicinity but may not have been gathered for purposes of earthquake evaluation. Their usefulness will need to be evaluated by the TI. In some circumstances for site-specific studies, certain critical data may need to be gathered in order to reduce uncertainties in the SSC and seismic hazard analysis. For example, if a fault has been identified in the immediate site vicinity, but no data are available to evaluate its activity, the peer reviewers can assist the sponsor in deciding whether new data collection is required.

The methods and procedures that will be used to characterize seismic sources and to quantify uncertainties should be identified. These can then

be reviewed by the peer reviewers prior to making assessments.

3. Conduct of SSC Analysis The TI takes the responsibility for conducting the source characterization assessment and is responsible for representing the informed technical community's full range of credible interpretations of the source characteristics. To achieve this, the TI must be familiar with and document all published or otherwise available interpretations of seismotectonics that might affect SSC; acknowledged experts should be contacted to obtain their views on possible interpretations; and any available site-specific data should be factored into the assessment. Throughout this process, the TI should be receiving input from the peer reviewers on possible alternative interpretations. The entire SSC analysis should be documented fully in a draft report, including the technical basis for all assessments and their uncertainties. The draft report should include sensitivity analyses that display the results of alternative SSC interpretations.

4. Interactions with Peer Reviewers The peer reviewers should meet frequently with the TI to review all aspects of the analysis. Their role is to inform the TI of available data and interpretations being made that might have an impact on the SSC analysis, to express their own interpretations as experts, to examine and suggest refinements to methods and procedures being followed by the TI, and to ensure that a wide range of technical interpretations is being represented. In reviewing the TI's work, the peer reviewers may recommend to the sponsors that significant new data be gathered or new analyses be undertaken that will strengthen the technical basis for the conclusions drawn. If such data are gathered, the peer reviewers may assist the TI in designing the data collection or analysis effort and in reviewing the results.

5. Peer Review of Draft Report Upon completion of draft documentation of the SSC analysis, a report is submitted to the peer reviewers. Here, standard peer review procedures should be followed to prepare written comments that relate to the process, assessments, and documentation prepared by the TI. The clear advantage of a participatory peer review process is that by this stage of the project, the peer reviewers are thoroughly familiar with the data bases available, the SSC procedures being used, and the technical basis for the assessments made. Likewise, the TI is aware of potential concerns that the peer reviewers may have and has sought to address them in the draft report. Much of the review will focus, then, on how effectively the documentation captures the assumptions made and analyses carried out. One or more meetings should be held with the peer reviewers and the TI to clarify the comments made on the report to minimize the need for subsequent review/revise cycles.

6. Resolution of Review Comments and Finalization of Report The normal resolution process for peer review comments should be followed, which includes addressing each comment, revising the report and, if necessary, the analysis, and preparing written documentation describing the manner in which each comment has been addressed. Again, interaction with the peer reviewers will have served to clarify the basis for each comment and provide for a clear resolution of remaining concerns. The final SSC report is then submitted to the sponsor.

The final step in the peer review approach is a written endorsement of the SSC study by the peer review panel. Their endorsement should extend to those aspects of the study where they were able to provide a direct, substantive review of the procedures and results.

4.4.4 Modifications to the Recommended Approaches

The recommended approaches (TFI and TI approaches) discussed above for incorporating SSC expert judgments are appropriate for most seismic hazard applications that SSHAC anticipates for the sponsors of the study. However, it is possible that the approaches may be modified to accommodate different needs. For example, the approaches discussed above are resource-intensive (Level 3 and 4 studies) and are, therefore, most appropriate for large-scale studies for critical facilities.

To provide some guidance for alternative applications, the following discussion considers how the approaches might be modified for more modest studies and for regional or site-specific analyses.

4.4.4.1 Resource-Intensive versus Modest Resources

This issue recognizes that some seismic hazard analyses will need to be conducted under limited budgets and time constraints (particularly, for less-critical facilities or for applications where there is interest in only higher probability levels). The trade-off in conducting more modest studies is not in the quality of the SSC assessment, but in the ability to verify the level of quality using multiple representatives (experts) from the SSC community. For example, a key attribute of the recommended approaches is assurance that a full, documented representation of the range of scientific interpretation is incorporated into the SSC assessment. Any seismic hazard analysis, of any scale, must strive to achieve this goal. However, the more modest SSC analyses will provide less assurance to an independent third-party reviewer that, in fact, the goal has been achieved.

In general, the levels of analysis discussed in Section 3.1 (Levels 1 through 4) scale with resources available. Therefore, a TFI approach, which is Level 4, requires considerable resources in most cases. However, we here discuss some of the ways that both a TFI and a TI approach might be modified to account for modest resources.

TFI Approach

- Reduction in the number of experts who will conduct the analysis (say 3 to 4 experts).
- Reduction in the number of expert interactions. The number of workshops could be reduced, or eliminated altogether and replaced with smaller working meetings to review data bases, interpretations, and preliminary results
- Reliance on readily available data bases. Published and readily retrievable data bases could be made available to the experts (e.g., existing earthquake catalogs). Data processing might be reduced.

The reduction in scope discussed above would result in less of the desirable attributes of the TFI approach itself: use of multiple experts to represent the technical community, intensive interaction of the experts to resolve unintended disagreements and review/challenge interpretations, and formal elicitation of a representative sample of the informed technical community.

The focus of the TFI approach would remain one of using the judgments of experts directly to characterize seismic sources and to quantify the uncertainties.

TI Approach

- Reduction in the number of peer reviewers and less participation. A "late-stage" peer review process could be adopted whereby the reviewers do not interact with the TI during the course of the study, but review and comment on the draft report.
- No new data collection or data processing. The TI would base the SSC assessment on available data only and would not gather new data. Data processing might be reduced.

Other reductions in resources could, perhaps, be accomplished in the analysis itself. For example, the number of sensitivity analyses could be reduced. In no case, however, can the quantification or incorporation of uncertainties be significantly reduced in scope. Even the more modest seismic hazard studies must attempt to incorporate uncertainties.

4.4.4.2 Regional versus Site-Specific Studies

The application of the seismic hazard analysis can be for regional seismic hazard assessments (e.g., contours of hazard levels over regional scales) or for site-specific applications at particular facilities. The recommended approaches are appropriate for either application. However, experience leads us to provide preferences for particular applications.

For the following reasons, large, regional studies (say, of the eastern U.S.) might best be conducted using the TFI approach; site-specific studies might best be conducted using the TI approach. In terms of source characterization, seismic sources are commonly identified based on regional data sets (seismicity maps, fault maps, tectonic maps, geophysical maps). Because local small faults and sources do not have much significance in regional hazard results, detailed characterization of small-scale features is usually not attempted. Experts for regional studies, therefore, are required to have a strong knowledge of seismotectonics over geographically extensive areas. Although they may have detailed knowledge of specific local areas, their regional knowledge is most important. In contrast, experience with site-specific hazard studies has shown that the hazard is often dominated by a few local sources. Further, a few key characteristics of these local sources are often the most important (e.g., fault slip rate, geometry). As a result, experts for site-specific studies are required to have a detailed knowledge of the local seismotectonics and otherwise "minor" local sources. The need for this local knowledge limits the number of possible experts. Commonly, a TI takes the responsibility of assembling all of the available data sets and making the assessments. In doing so, the TI can draw on site-specific knowledge and expertise from, for example, scientists who have worked in the site area. The peer review of these TI assessments should entail a review of both the methodologies used by the TI as well as the use that has been made of local site-specific data sets. Hence, even if the peer reviewers do not have the site-specific expertise of the TI, they can contribute to the overall quality of the hazard analysis.

In past studies, there have been significant differences in the amount of new data that have been collected to provide a basis for the SSC assessment, ranging from no new data collection to extensive programs, including geologic field studies (designed to focus on significant seismic sources and to reduce uncertainties in their characterization). Differences in data collection appear to be tied to the regional versus site-specific issue. Regional hazard studies typically have a limited program of new data collection, although processing of existing data can require significant effort. Conversely, site-specific SSC studies usually include at least a limited, focused program of new data collection. For some site-specific studies (e.g., Diablo Canyon Long-Term Seismic Program), "scoping studies" (preliminary hazard analyses based on existing knowledge) have been carried out with the specific purpose of identifying the data collection activities that will have the most importance to the hazard results.

Another circumstance to consider is the case where a site-specific study is being conducted within a region for which a regional seismic hazard study has already been completed. An example might be a facility site in eastern Oregon, where a regional seismic hazard map has recently been completed for the state. Because the focus of the regional map was, in fact, on the regional variation of seismic sources, it is unlikely that local faults and sources have been considered. In these cases, either the TFI or the TI approach could be used first, to review the basis and applicability of the regional sources identified for the regional study, and second, to identify and characterize the local seismic sources of importance to the site. If the results of the site-specific study differ significantly in seismic sources or characteristics (and thereby hazard values) than the regional study, it is important to document the reasons for those differences. The differences may be due to an evolution in understanding of earthquake sources (e.g., the identification of a previously unknown source) or to differences in the manner in which uncertainties have been treated. A goal in seismic hazard analysis should be, over the long term, stability in estimates of seismic hazard both regionally and locally. As the science evolves and new findings are incorporated into the hazard analysis, the differences in the resulting hazard values should be understood and documented.

5. METHODOLOGY FOR ESTIMATING GROUND MOTIONS ON ROCK

5.1 Introduction

This chapter discusses the estimation of ground motions for use in probabilistic seismic hazard analysis. The chapter is divided into two major portions. The first portion comprises Section 5.2 through 5.5 and discusses the ground-motion measures considered in PSHA, the explanatory variables in current use, the various methods for prediction of ground motions, and the treatment of uncertainty. The recommendations from this portion are summarized in Section 5.6. The second portion comprises Section 5.7 and discusses issues of expert elicitation regarding ground motion models. The calculation of site response is not discussed, in spite of its importance, because this issue is outside the scope of the SSHAC study.

Computations of seismic hazard require a specification of earthquake occurrence and size, discussed in the previous chapter, and the ground shaking from earthquakes as a function of, at least, magnitude and distance. Unlike most deterministic studies, PSHA requires the ground motion as a continuous function of magnitude and distance. Furthermore, in general it is not enough to give the expected median value of the ground shaking; the aleatory uncertainty of the shaking, as well as the epistemic uncertainty in the median value, must also be specified. This chapter discusses the estimation of ground shaking for use in PSHA.

In many situations, there are not enough recordings of ground motion to allow a direct empirical specification of ground shaking. As a result, estimates are based on a variety of methods, using different assumptions and models that are calibrated and verified using various data sets. This can lead to widely differing motions for a particular magnitude and distance. In the face of such diversity, an important task in our study is to recommend a methodology for obtaining ground motions that adequately capture the uncertainty in the estimates and is defensible in a regulatory arena.

The Committee recognized at an early stage that it could not recommend a particular model or even a particular class of models. It is very likely that the models will change with time, as new data become available and as methods are refined to account better for existing data and for improvements in our understanding of how the earth works. We felt it more valuable to recommend *procedures* for obtaining the ground motions, procedures that will be as applicable ten years from now as they are today.

The scope of this chapter includes ground-motion estimation on hard-rock sites throughout North America, for distances, magnitudes, and frequencies of relevance for engineering design and structural response. Generally the magnitudes of interest will be above 5.0, the distances will be up to several hundred kilometers, and frequencies of response will range from 0.5 Hz to several tens of Hz.

Although SSHAC considers site effects to be one of the most important factors affecting the amplitudes and durations of ground motions (aside from the earthquake magnitude, of course), the project scope was restricted to motions on rock sites. We assume that site-specific applications will consider the expected modifications of the ground motions for the local site conditions. Local site conditions include the geologic materials below the surface as well as topographic irregularities of the ground near the site. The model for site effects should consider soil nonlinearity, if appropriate.

In discussing ground-motion prediction, it is common to divide North America into two regions—western and eastern. It is also important to divide North America on the basis of the availability of ground-motion data, because this determines the preferred methods for ground-motion prediction. In this sense, what is often loosely referred to as western North America is actually coastal California, from which most strong-motion data have been obtained. Ground-motion prediction in other parts of California or western North America suffers from the same (or

even a greater) lack of data as does eastern North America (in fact, a large number of useful ground-motion recordings have been obtained from southeastern Canada and the northeastern United States). Predictions of ground motions in regions lacking sufficient data for a direct empirical estimation must be based on similar methods, using different parameter values to represent differences among geographic regions. For this reason, we emphasize methods more than regions in our discussions.

We have organized this chapter along traditional lines, first discussing measures of ground motion, followed by sections on explanatory variables and methods for obtaining ground motions. We have not attempted to give a comprehensive and in-depth treatment of these subjects. This is not a textbook for predicting strong ground motion; we have provided references for those interested in the details of the prediction methods.

The heart of SSHAC's ground-motion contribution to PSHA is given in the final two sections of this chapter. Section 5.5 discusses the definition and estimation of uncertainty in ground-motion predictions, and section 5.7 contains SSHAC's recommendations for obtaining estimates of strong ground motion for PSHA .

5.2 Ground-Motion Measures

Although a time series of the ground shaking is needed for an exact analysis of nonlinear, dynamic behavior of structures or soil deposits, most PSHA studies characterize ground shaking in terms of a few ground-motion measures. The most common measures are peak ground acceleration (*PGA*) and a few (typically 6) ordinates of the response spectrum. This characterization is sufficient for most applications.

Peak ground acceleration is defined as the maximum absolute amplitude of a ground acceleration time series. It is easy to obtain from analog records and is used to define lateral forces and shear stresses in equivalent-static-force procedures (e.g., those specified in building codes) and liquefaction analyses. Being controlled by the highest frequency content in the spectrum, however, *PGA* is very sensitive to processes that can alter the high-frequency content, such as local geologic conditions and instrument response, and furthermore is not easily related to any particular range of ground-motion frequencies; in any earthquake, the *PGA* can be controlled by different frequencies at different distances from the earthquake. In addition, the frequencies that dominate the peak acceleration in a particular record are often not in the range of those most important for structural response. Peak accelerations in excess of 5 *g* have been measured from rockbursts in mines, but these motions were dominated by frequencies near 400 Hz. For these reasons we recommend that *PGA* not be used to determine design spectra for most applications. A much more useful measure of ground motion is the response spectrum.

Response spectra describe the response of a single-degree-of-freedom damped elastic oscillator to ground shaking. A number of different measures have been used, referred to by a confusing variety of symbols and terms. The one most commonly used for PSHA is the peak spectral acceleration, *PSA*, defined as:

$$\mathrm{PSA} = \left(\frac{2\pi}{\mathrm{T}}\right)^2 \mathrm{S_d} \qquad (5.1)$$

where S_d is the maximum displacement of the mass of an elastic, viscously-damped, single-degree-of-freedom oscillator with undamped natural period T, relative to its point of attachment to the ground. In most applications, the damping is taken to be 5 percent. When PSA is plotted as a function of frequency or period, the result is a response spectrum. Because the response of many structures can be well-approximated by that of a single-degree-of-freedom simple harmonic, damped oscillator, the characterization of ground shaking as a response spectrum is immediately useful. Once a response spectrum is defined, the maximum acceleration, and thus the force, to which a structure is subjected is easily determined by scaling the appropriate value off the spectrum. Because of its simplicity, the response spectrum has been universally adopted as the standard method of defining earthquake motions for purposes of performing dynamic analyses of

simple elastic and inelastic structures. Furthermore, we anticipate that future editions of building codes will use response spectral ordinates at a few selected periods, rather than peak acceleration, as a basis for seismic zonation (Algermissen et al. 1991).

In spite of their usefulness, response spectra have some limitations. First, they provide the response of a linear oscillator. Studies of nonlinear response require a more complex representation of the ground shaking (e.g., Kennedy et al. 1984; Sewell 1988; Krawinkler et al. 1991; de Bejar and Ganapathi 1992). Second, although PSA can be calculated for any frequency, a non-zero *PSA* does not imply ground-motion energy at the frequency of the oscillator. For example, if shaken by ground motion with frequencies no higher than, say, 5 Hz, the response of a 100-Hz oscillator will simply reproduce the ground acceleration. Finally, because of the process by which they are constructed, response spectra do not have the same properties as Fourier spectra. In particular, the ratio of two response spectra is not the same thing as the ratio of Fourier spectra. For example, ratios of response spectra from the same event do not cancel the effect of the source.

The ground shakes in both horizontal and vertical directions, and in addition spatial variations can produce rotations of the ground. Generally, *PSA* are only computed for the horizontal ground shaking. It is not sufficient, however, simply to specify that the *PSA* is for horizontal shaking, since the shaking can be in two spatial directions. Usually the specification is for either the geometric mean or the larger of two horizontal orthogonal components of motion placed randomly with respect to the orientation of the fault that produced the motion. It is important to be specific about the particular definition, for there are systematic differences in the motion between various definitions (e.g., Boore and Joyner 1988).

Less emphasis is usually placed on the vertical component of motion. The vertical component is often estimated from the horizontal component using a rule-of-thumb, for example, where the vertical is about 2/3 of the horizontal. Such rules should be used with caution, however, for the actual ratio may depend on the frequency of motion, the local site conditions, the focal mechanism, and the distance from the event (e.g., Atkinson 1993a; EPRI 1993). If important, we recommend that vertical motions be obtained from independent analyses, in the same manner as for the horizontal motions (e.g., Abrahamson and Litehiser 1989).

5.3 Explanatory Variables in Ground Motion Models

5.3.1 Introduction

This section discusses the quantities that serve as input to ground-motion attenuation models. The discussion focuses on the current state of practice and on anticipated developments over the next five years. These explanatory variables in ground-motion attenuation equations fall into three general categories, as follows: (1) size and other characteristics of the earthquake (typically magnitude), (2) location of the site relative to the earthquake (typically distance), and (3) site characteristics.

5.3.2 Background

In general, the addition of a new explanatory variable X in the ground-motion model is justified from the point of view of seismic hazard analysis if the following three conditions are satisfied:

1. Introduction of X in the ground-motion model results in a significant reduction in the scatter of the ground-motion residuals (observed minus predicted amplitudes), as measured, for example, by a 10 percent reduction in the residual standard deviation.

2. There is the ability to characterize the probability distribution of parameter X for future earthquakes affecting a given site.

3. The probability distribution of parameter X in future earthquakes affecting a given site must be significantly different from the distribution in the sample data used in the development of the ground-motion model. If the two distributions are similar, the explicit introduction of parameter X in the ground-motion model and in the seismic-hazard

integration will not have an effect on the hazard. This condition is somewhat less critical than the other two. One may be justified in introducing parameter X because it makes the model more robust, or because it may allow for future site-specific updating of the hazard if the site-specific probability distribution of X becomes better known.

There is a relationship between the number and type of explanatory variables in an attenuation equation and the associated uncertainties (both aleatory and epistemic). This relationship will be discussed in Section 5.5.

5.3.3 Characterization of the Earthquake Source

5.3.3.1 Magnitude

Magnitude is the most commonly used measure of earthquake size for the purpose of seismic hazard analysis. There are a large number of magnitude scales in use. It is imperative that the ground-motion attenuation equations and the source characterization use the same magnitude scale.

Most magnitude scales are instrumental. Magnitude is calculated from the peak amplitudes and distances from the earthquake sources to seismographs of a certain type that recorded the earthquake (the process is analogous to applying an attenuation function in reverse, solving for magnitude given amplitude and distance). Moment magnitude, unlike instrumental magnitudes, has the advantage that it is related to a well-defined physical characteristic of the earthquake source (i.e., the seismic moment). In practice, seismic moment is not observed directly and, like instrumental magnitudes, it must be calculated from indirect observations (e.g., seismograph recordings, geologic or geodetic measurements), and the calculation of seismic moment from these observations often requires assumptions about seismological models and their parameters.

Magnitudes for pre-instrumental earthquakes are determined from macro-seismic measurements such as epicentral intensity, felt area, or the extent of liquefaction, using empirically derived conversions. Issues of magnitude conversion are discussed in numerous references (see Johnston et al. 1993; EPRI 1986).

Current practice as to the choice of magnitude scale for seismic hazard analysis is different for different regions of North America and for different earthquake sizes.

- East of the Rocky Mountains, Nuttli's (1973) m_{bLg} (also called m_{Lg} or m_N) magnitude is the most commonly used magnitude (e.g., this is the primary magnitude in the EPRI 1986, catalog). This is the magnitude in current use by seismograph networks in the region. Pre-instrumental earthquakes have been converted to m_{bLg} using empirically-based relations that use intensity as the fundamental observable (e.g., EPRI 1986). The choice of alternative conversion relations for pre-instrumental earthquakes is one potential source of differences among seismic hazard studies (see Toro et al. 1992). There are several deficiencies in current procedures for calculating and reporting m_{bLg}. For instance, no distinction is made in some catalogs between m_{bLg} and teleseismic m_b. Also, no account is taken of the instrument types (for instance, American stations typically calculate m_{bLg} using short-period WWSSN seismographs, which peak near 1 Hz, while Canadian stations use ECTN seismographs, which have a broader bandwidth). There are also variations in observatory practice, as some stations use Nuttli's (1973) equation, while others use variants of that equation. All these factors lead to moderate but systematic regional biases in m_{bLg} estimates.

- For western North America, the local magnitude M_L or an approximation to that magnitude has been used since its introduction by Richter in 1935 (Richter 1935). In the last 15 years, however, Hanks and Kanamori's (1979) moment magnitude **M** has been commonly used and is now the preferred magnitude for moderate and large earthquakes. The moment magnitude generally correlates well with other magnitudes over limited ranges of earthquake size. For magnitudes between about 3 and 6,

M is approximately equal to the M_L magnitude used by seismograph networks in California (Hanks and Boore 1984). For large, but not great, earthquakes, **M** is approximately equal to the surface-wave magnitude M_S. The issue of determining **M** for historical earthquakes is of less importance in the west because the historical catalog is shorter than in the east, and because there are more instrumental data as a result of higher activity rates.

There is a trend towards the use of **M** in central and eastern North America. This trend is motivated by several factors, as follows: (1) the deficiencies in the calculation and reporting of m_{bLg} mentioned above, (2) the preference for predicting ground motions using the stochastic and physical ground-motion models, in which the seismic moment is a fundamental model parameter (e.g., Atkinson and Boore 1995; EPRI 1993)[1], and (3) the benefits of using one magnitude scale for all of North America, thereby eliminating a non-physical distinction between east and west.

In addition, moment magnitude is the magnitude used for the GSHAP (Global Seismic Hazard Analysis Program 1993) worldwide seismic-hazard study.

The conversion to moment magnitude for seismic hazard studies is justified if the overall uncertainty in the calculated seismic hazard is reduced; this may be verified quantitatively through statistical analysis of events for which both mblg and **M** are available. This uncertainty involves both uncertainty in the conversion of the earthquake catalog to moment magnitude (especially for large earthquakes with no instrumental data) and in the attenuation equations. Regarding the first issue, recent work on the characterization of large intra-plate earthquakes (Johnston 1995a,b,c, in press) has provided estimates of moment magnitude for the tectonically-stable region of North American earthquakes above **M** 3.9 for pre-instrumentally recorded earthquakes (for which only intensity data are available) and above **M** 3.5 for instrumentally-recorded events. The scatter in these relationships is generally lower than for conversions to m_{bLg}. Regarding the second issue, attenuation functions in terms of m_{bLg} and of **M** have comparable scatter (as characterized by the residual standard deviation) for high-frequency ground-motion measures (i.e., f > 2.5 Hz), and the M-based equations have lower scatter for low-frequency ground-motion measures (EPRI 1993; Atkinson 1995). In conclusion, there are significant advantages in converting to moment magnitude as the measure of earthquake size for seismic-hazard analysis and we recommend a gradual transition towards the use of **M** for seismic hazard studies in all regions of the United States. Appendix C contains a more complete discussion of these issues.

Another interesting alternative to the use of m_{bLg} is the high-frequency magnitude m recently proposed by Atkinson and Hanks (1995). The main advantages of this magnitude scale are that it is (1) more directly related to high-frequency ground motions than are the other magnitude scales, and (2) it correlates very well with felt area, thus allowing reliable estimation of magnitude for large, pre-instrumental events. Although we consider this scale to have much promise for use in PSHA, we cannot recommend its use at this time; it must be better understood and accepted before it is used as the magnitude scale for seismic hazard analysis.

5.3.3.2 Other Source Characteristics

Another source characteristic that affects ground motion is the tectonic regime where the

[1] Physical and stochastic models use seismic moment as the basic measures of the size of the earthquake. Thus, in order to predict ground motions for a given m_{bLg} using these models, one needs a relationship for seismic moment given m_{bLg}. Often, these relationships are obtained using the model itself, by predicting the amplitude recorded at a hypothetical seismograph at a certain reference distance, and then applying Nuttli's (1973) equation to calculate m_{bLg} as a function of amplitude and distance. The resulting relationship is sensitive to the assumptions of the model (particularly Q), the choice of seismograph, and the choice of reference distance. An alternative is to derive the relationships between m_{bLg} and seismic moment empirically. Unfortunately, these two approaches can lead to large differences for large earthquakes. The empirical data can be fit well by a linear relation in the range for which data are available. The model-based relationships agree with the empirical relationships in this region, but predict curvature for larger earthquakes. Sufficient data are lacking to resolve the differences at the larger magnitudes. These differences can translate into large differences in predicted ground motions.

earthquake occurs (i.e., intraplate, plate margin, or subduction zone). This characteristic is not typically included explicitly (as a parameter) in attenuation equations because most attenuation equations are applicable to a single tectonic regime. Instead, tectonic regime is used implicitly as an explanatory variable by the selection of the attenuation equations applicable to the region of interest. If a site is affected by earthquakes from different tectonic regimes, it may be required to use different attenuation equations for the source zones associated with the different types of earthquakes.

Other source characteristics that affect ground-motion amplitudes are focal mechanism (strike slip, reverse, or normal), and source depth. Focal mechanism is typically not used as an explanatory variable, although it has been used in some attenuation functions for California (e.g., Campbell and Bozorgnia 1994; Sadigh 1993). Depth is seldom included explicitly as an explanatory variable.

5.3.4 Characterization of Site Location Relative to the Earthquake

5.3.4.1 Distance

Figure 5-1 illustrates the most common definitions of distance used in attenuation functions and in seismic hazard analysis.

For small and moderate earthquakes, the dimensions of the earthquake rupture are negligible compared to the distance from the earthquake to the site (except, perhaps, for earthquakes in the host source zone). In this case, two definitions may be used: hypocentral (or focal) distance, and epicentral distance. These two definitions of distance are consistent with the use of areal source zones in seismic hazard analysis. Consistency must also be maintained in the treatment of depth. If the attenuation function uses epicentral distance, the areal source zones must be specified as having zero depth. If the attenuation function uses hypocentral distance, the seismic source must have a non-zero depth (or, preferably, a probability distribution of depth).

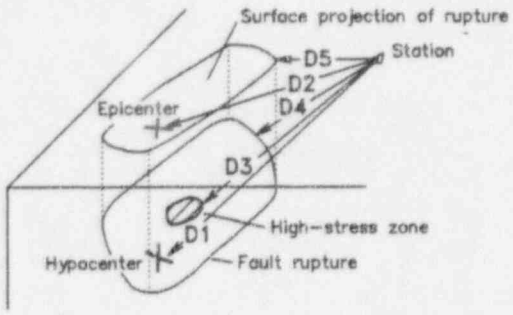

Distance Measures (from recording station)

D1 – Hypocentral
D2 – Epicentral
D3 – Closest distance to high-stress zone
D4 – Closest distance to fault rupture
D5 – Closest distance to surface projection of rupture

Figure 5-1 Diagram illustrating different distance measures used in predictive relationships (from Shakal and Bernreuter 1981).

If the dimensions of the earthquake rupture are of the order of ten kilometers or more, point-source idealization of the seismic-energy release may become inappropriate. In these situations, some form of closest distance to the rupture should be used. Figure 5-1 illustrates the shortcomings of the hypocentral and epicentral definitions of distance when the rupture is long. Of the three definitions of distance that consider the spatial extent of the rupture, the distance to the slipped fault (also called distance to the seismogenic rupture), and the closest horizontal distance from the station to the point on the earth's surface that lies directly over the rupture are commonly used. These definitions of distance are consistent with the use of fault seismic sources with extended ruptures. Consistency must also be maintained between the definition of distance and the geometric representation of earthquake occurrences. Attenuation equations that use distance to the surface projection of the rupture require (as a minimum) line models of the fault trace. The other definitions require three-dimensional models of the fault plane.

5.3.4.2 Other Characteristics

At least one seismic hazard study has considered the location of the site relative to an earthquake with a reverse focal mechanism, under the

assumption that ground motions are different for the up-thrown and down-thrown blocks (PG&E 1988). In addition, some attenuation equations under development predict different amplitudes for the up-thrown and down-thrown blocks (Somerville and Abrahamson 1995, in press).

Directivity is the theoretical tendency for higher ground-motion amplitudes when the rupture propagates toward the site. Some empirical studies show evidence for directivity effects (e.g., Boatwright and Boore 1982; Campbell 1987), but these effects are usually not seen at the frequencies of interest in this study (i.e., $f > 0.5$ Hz). Directivity effects are often obvious in data from large earthquakes at lower frequencies (e.g., Benioff 1955; Gutenberg 1955; Kanamori et al. 1992; Somerville and Graves 1993; Velasco et al. 1994) and may be important in seismic-hazard studies for low-frequency structures such as base-isolated structures, suspension bridges, or tall buildings. Attenuation equations in current use do not include directivity effects as an explanatory variable, although to the extent that such effects are in the data, they will be included implicitly as scatter in empirically-based equations. Theoretical models often show directivity effects, and simplified theoretical models, such as the commonly-used stochastic model, might incorporate directivity implicitly in deriving model parameters from empirical data (e.g., Boore and Joyner 1989).

5.3.5 Characterization of Site Response

Site effects are best treated on a site-specific basis, because these effects may substantially change the amplitude and frequency content of ground motion and because data on the dynamic characteristics of the site are always gathered for important facilities. Guidance on methods for site-specific evaluation of site effects is beyond the scope of this report; the reader is referred to EPRI (1993), Martin (1994), and other engineering literature on site-response calculations, including the treatment of soil nonlinearity. Nevertheless, a short discussion of the generic characterizations of site response is included here because these may be of use for preliminary studies or for studies involving low-risk facilities.

Some attenuation equations include terms describing site conditions (see Joyner and Boore 1988, for a review). Initially those terms consisted of dichotomous variables (rock or soil; e.g., McGuire 1978; Joyner and Boore 1981). Other studies distinguish among different soil depths by considering the depth to basement rock (Trifunac and Lee 1979; Campbell 1987). Boore et al. (1993 1994a) consider the combined effect of soil depth and soil impedance by using the average shear-wave velocity over the top 30 meters of the soil column.[2]

Other studies have developed amplification factors for various soil types and depth categories (Bernreuter et al. 1989; Boore and Joyner 1991; EPRI 1993). These amplification factors are used to modify the rock attenuation functions or the rock hazard results.

Note that it is important to be precise about what constitutes a rock site. For example, for the SSHAC elicitation of ground motion (Appendix B), a rock site was defined to be one whose time-weighted shear velocity in the upper 30 m is 2800 m/s. A hard-rock site such as this may be appropriate for glaciated portions of eastern North America, but for many other sites what is commonly taken for rock will have much lower shear velocities in the upper 30 m. In such cases, careful consideration of differences in local rock velocities must considered before importing results from other regions; a "site effect" may have to be developed for the rock. For example, a sample of California sites that were classified as rock (based on their geological description) had a geometric-average shear-wave velocity of 650 m/s.

5.3.6 Introduction of Other Explanatory Variables into the Hazard Analysis Calculations

Magnitude and distance are the only ground-motion explanatory variables used in a majority of seismic hazard studies and the available calculational methods are designed to integrate over the aleatory (i.e., random) distribution of

[2] Actually, the quantity used is the harmonic or "time-weighted" average of the velocity, i.e., 30 m divided by the travel time through the upper 30 m of the profile.

earthquake magnitudes and locations. These calculational methods also consider epistemic uncertainty (i.e., ignorance) about the true form of these distributions.

Additional explanatory variables have been used in ground-motion models for recent seismic hazard studies. This practice will become more widespread as more data become available and ground-motion models become more sophisticated. If the value of an explanatory variable applicable to the site (or to a site-source pair) is not known (either because the value is anticipated to vary from event to event or because of incomplete knowledge), it is important to consider that uncertainty explicitly. If the value is expected to vary from event to event (i.e., aleatory uncertainty), one should integrate over the values of the explanatory variable in the calculation of the exceedance probabilities (as is done with magnitude and distance). Because the standard calculational methods for seismic hazard analysis do not integrate over variables other than magnitude and location, it is usually easier to perform this integration prior to the hazard analysis. This integration will change the median prediction of the ground motion and the distribution about that median. If the uncertainty about the explanatory variable is epistemic, it is easier to incorporate it in the conventional seismic-hazard analysis.

5.4 Methods for Predicting Ground Motions

This section contains a brief discussion of the various common components that must be included in ground-motion prediction models (other than those based strictly on data). Any particular model can be built from the various elements that are cascaded together to make the predictions.

A matter of terminology must be cleared up now. By "methods" we mean a general class or way of predicting ground motions. A particular application of a method to derive ground motions, either explicitly in the form of a table of values or equations, or implicitly as a procedure or algorithm from which ground motions can be computed, is referred to as a "model." These "models" often are associated with the names of the authors who derived the model. For example, Toro and McGuire (1987) and Boore and Atkinson (1987) are two models using the stochastic method.

By focusing on the components of models, we have removed the temptation to judge, rank, recommend, or advise against particular models. This is more properly the task of the elicitation process discussed in Section 5.7. References to specific models and some discussion of their advantages and disadvantages can be found in the discussions of the workshops held during the course of this project (Appendices A and B).

We first start with a description of methods that rely on data. Next, we discuss aspects of methods that must rely on a mix of theory and data; models based on these methods are currently the most commonly used procedures for predictions of ground motions outside of coastal California. We conclude with remarks regarding the use of scaling spectral shapes to obtain ground motions. This last subsection is the only one in this section containing any significant recommendations—in general we recommend that the use of scaled spectral shapes be avoided. Even though the spectral shapes are usually determined from analysis of empirical data, we placed this section last because the peak acceleration needed to scale the shapes can come from any method—empirical or theoretical.

For those readers seeking more details about various methods, we recommend consulting the original research papers, as well as reviews such as Joyner and Boore (1988), Atkinson and Boore (1990), Reiter (1990), and Boore et al. (1994b). In addition, the quadrennial reviews of strong motion seismology published by the American Geophysical Union can be very useful (e.g., Joyner 1987; Anderson 1991).

5.4.1 Empirical Methods

This topic divides naturally into those methods that use instrumental data and those that use intensity data. We discuss the use of instrumental data first.

Methods that Use Recorded Ground-Motions

In some site-specific and deterministic applications, ground motions can be determined directly by choosing a suite of motions from earthquakes of similar size, fault type, and distance from the site. If the data are available, nothing more need be done. This use of observed data is not particularly relevant for PSHA, which requires the prediction of ground motions for a continuum of magnitudes and distances. In this situation, the obvious choice is to use regression analysis to fit a functional form to the set of ground-motion recordings. Details about this method can be found in a number of places, including Boore and Joyner (1972), Campbell (1985), Joyner and Boore(1988), Boore et al. (1993 1994), Sadigh (1993), Campbell (1993), and Campbell and Bozorgnia (1994). Even with similar data available, differences can arise in the results because of different data winnowing, different choices for the explanatory variables, and different assumptions regarding the functional form. The equations are directly useful when the PSHA is being performed in the region for which the data were obtained, and the results are also useful as a means for checking and calibrating theoretical methods. The method suffers from several weaknesses, all related to the lack of data for various magnitudes and distances. It is well known that few data are available at close distances for large events, but perhaps less well known is the limitation at distances beyond about 100 km. The ground motions at these distances are small enough that not all operational instruments are triggered, even for large events. This can lead to biases in the regression results and uncertainty in the form of the attenuation equations for the greater distances. For use in PSHA, the equations must be capable of predicting motions beyond the distance at which the bias might appear, and therefore the attenuation equations determined strictly from empirical data may have to be extended using theoretical models or data from small earthquakes recorded on sensitive seismological networks.

Methods that Use Intensity Data

Because it is often the only information related to the ground shaking from large earthquakes in regions characterized by low seismicity, seismic intensity has been used in the past to predict ground shaking in future earthquakes (see, e.g., Trifunac and Brady 1975; Veneziano 1988). The present consensus that emerged from the ground-motion workshops (see the Appendices) is that intensity should no longer be used as the principal means for obtaining ground motions. The principal weaknesses are that intensity is a fairly crude, qualitative measure of ground shaking, and the correlations between intensity and ground shaking, needed to derive equations for spectral or peak accelerations, are very poor. This leads to significant uncertainty in the motions. Also, these relationships between intensity and instrumental ground motion are region-dependent, magnitude-dependent, and distance-dependent. In addition, there are pitfalls in the process of substituting one regression into another, as required in order to construct an attenuation equation for PGA or response spectra from intensity attenuation equations (e.g., Cornell, Banon, and Shakal 1977; Veneziano 1988; Risk Engineering 1991). These limitations might be overcome to some extent if ground motion and earthquake magnitude never entered the picture. This would be the case if seismicity were expressed in terms of epicentral intensity and if the product of the analysis were hazard curves for various intensities. This scenario, however, is not applicable to the PSHA studies of interest to SSHAC because a characterization of seismic hazard in terms of intensity is of no practical use in setting design levels or in seismic PRA studies. This scenario is applicable to earthquake loss studies, which routinely use loss functions in terms of intensity.

Intensity data and intensity attenuation equations are useful, however, as consistency checks for the predictions (especially for large magnitudes) obtained from attenuation equations based on modeling and/or instrumental data.

The use of intensity has recently experienced a renewal of interest as a means for estimating seismic moment (e.g., Hanks and Johnston 1992; Johnston 1995b; Bollinger et al. 1993) and high-frequency Fourier spectral levels of the ground shaking (Atkinson and Hanks 1995); see also section 5.3.3.1. These approaches make use of

data from events for which both intensity observations and instrumental magnitudes are available. Thus, this process ensures consistency with relationships such as that of Nuttli and Herrmann (1984). These estimates of source strength can then be used with theoretical models of the ground shaking to obtain predictions of ground shaking that potentially have less uncertainty than those produced by the direct use of intensity. We think that this very promising use of intensity has an important role in PSHA in many regions of North America.

5.4.2 The Components of Non-Empirical Methods

Introduction

For most applications, sufficient data are lacking to be able to use the purely empirical methods. Instead, ground motions are predicted by methods that generally combine theoretical and empirical aspects: the theory gives functional forms with parameters that are determined from data, if possible, or are specified by analogy with other regions or from experience.

Hind- vs. Fore-Sight Prediction. Before continuing, we point out that for the ground-motion component of PSHA, the fundamental task is to predict the statistical distribution of ground motion for future earthquakes as a function of magnitude and distance (at the least, this would be the mean and the standard deviation of the ground-motion measure). Many seismological studies are focused on hind-sight predictions of ground shaking in individual earthquakes for which records of ground motion have been obtained. The purpose of these studies is usually to infer the details of the seismic source—the geometry of the source and the distribution of slip across the fault as a function of space and time. When done for enough earthquakes, the results of such studies can be used to develop statistical distributions of the source properties that might be used for the type of ground-motion predictions needed by PSHA. Furthermore, the methods used to construct synthetic seismograms can be used in a forward sense to predict motions from future earthquakes, but to be useful for future earthquakes, many simulations are required to estimate the parameters of the statistical distribution of ground motion. This can be a computationally intensive exercise. This approach will become more viable as data from more earthquakes are accumulated and as computing power increases.

Pieces of the Puzzle. The methods usually break the task of estimating ground motions into three pieces: the source, the path, and the site (the division between the latter two is somewhat artificial—the site is usually that part of the path within a few kilometers of the point at which ground motion is predicted or observed). Later in this subsection we discuss the various ways that these pieces are treated.

The Issue of Complexity. It is clear from looking at observed ground shaking that the motions are usually chaotic, particularly at high frequencies. Simple models in which a fault with uniform slip is embedded in a constant-velocity half space do not produce enough complexity in the motions (although such models have been used for many years in studies of motions at periods much longer than of concern in engineering). A key issue is how the various methods incorporate the complexity and randomness in the motions. Some methods attempt to model the actual complexities of the earth, while others might be classified as phenomenological models with functional forms that are guided by insights from physical models and parameters that can be adjusted to fit data. A fundamental precept of the latter methods, whether or not explicitly stated, is that the earth's dynamics and structure are too complicated to model deterministically, particularly for a future event.

Low vs. High Frequency. Before embarking on a discussion of the pieces of the puzzle, we remind the reader that the applicability and necessity of a method are often related to the frequency of the motions that are being predicted. As mentioned earlier, simplified models of the source and path are adequate if very long periods are being modeled, motions whose wavelengths are much greater than the size of the earthquake source and most of the earth's heterogeneities. This frequency regime is of no interest for PSHA. On the other hand, motions with periods of several seconds

may be of interest, and for these motions deterministic models of the propagation path can be adequate, although even for this case unknowable complexities of the source may require a probabilistic treatment of the source.

Source

Introduction. The seismic source is described fundamentally by the specification of the slip across the fault plane as a function of space and time. This information is never available for future earthquakes, and therefore the methods for treating the source use various approximations to obtain the seismic radiation from faults. One class of methods sums subevents over a finite fault, and another represents the fault by an equivalent point source. The latter can still be applied to faults of moderate size as long as the source characterization and the distance measure used in the hazard calculations are consistent with the method for predicting the ground motions.

Summations Over a Finite Fault. In methods treating the source as a fault with finite, nonzero extent, the motions are generally calculated by summing and lagging the motions from many subevents distributed over a fault plane with a particular orientation and size. In the sense of a Monte Carlo study, the motions for many realizations of the subevents can be combined to provide the ground-motion distribution needed for PSHA. These subevents can be defined in a number of different ways. Some methods use records from actual earthquakes as the subevents (e.g., Somerville et al. 1991; see also Appendix E). A number of others generate a random slip distribution with prescribed properties (many studies have done this, a recent example is Herrero and Bernard 1994). Finally, some methods do not attempt to simulate a physical distribution of slip over the fault plane, but instead use simple subevents, with the needed complexity contributed by adding together the motions from a distribution of these subevents (many small ones and a few whose dimensions are comparable to that of the target event; e.g., Zeng et al. 1993).

Equivalent Point Source. This popular subset of models of the Stochastic Models method, often referred to, in the literature, as "stochastic models," describes the radiation from a fault in terms of ground-motion spectra whose amplitude and shape are given by smooth, relatively simple functions, and whose phase is quasi-random, such that the motions are distributed in time over a duration related to the size of the source and the distance from the source to the site (e.g., Hanks and McGuire 1981; Boore 1983; Toro and McGuire 1987). The amplitude spectra can be obtained by fitting functional forms to data (e.g., Atkinson 1993b), but more often are taken from physically-motivated seismological models of the source. The most common is the single-corner frequency, ω^2 model; this model is usually, but not completely accurately, referred to as the "Brune" model (after Brune (1970)). Joyner (1984) has published a two-corner extension of the model to account for a breakdown in the self similarity of seismic sources. No attempt is made in the equivalent point-source models to account for the distribution of motions around a fault of a particular orientation; usually a simple scalar factor, taken from studies such as Boore and Boatwright (1984), is used to represent the effect of rays leaving the source in many directions. Because a randomization is not needed for many subevent distributions over a given fault plane and for many orientations of the fault plane, the computational requirements of the equivalent point-source methods are almost trivial.

Path

As seismic energy leaves the source, it is subject to modification enroute to the site. In some methods, this modification is captured by using empirical Green's functions—recordings of small events at a site that have traversed the specified path. Such Green's functions are of little use for PSHA, however, for seldom are sufficient Green's functions available for a specific site of interest. Instead, most often the modification is parameterized by the multiplication in the frequency domain of two factors. The first is a simple geometrical spreading (often frequency-independent) to model the elastic wave-propagation effects. The simplest form of this model assumes propagation in a uniform whole space and predicts a decay proportional to 1/r; a

straightforward modification makes this decay proportional to $1/\sqrt{r}$ at distances greater than 100 km to account for multiple reflections of waves trapped in the earth's crust (Herrmann 1985). The second is a representation of anelastic attenuation, usually given by a function of the form exp (-π f r/ Q c), where c is the shear-wave velocity and the quality factor Q can be frequency dependent:

$$Q = Q_0 (f / f_0)^{\eta} \qquad (5.2)$$

where usually $f_0 = 1$. The parameters for the anelastic attenuation are usually obtained from studies of wave propagation in the region of interest.

More realistic models of geometrical spreading are now in use. For example, somewhat more complicated functional forms than 1/r can be fit to existing data (e.g., Atkinson and Mereu 1992). Another class of models uses wave propagation in a layered earth to account for the path complexities. These include computationally-rapid high-frequency approximations that try to capture the essential modifications due to the earth's layering (e.g., Ou and Herrmann 1990; EPRI 1993), generalized ray theory (Somerville 1992), and full-wave calculations that compute the complete wave motions for a layered earth (Saikia 1994). The latter two methods are computationally intensive, and the basic assumption of plane layering is in many cases a very gross approximation to reality. Furthermore, even if the earth were plane-layered, it is unlikely that the properties of the layering would be known in sufficient detail to account for propagation of high-frequency waves. For these reasons, the motions computed for full-wave methods, although mathematically precise, may not give *a priori* predictions of ground motion that are any more accurate than the methods that treat the path in a much simpler manner. Recognizing that full-wave methods may not produce sufficient complexity in the motions, some recent studies are adding an empirically-determined filter to the simulations to account for the scattering of waves due to geologic complexities not included in the computational models (e.g., Horton 1994).

At long periods (which are important for base-isolated structures, suspension bridges, or tall buildings) the full-wave methods become more attractive because the path is more deterministic at this larger scale (see Helmberger et al. 1993). A hybrid approach, which uses a full-wave method for long periods and a more stochastic method for short periods has been utilized for this purpose (Saikia and Somerville 1995).

Anelastic attenuation has only a moderate effect on PSHA for frequencies below 10 Hz, especially in the central and eastern U.S. where anelastic attenuation is lower. Although geometrical spreading and anelastic attenuation are treated as separate phenomena, they are strongly coupled in practice, because ground motion is affected by both. Given the amount of data and their scatter, it is typically impossible to resolve the two effects; all that is known is their combined effect.

Site

As mentioned earlier, site effects can have a first-order effect on ground motions. This can be the case even for sites that are nominally founded on hard rock. In fact, at least one study suggests that there is more ground-motion uncertainty in rock sites than in soil sites (Abrahamson and Sykora 1994). It is very important in elicitation of ground motions to be specific about what is meant by "hard-rock"; in our workshops, we specified this to be sites underlain by material for which the time-weighted shear velocity exceeded 2800 m/s.

Examples

Because each piece of the ground-motion puzzle can be treated in so many different ways, many ground-motion models can be obtained. It is because of this diversity that a well-developed elicitation process is necessary to obtain the ground motions needed for PSHA. We devote Section 5.7 and a good part of Chapter 3 to this need. Some discussion of the strengths and weaknesses of a number of specific models for the predictions of ground motions in the central and eastern U.S., circa 1994, is contained in the discussions of the ground-motion workshops held

during the course of this study (Appendices A and B).

5.4.3 The Use of Fixed Spectral Shapes Anchored to PGA

Introduction

One of the approaches for obtaining a representation of seismic hazard as a function of structural frequency is to perform the seismic hazard calculations for PGA only and then use an independently obtained spectral shape or a standard spectral shape to convert the PGA to spectral accelerations or velocities at all frequencies of interest. This approach maybe viewed as assuming a ground-motion model in which spectral acceleration at any frequency f is given by C(f)xPGA(M,R), where C(f) is independent of magnitude and distance. This approach was commonly used in the past.

This section begins by describing the main factors that affect the spectral shapes of earthquake ground motions. Considerations about these factors are then used to support the recommendation that fixed spectral shapes not be used in seismic-hazard analysis (except for studies of limited or very-limited scope).

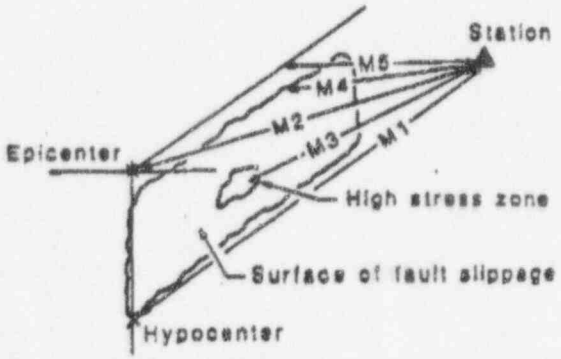

Figure 5-2a Diagram illustrating different distance measures used in predictive relationships (from Shakal and Bernreuter 1981).

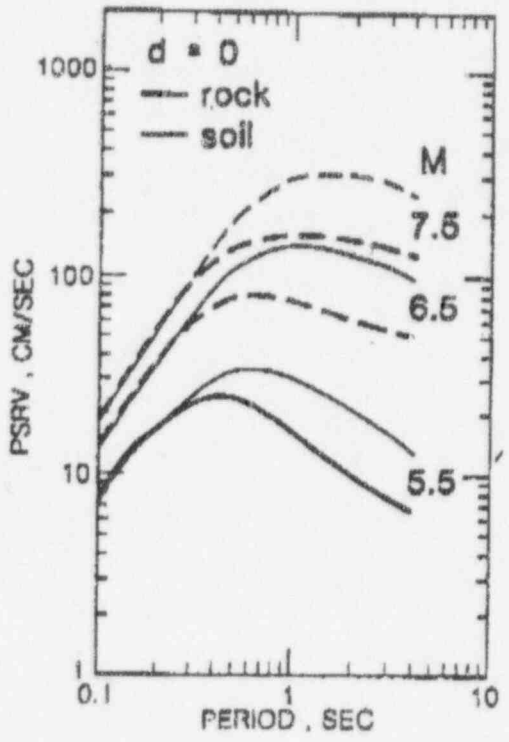

Figure 5-2b Predicted pseudovelocity response spectra for zero epicentral distance and several values of moment magnitude. Source: Joyner and Boore 1988.

Factors Affecting Spectral Shapes

Magnitude

Larger earthquakes generally break a larger portion of the earth's crust and have longer durations than smaller earthquakes. As a result, larger earthquakes are more effective at producing lower-frequency ground motions and have a higher proportion of low-frequency energy relative to the high-frequency energy than smaller earthquakes. Figure 5-2 shows the effect of magnitude on spectral shapes. The effect of magnitude on spectral shapes is also obvious by examining the magnitude coefficients in attenuation functions for spectral acceleration and for PGA (e.g., Boore et al. 1993; Toro et al. 1995): the coefficients for 1 Hz spectral acceleration are approximately twice as large as those for PGA.

As an illustration, consider the calculation of seismic hazard of a hypothetical 1 Hz structure located in a region with maximum magnitude in the range of 6.0 to 6.5. If we use a spectral shape based on earthquakes with magnitudes near 7.0, we implicitly assume that the ratio *PSA*(1 Hz)/*PGA* is the same for the earthquakes affecting the site as for the magnitude 7.0 earthquakes used to construct the standard spectral shapes commonly used (we will return later to the standard spectral shapes). We know from seismological theory and observation, however, that this ratio is lower for the earthquakes that threaten our hypothetical structure, due to the effects of magnitude.

High-Frequency Energy

Earthquake ground motions at rock sites in the western United States have little energy at frequencies higher than 20 Hz (Hanks 1982). In contrast, a number of ground-motion records obtained at hard-rock sites in the central and eastern U.S. have significant energy at frequencies as high as 50 Hz. This high-frequency energy affects the spectral accelerations at high frequencies (f>20 Hz) as well as the PGA, but it does not affect spectral accelerations at lower frequencies. These differences are generally explained as the result of less damping in the upper crust (Hanks 1982; others), but alternative interpretations have been proposed (e.g., Papageorgiou 1988). The shape of the power spectrum at high frequency is often parameterized by the frequency f_{max} (Hanks 1982) or the attenuation time κ (Anderson and Hough 1982).

As a result of these differences in high-frequency energy, the high-frequency portion of the response spectrum is very different for these two types of earthquakes as illustrated in Figure 5-3.

Considering a hypothetical structure with a 1-Hz resonant frequency in the eastern United States, we note that earthquakes in the central and eastern U.S. have similar 1-Hz amplitudes and higher PGA than California earthquakes of the same moment magnitude. Thus, the *PSA*(1 Hz)/PGA ratio is lower (nearly 50% lower, according to Figure 5-3) for the the central and eastern U.S. earthquake. As a result, using a California spectral shape (with its higher *PSA*(1Hz)/PGA ratios) with a proper central and eastern U.S. PGA attenuation function will lead to overestimation of the hazard for the hypothetical 1 Hz structure in the central and eastern U.S.

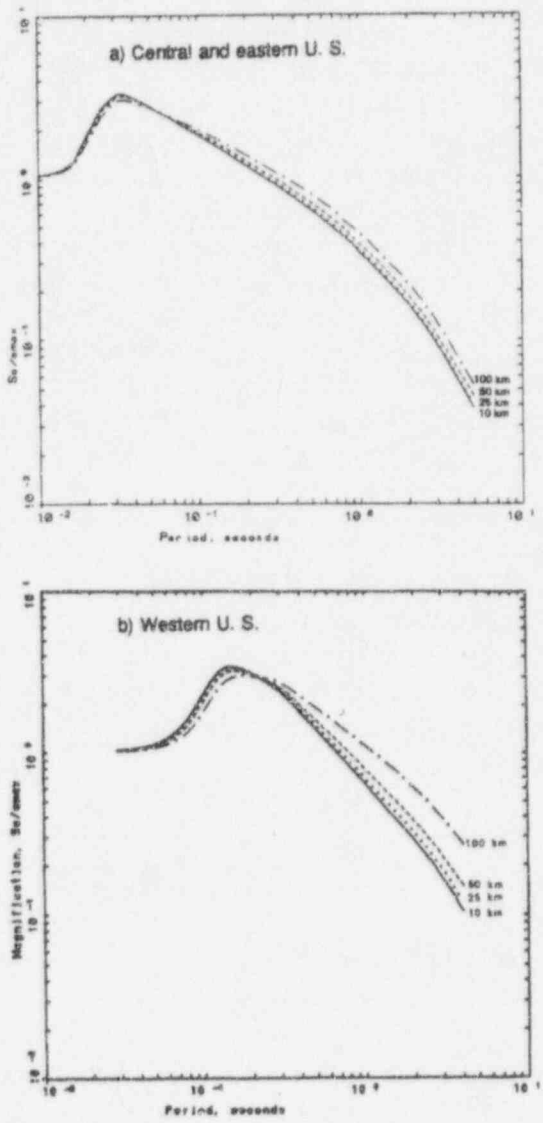

Figure 5-3 Comparison of spectral shapes for the central and eastern U.S. (CEUS) and the western U.S. (WUS) for moment magnitudes 6.5. Source: Silva and Green (1989).

Distance also has an effect on high-frequency energy, due to anelastic attenuation (i.e., damping) in the earth's crust. High-frequency waves go through more cycles as they travel a certain distance, so they undergo more attenuation. This phenomenon becomes important

at distances of more than 50 km and has only a minor effect on seismic hazard.

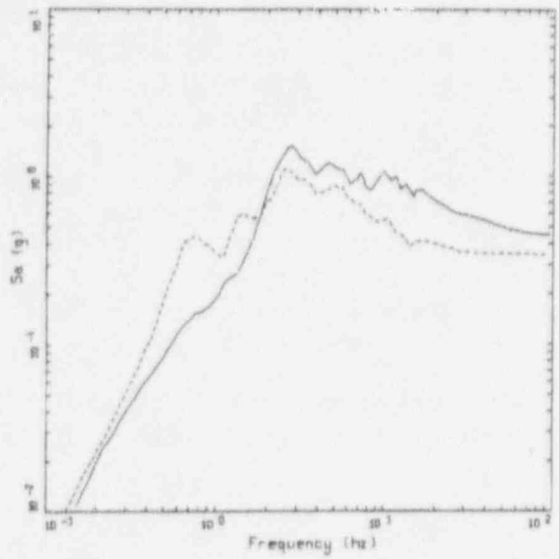

Figure 5-4 Response spectra (5% damped) recorded at Gilroy 1 (rock: solid line) and Gilroy 2 (soil: dashed line) during the 1989 Loma Prieta earthquake (the log average of two horizontal components is plotted). Source: EPRI 1993.

Site Response

Site responses may have a dramatic effect on spectral shapes. Figure 5-4 illustrates the effect of site response on spectral shape by comparing the spectra from records obtained at two nearby stations: one on firm soil, the other on rock.

Deep-soil sites tend to amplify low-frequency motions (for which amplification effects are dominant) and to dampen high-frequency motions (for which damping effects are dominant). The net effect on peak acceleration is typically small, but the effect on spectral shapes may be dramatic. Shallow-soil sites have little effect on low-frequency energy and tend to amplify high-frequency energy and PGA (for which elastic resonance effects due to trapped energy are dominant).

The data used to develop some standard spectral shapes (e.g., the NRC Regulatory Guide 1.60, U.S. Atomic Energy Commission 1973 and NUREG 0098, Newmark and Hall 1978) contain a large number of records from deep-alluvium sites in California. As a result, these spectral shapes may overestimate the PSA(1 Hz)/PGA ratios for rock sites.

Use of Spectra Associated with Certain Percentiles

A common practice in the development of spectral shapes is to select representative records, normalize or scale all records to a common PGA, compute their response spectra, and calculate the normalized spectrum associated with the 84 percentile. That is, at each frequency the calculated value of the normalized spectral acceleration is higher than the normalized spectral accelerations in 84 percent of the selected records. This practice was followed in the development of the NRC Regulatory Guide 1.60 spectrum (U.S. Atomic Energy Commission 1973; see also Blume et al. 1973).

At frequencies associated with the PGA, there is no scatter among the spectra because all spectra have been normalized to a common PGA. As one moves to lower frequencies (i.e., away from peak acceleration), the scatter among the selected spectra increases due to the effects of magnitude, high-frequency energy, and site response discussed earlier. Thus, the 84 percentile spectrum deviates substantially from the median normalized spectrum. Returning to the hypothetical 1-Hz structure, the PSA(1 Hz)/PGA ratio is much higher for the 84-percentile spectrum than for the median spectrum[3]. This difference is not the result of justified conservatism in the face of uncertainty; rather, it is the result of conservatism combined with a sub-optimal procedure for the estimation of low-frequency spectral accelerations (see also Cornell 1993).

Not all standard spectral shapes in current use are associated with 84 percentiles. For instance, the NUREG-0098 median spectral shape (Newmark and Hall 1978) is often used.

5.4.4 Recommendations

It is recommended that the representation of seismic hazard as a function of structural

[3]Recall that the PSA(1 Hz)/PGA ratio in the median spectrum may already be too high, due to the effects mentioned earlier.

frequency be obtained directly through attenuation functions (or equivalent formulations) that predict spectral acceleration as a function of structural frequency, rather than using a fixed spectral shape anchored to a value of PGA. This recommendation implies that separate seismic-hazard calculations must be performed for each frequency, but this is not a problem with today's computational capabilities. This is also the only way to capture the manner in which the various seismic sources contribute differently to hazard at the different structural frequencies.

The recommended approach also requires more input than a PGA hazard analysis, because it requires attenuation equations for spectral accelerations at multiple frequencies. This is less of a problem at present than in the past for the following two reasons: (1) most ground motion records of interest are routinely digitized and processed (making the response spectra available), and (2) there has been significant progress in the use of seismology-based models of ground motion, which can readily provide estimates of spectral accelerations.

In cases where the limited scope of the study makes it necessary to use fixed spectral shapes, these shapes should be developed on a site-specific basis, using records that are representative of the seismic exposure, κ values, and soil conditions of the site. The site-specific spectral shape should be associated with a median normalized shape. Also, it may be desirable to use spectral acceleration at some intermediate frequency, rather than PGA, as the reference ground-motion measure used for anchoring.

In practice, it is nearly impossible to find a sufficient number of records (i.e., more than 10) that meet all the required conditions (the exception is portions of California and similar areas with high seismic activity and dense network of strong-motion instruments). In addition, the level of effort is not too different from that associated with a seismic hazard analysis for multiple spectral accelerations.

A simpler approach is to use existing attenuation functions or seismological ground-motion models to construct the site-specific spectral shape. These attenuation equations or ground-motion models used for this purpose must have sufficient basis (either empirical or based on sound theory) for predicting spectral accelerations at the magnitude, distance, and structural frequency of interest. This approach is also more economical than the approach described above.

Standard spectral shapes may be used in studies of limited scope only if they are shown to be consistent with the above conditions or if they are shown to be conservative, as long as this conservatism is not burdensome to the owner or operator of the facility.

Standard spectral shapes may be used in studies of very limited scope, without having to show their applicability.

5.5 Characterization of Uncertainty in Ground-Motion Predictions

5.5.1 Types of Uncertainty

This section presents a description of the types of uncertainty in ground-motion predictions. The taxonomy of uncertainty used by some ground-motion analysts is somewhat more elaborate than the one presented in Chapters 1 and 2. Although this taxonomy should not be mandated, it is included here because it provides useful insights into the causes of uncertainty and allows the quantitative calculation of that uncertainty.

The distinction between aleatory and epistemic uncertainties was introduced in Chapters 1 and 2; it is repeated here for emphasis.

Epistemic Uncertainty. Uncertainty that is due to incomplete knowledge and data about the physics of the earthquake process. In principle, epistemic uncertainty can be reduced by the collection of additional information.

Aleatory Uncertainty. Uncertainty that is inherent to the unpredictable nature of future events. It represents unique details of source, path, and site response that cannot be quantified before the earthquake occurs. Given a model, one cannot reduce the aleatory uncertainty by collection of additional information. One may be able,

however, to quantify the aleatory uncertainty better by using additional data[4].

From the point of view of the ground-motion analyst who is using physical models, the total uncertainty in predicted ground motions is often partitioned in a manner that may be considered orthogonal to the above partition (see Abrahamson et al. 1990), as follows:

Modeling Uncertainty. Represents differences between the actual physical process that generates the strong earthquake ground motions and the simplified model used to predict ground motions (Abrahamson et al. 1990, call this modeling+random uncertainty). Modeling uncertainty is estimated by comparing model predictions to actual, observed ground motions.

Parametric Uncertainty. Represents uncertainty in the values of model parameters (e.g., stress drop, anelastic attenuation) in future earthquakes. Parametric uncertainty is quantified by observing the variation in parameters inferred (usually in an indirect manner) for several earthquakes and/or several recordings.

It is important to recognize that the distinction between modeling and parametric uncertainty is model-dependent. For instance, one may reduce the scatter in the predictions by making the model more complete, thereby introducing new parameters in the model. Unless these new parameters are known *a priori* for future earthquakes and for the site of interest, there will be additional parametric uncertainty, thereby transferring some modeling uncertainty into parametric uncertainty, without varying the total uncertainty.

Both the modeling and parametric uncertainties contain epistemic and aleatory uncertainty. For instance, observed scatter that is not accounted for by the model and varies from event to event is aleatory modeling uncertainty, whereas statistical variability in the calculated bias that introduces uncertainty about the accuracy (or unbiasedness) of the model (due to limited data) is epistemic modeling uncertainty. Similarly, the event-to event variation in stress drop is aleatory parametric uncertainty, whereas the imperfect knowledge about the probability distribution of stress drops from future earthquakes (e.g., What is the median stress drop for **M** 7 earthquakes?) is epistemic parametric uncertainty. Table 5-1 illustrates this two-way partition of total uncertainty. The different types of uncertainty are illustrated by way of a more concrete example. The Hanks-McGuire ground-motion model (Hanks and McGuire 1981; Boore 1983, etc.) may be used to predict spectral acceleration for a given magnitude and distance and for given values of the model's physical parameters; i.e.,

$$\ln[Amplitude] = f(\mathrm{m,r}: \Delta\sigma, Q, f_{max}) \qquad (5.3)$$

where the stress drop $\Delta\sigma$, the quality factor Q, and the frequency f_{max}, are physical parameters of the model.

When one applies this model to well-studied events in well-studied regions (for which the parameter values have been determined), and compares predictions to observations, one observes some scatter and possibly some bias, addressed below, because the physical model contains only a crude representation of source and path effects. This scatter represents aleatory modeling uncertainty (i.e., observed scatter not explained by the physical model). In order to include this scatter in our predictions, the physical model above is used to construct an aleatory model of the form

$$\ln[Amplitude] = f(\mathrm{m,r}: \Delta\sigma, Q, f_{max}) + \varepsilon_{\text{aleatory modeling}} \qquad (5.4)$$

where $\varepsilon_{\text{aleatory modeling}}$ is a zero-mean random quantity that represents the observed scatter not explained by the physical model.

[4] An example may clarify these definitions. Consider a Gaussian random quantity X with mean μ and standard deviation σ. The value of μ represents the deterministic component of X. The value of σ represents aleatory uncertainty in X. The probabilistic modeler's uncertainty about the true values of μ and σ (due to a small statistical sample or to alternative hypotheses about the nature of X) represents epistemic uncertainty in X.

Table 5-1 Partition of Uncertainty in Ground-Motion Prediction

		Seismic-Hazard Analyst	
		Epistemic	Aleatory
	Modeling	Uncertainty about the true model bias (i.e., to what extent model has tendency to over- or under-predict observations)	Unexplained scatter due to physical processes not included in the model
	Parametric	Uncertainty about the median stress drop for the central and eastern U.S., depth distribution, etc.	Event-to-event variation in stress drop or focal depth, etc.

Because we need to make predictions for future earthquakes (for which the stress drop and other model parameters are not yet known), those predictions will contain additional uncertainty: aleatory parametric uncertainty. Knowing the aleatory, event-to-event and site-to-site, variation in the parameters, one can calculate the associated aleatory uncertainty in ground-motion amplitude by using the methods of derived distributions (Benjamin and Cornell 1971). Referring to Section 2.2, we note that the aleatory modeling and aleatory parametric contributions are combined into ε_a in Equation 2-5.

Epistemic uncertainty in the above prediction also comes in two forms, as follows. The limited data, and the scatter in these data, do not allow us to quantify any systematic biases in the physical model's predictions for given parameter values. Small biases are obscured by the scatter in the observations, unless one has observations from many events and sites. It may also happen that most of the data fall outside the (m,r) ranges of engineering interest. This uncertainty is epistemic modeling uncertainty. In addition, the aleatory distributions of the model parameters are not known exactly (e.g., What is the median stress drop for ENA earthquakes?); this introduces epistemic parametric uncertainty. Referring to Section 2.2, we note that the epistemic modeling and epistemic parametric contributions are combined into ε_e in Equation 2-5.

Appendix F describes how these concepts are used in the context of the various types of ground-motion models and how the various components of uncertainty are estimated in practice.

The distinction between epistemic and aleatory uncertainties is common practice in PSHA (see Chapter 2) and should be maintained throughout the process of characterizing uncertainty in ground-motion predictions. The distinction between modeling and parametric uncertainties is a useful tool for the quantitative determination of the epistemic and aleatory uncertainties in the context of physical models, but it is not required. In fact, this latter distinction is internal to the ground-motion modeling and is not carried downstream into the seismic hazard calculations.

Site-Specific Perspective on Uncertainty. In ground-motion studies, the limited availability of data forces the investigator to use data from large geographical areas[5]. Any undetected geographic trends in the data are implicitly counted as part of the aleatory parametric uncertainty. (Also, undetected geographic variations in the model bias are implicitly counted as aleatory modeling

[5]In physical models, data on different parameters may be collected at different geographical scales (e.g., one may use stress-drop data from all of ENA, while using Q data from a smaller region).

uncertainty). Thus, the scatter one obtains from regional data applies to a site chosen at random.

In seismic hazard studies for a given site, one is interested in ground motions from certain seismic sources (with their particular distributions of stress drop and source depth and their preferred focal mechanisms), which are propagated through a certain portion of the earth's crust (over a radius of say, 100 km, with its particular Moho depth, Q, and upper-crust velocity profile), and are further modified by the local geologic conditions beneath that specific site (with its particular amplification, resonance, κ, and degradation properties). If one takes this perspective, much of the geographic variations that were implicitly counted as aleatory uncertainty should be counted instead as epistemic, thereby decreasing the aleatory uncertainty and increasing the epistemic uncertainty.

If site-specific information on any of these parameters is obtained (from geophysical or geotechnical studies, regional Q studies, weak-motion recordings, etc.), the epistemic parametric uncertainty will be reduced accordingly (and there will likely be a change in the central value of the parameter).

This site-specific perspective is, in principle, the proper perspective for all seismic-hazard studies, regardless of the level of effort, and regardless of the availability of site-specific data. In practice, however, it may be difficult to quantify *a priori* how much of the aleatory parametric uncertainty in a parameter is associated with geographic variation and should be treated as epistemic.

This site-specific perspective is particularly important for site effects, because the site-response parameters (shear-wave velocity profile, stiffness-and damping-degradation curves, and κ) have a significant effect on ground motions. Also, these parameters are determined as part of the site-characterization studies for important facilities. Thus, one would expect a significant reduction in epistemic uncertainty (and a significant, but unknown *a priori*, change in the median ground-motion prediction) when site-specific site-response information becomes known and is used to update ground-motion predictions. This is true for both soil and rock sites, as there appears to be significant differences in the response of different rock sites.

5.5.2 Propagation of Parametric Uncertainties

Let $X_1, X_2, \ldots X_n$ be aleatory quantities representing uncertain parameters of the ground-motion model (e.g., stress drop, Q). Aleatory uncertainty in the values of $X_1, X_2, \ldots X_n$ for a given event and site is represented by probability distributions with parameters[6] (e.g., means and standard deviations) $\Theta_1, \Theta_2, \Theta_3, \ldots \Theta_m$. Uncertainty in the values of $\Theta_1, \Theta_2, \Theta_3, \ldots \Theta_m$ represents epistemic uncertainty, and is also represented by probability distributions. Also let

$$\ln[\text{Amplitude}] = f(m,r;X_1,X_2,\ldots X_n) + \varepsilon_{\text{epistemic modeling}} + \varepsilon_{\text{aleatory modeling}} \quad (5.5)$$

represent the ground-motion model, including modeling uncertainty. The propagation of parameter uncertainties into uncertainties in ground-motion amplitudes (i.e., finding the distribution of ln[Amplitude] as a function of only m and r) is one of derived probability distributions (see Benjamin and Cornell 1970). It may be performed using logic trees, Monte Carlo simulation, or other appropriate methods. Because epistemic and aleatory uncertainties must be kept separate, this propagation must be performed in a nested manner. The innermost step consists of calculating the distribution of ln[Amplitude(m,r)]for given values of $\Theta_1, \Theta_2, \Theta_3, \ldots \Theta_m$, and $\sigma_{\text{aleatory modeling}}$, by integrating over all possible values of the aleatory quantities given $\Theta_1, \Theta_2, \Theta_3, \ldots \Theta_m$, and $\sigma_{\text{aleatory modeling}}$. In practice, one often calculates the mean and standard deviation of ln[Amplitude(m,r)] rather than the full distribution. The outer step is to calculate the epistemic distribution of the mean and standard deviation calculated above, when $\Theta_1, \Theta_2, \Theta_3, \ldots \Theta_m$, and $\sigma_{\text{aleatory modeling}}$ are allowed to vary based on their respective probability distribution.

6We use the word parameter with two different meanings in this paragraph. The first time we mean a physical (or perhaps empirical) parameter such as stress drop (aleatory), the second time we mean a parameter of a probability distribution such as median stress drop (epistemic).

The result of this exercise is a model of a form that may be used in the seismic-hazard calculations, i.e.,

$$\ln[\text{Amplitude}(m,r)] = g(m,r) + \varepsilon_{\text{epistemic}} + \varepsilon_{\text{aleatory}}$$

where $\varepsilon_{\text{epistemic}}$ and $\varepsilon_{\text{aleatory}}$ have zero mean. We can write the following expressions for the quantities appearing in the above equation.

- Mean value of ln[Amplitude(m,r)] and epistemic standard deviation of ground motion amplitude.

$$g(m,r)=E_{\Theta}\{E_{X|\Theta}[f(m,r,x)]\}$$

$$\sigma_{\text{epistemic}} = \left[\text{Var}_{\Theta}\left\{E_{x|\Theta} f(m,r,x)\right\} + \sigma^2_{\text{epistemic modeling}}\right]^{1/2} \quad (5.6)$$

- Mean value and epistemic standard deviation of the variance of $\varepsilon_{\text{aleatory}}$:

$$\sigma^2_{\text{aleatory}} = E_{\Theta}\left\{\text{Var}_{x|\Theta}[f(m,r,x)]\right\} + \sigma^2_{\text{aleatory modeling}}$$

$$\sigma_{\sigma^2_{\text{aleatory}}} = \left[\text{Var}_{\Theta}\left\{\text{Var}_{x|\Theta}[f(m,r,x)]\right\} + \text{Var}_{\text{stat.}}\left[\sigma^2_{\text{aleatory modeling}}\right]\right]^{1/2} \quad (5.7)$$

In the above equations, $E_X[\,]$ and $\text{Var}_X[\,]$ represent the expectation and variance operators, Θ represents the vector of $\Theta_1, \Theta_2, \Theta_3, \ldots \Theta_m$, and X represents the vector of $X_1, X_2, \ldots X_n$, and we assume that $\varepsilon_{\text{aleatory modeling}}$ and $\varepsilon_{\text{epistemic modeling}}$ have zero mean. The last term in the last equation represents statistical uncertainty in the value of $\sigma^2_{\text{aleatory modeling}}$. The above equations show the nested nature of these calculations, with the conditional integrations over $X|\Theta$ on the inside and the integrations over Θ on the outside.

One may choose not to integrate over a few of the more important epistemic uncertainties, leaving those uncertainties explicitly in the model (in which case one would have to provide the conditional values of the above four quantities). Those uncertainties would then be considered explicitly in the seismic hazard calculations. This has the advantage that one can calculate the sensitivity of the seismic hazard to those uncertain quantities.

One may also choose to calculate the full epistemic distribution of $\varepsilon_{\text{epistemic}}$ and $\varepsilon_{\text{aleatory}}$, rather than their first and second moments. These calculations would follow the same nested structure shown above, although the computations would be somewhat more demanding.

5.6 Recommendations

Based on the discussion in Sections 5.2 through 5.5, the following recommendations are presented (recommendations on elicitation are contained in Section 5.7).

Ground Motion Measures

1. We recommend that PGA alone not be used for most applications. A much more useful measure of ground motion is the response spectrum.

2. We recommend that vertical motions, if important for the study, be obtained from independent analyses, in the same manner as for the horizontal motions.

Explanatory Variables

1. All data used to construct the attenuation functions must be in the same magnitude scale and this scale must be the same as the scale used to define seismicity parameters.

2. Seismic hazard studies in the western United States must use the moment magnitude scale. Although seismic hazard studies in CEUS may use either moment magnitude or m_{blg} in the near term, we strongly recommend that an effort should be made to convert the CEUS earthquake catalogs to moment magnitude. At the same time, the collection of macro-seismic data from current earthquakes should not be discontinued, in order to improve our understanding of the relationship between macro-seismic effects and instrumental magnitudes. Because of its potential utility in PSHA, we also recommend a detailed evaluation of the newly-proposed high-frequency magnitude m.

3. **Distance** The definition of distance in the attenuation equations must be consistent with the geometric model of earthquake occurrences. Attenuation functions in terms of distance to a point source (or a projection thereof) are consistent with areal seismic sources and are appropriate for source dimensions less than 5 km. Attenuation functions in terms of distance to the rupture surface (or a projection thereof) are consistent with fault-type seismic sources and are required for source dimensions of 5 km or greater. Consistency is also required for source dimensions of 5 km or greater. Consistency is also required in the treatment of depth, which is important for the host seismic source or for faults located near the site.

4. **Site Response** If the scope of the study does not warrant a site-specific site response analysis, it is necessary to use attenuation equations that are applicable to the conditions at the site or to use appropriate amplification factors. The explanatory variables that characterize site conditions should consider both the depth and dynamical properties of the site; the use of a soil/rock dichotomous explanatory variable is not sufficient.

5. For site-specific analyses, sufficient resources should be made available to adequately characterize local geologic and geotechnical properties.

6. **Additional Explanatory Variables** In most situations, there is no need to introduce additional ground-motion explanatory variables to represent earthquake characteristics or location of the site relative to the earthquake. Additional explanatory variables may be introduced with adequate justification, as discussed in Section 5.3.2. If additional explanatory variables are introduced, their aleatory and epistemic uncertainties must be modeled explicitly.

Methods for Predicting Ground Motion

1. A methodology for obtaining ground motions that adequately capture the uncertainty in the estimates and is defensible in a regulatory arena is recommended in Chapter 5.

2. We recommend that the use of scaled spectral shapes be avoided.

Recommendations on Uncertainty

1. The estimates of total uncertainty in ground motion must be realistic and must include all sources of uncertainty. In particular, one must avoid the following two frequent situations: (1) very narrow estimates of uncertainty as a result of ignoring the existence of other models or the possibility of alternative interpretations of the existing data, or (2) very broad estimates of uncertainty (which would allow for unreasonable ground-motion amplitudes or which predict much more scatter or bias than is observed in the data in (m,r) regions where data are available).

2. The partition of total uncertainty into aleatory and epistemic, though sometimes arbitrary, must be made carefully and in a manner that is consistent with current practice.

3. Ground-motion analysts are encouraged to use quantitative procedures for the development of uncertainty estimates and to follow the framework discussed here. This facilitates the exchange of information and should help resolve some of the differences between experts' estimates of uncertainty. It is recognized, however, that there are limits to the applicability of purely data- and model-driven procedures and that some subjective inputs are always required.

5.7 Specific Expert-Elicitation Guidance for Obtaining Ground-Motion Values

The ground-motion information needed for PSHA is the probability distribution of the ground-motion measure of interest, conditional on earthquakes of magnitude **M** occurring at distance R, for all **M** and R within a specified range. Usually the probability distribution is specified by giving the median value of ground motion and a parameter related to the breadth of the distribution function. In most cases, the ground-motion

measure is assumed to follow a lognormal distribution. The goal of the elicitation, then, is to obtain the median value and the breadth parameter for any M and R within the specified range. The elicitation must consider the uncertainty in the desired parameters.

The procedure for performing the elicitation will depend on the particular project. For projects that do not involve critical facilities whose failure might cause a substantial hazard to the nation and its citizens, the elicitation might involve nothing more than an analyst choosing an equation from the literature. Elicitation at the other end of the spectrum involves an intensive effort that employs a group of experts. We will concentrate here on this latter case. At the end of this section we have a short discussion on the use of the contractor/peer review process outlined in the previous section.

In this section we will recommend a procedure for obtaining ground motions that should be as applicable ten years from now as it is today, even if a new generation of ground-motion models is available by then. The bulk of this section will be devoted to a detailed discussion of how to use multiple experts to obtain ground-motion values. We imagine that such an exercise will be undertaken every 5 to 10 years, and will be focused on ground motions to be used in regional studies. For many projects, the results of such an elicitation can be used with or without modification. For site-specific studies, more detailed knowledge of important parameters such as crustal structure, Q, kappa, and local basin and soil properties might be available, and it would be appropriate to modify the ground-motion values to account for these site-specific properties.

This section begins with a short discussion of the ground-motion elicitation procedures used in the EPRI and LLNL studies. It then presents a brief summary of SSHAC's recommendation for performing ground-motion elicitation. This is followed by a more detailed discussion of the elicitation process. Further details and supporting information are contained in several Appendices.

This section is designed to augment, not substitute for, the general guidance in Chapter 3. Every element of the general guidance applies to ground motion specifically and is an integral part of the ground-motion elicitation process.

5.7.1 Historical Perspective

For ground-motion elicitation, EPRI and LLNL used fundamentally different procedures. EPRI used one analyst (a few individuals in one consulting company), who conducted several information-gathering workshops and then decided on particular models to be used in the analysis. Weights were assigned by the analyst to three specific models, and the hazard calculations were performed for each of these three ground-motion models; the weights were used in aggregating the results.

LLNL used two different procedures. In their initial work (mid 1980's) they asked each member of a group of experts to assign weights to a set of ground-motion models. As in the EPRI work, the hazard analysis was performed for each model in combination with the many different seismicity models. The elicitation procedure involved an information-gathering and dissemination workshop, but, by design, the interaction among experts and the analyst team was minimal. LLNL adopted the role of a "weak" integrator, for they did not want to influence the experts in their choice of models.

In the early 1990's, LLNL again elicited ground-motion information from experts, but in this case what they did with the information—and to some extent how they obtained the information—differed from their first elicitation. After a one-day information dissemination workshop, LLNL asked the experts, in individual interviews, to provide ground-motion estimates and associated uncertainties for a set of magnitudes and distances (what we will refer to hereafter as points in (M, R) space). These estimates were combined to produce a "composite" ground-motion distribution which was used to compute the hazard. LLNL again adopted the role of a "weak" integrator, and interaction among the experts, although intended, was minimal. Lessons learned from LLNL's experience is contained in R. Mensing's paper (Appendix D).

SSHAC's Recommendations

On the basis of past experience with the LLNL and EPRI studies, as well as our experience in two ground-motion workshops convened to test some of our ideas (Appendices A and B), we recommend the following:

Composite ground-motion estimator

SSHAC recommends that ground-motion measures be estimated for a selected set of specific points in (M,R) space (as LLNL did in their second elicitation). The sections below describe a process for eliciting these estimates. Here we discuss the question of how to use the limited set of explicitly elicited (M,R) pairs to produce the information necessary for the hazard calculation, namely a functional form that can be used to estimate the ground motion for all (M,R) pairs.

There are at least three basic ways to generate ground-motion estimates for an arbitrary point in (M,R) space:

1. If the (M,R) pairs were constructed using explicit numerical weights on multiple models, the natural process is to form a composite model equal to a weighted average of the multiple models, and to use this composite model to calculate ground motion (the explicit numerical weighting approach is discussed below; one of its advantages is the attractiveness of this well-defined interpolation process).

   SSHAC strongly discourages the use (without scientific justification) of individual models with the same weights for all points in (M,R) space. Thus, the TFI elicitation process may often result in weights that are different (M,R) pairs. In this case a reasonable approach would be to vary the weights smoothly (e.g., linear interpolation) around the (M,R) space in such a way that the composite model fits through the elicited discrete set of points.

2. Another process is to use one specific model, such as a stochastic model, and adjust the parameters to provide a representation of the median that is a good fit through the elicited (M,R) pairs.

3. A third approach, used by LLNL in 1992, is to perform a regression analysis to fit a response surface empirically (i.e., some convenient parametric mathematical function) to the means of the (M,R) pairs.

   The standard deviations (aleatory and epistemic) can be dealt with similarly; the variation over the (M,R) space is less strong, however.

Use of a TFI

For reasons given elsewhere (see Chapter 3) we recommend for multiple-expert applications that ground-motion elicitation be done by a TFI—one entity responsible for producing a composite, ground-motion estimator based on input and interaction among experts and between the TFI and the experts. In a very real sense, the TFI will have intellectual responsibility for the product. The TFI process is explicitly contrasted with other alternate modes of using models, including using only one model, using multiple models with explicit numerical weights, and using one core model with other models for support. We have found that the TFI process, based explicitly on the principle that there is "no one correct model," reduces the participants' tendencies to view themselves as advocates and emphasize their role as scientists and evaluators with different scientific hypotheses.

We explicitly recommend against the use of a "weak" integrator, who simply mechanically combines the expert's opinions and weights without feedback between the integrator and the expert. Furthermore, we strongly endorse an elicitation process that involves significant feedback, iteration, and group discussion among experts and the TFI. Most of the rest of the chapter contains an extended discussion of the TFI approach to elicitation.

Use of a Technical Integrator (TI) to Develop the Ground-Motion Analysis

It is also feasible to use a Technical Integrator (TI) approach for developing the ground-motion part of the PSHA analysis. We will not develop

detailed guidance here for this option, which is described in more detail elsewhere in this chapter. However, much of the guidance that follows, although directed towards the TFI approach, is useful for the TI approach as well. Moreover, acting as a TI-like evaluator of a range of scientific viewpoints is one of the roles that each expert is asked to play in the TFI process.

5.7.2 The TFI Team

The TFI has two primary roles:

1. Structures and facilitates a high level of interaction among ground-motion experts.
2. Integrates data, models, estimates and expert evaluations to produce a "final" full probabilistic characterization of ground motion as a function of magnitude, distance and frequency.

Figure 3.3, described in Chapter 3, illustrates the different types of disagreement that may occur among a group of ground-motion experts. The figure also illustrates unintended disagreements due to incomplete communication and misunderstandings.

Following the general discussion in Chapter 3, the TFI for ground motion should be a small team that includes at least two essential types of expertise:

- Functional knowledge of ground motion (science, data, models and interpretations)
- Knowledge and expertise in elicitation methods

The functional knowledge is essential in clarifying, facilitating and leading scientific interchange and in summarizing points of agreement and disagreement. The elicitation expertise is essential in designing the interaction process and in structuring and conducting the information elicitation.

It is also extremely useful to have someone on the TFI team with detailed probabilistic seismic hazard expertise. Such expertise can help guide the facilitation process by focusing it on those elements and data that most affect the final hazard calculation. It is rare to find the three types of expertise in one individual; thus, the typical minimal size team would be two or three. In the SSHAC workshops, the TFI team was four to five persons for experimental reasons, but this is on the high side. The TFI team must work together very closely and meet often, so that increasing the size of the team makes logistics difficult. Additionally, the larger the team, the harder and more challenging it will be to achieve TFI consensus.

Another essential piece of the ground-motion TFI team are the resource experts, or "Implementers," described in Chapter 3 who handle logistics, mailings, process expert information, take technical notes, etc. At least one Implementer must be a ground-motion expert. It is worth repeating that it is essential that the Implementers report directly to the TFI because of need to respond quickly to logistical and technical needs.

5.7.3 The TFI Process

Figure 5-5 provides a road map of the ground-motion elicitation process. This process is explicitly based on the elicitation guidance in Section 5.7. There are 6 stages in the process, and in most stages, there are group interactions. Each group interaction is preceded and succeeded by TFI interaction with individual experts. This section is designed to supplement, not replace, Chapter 3, which provides detailed facilitation and integration guidance.

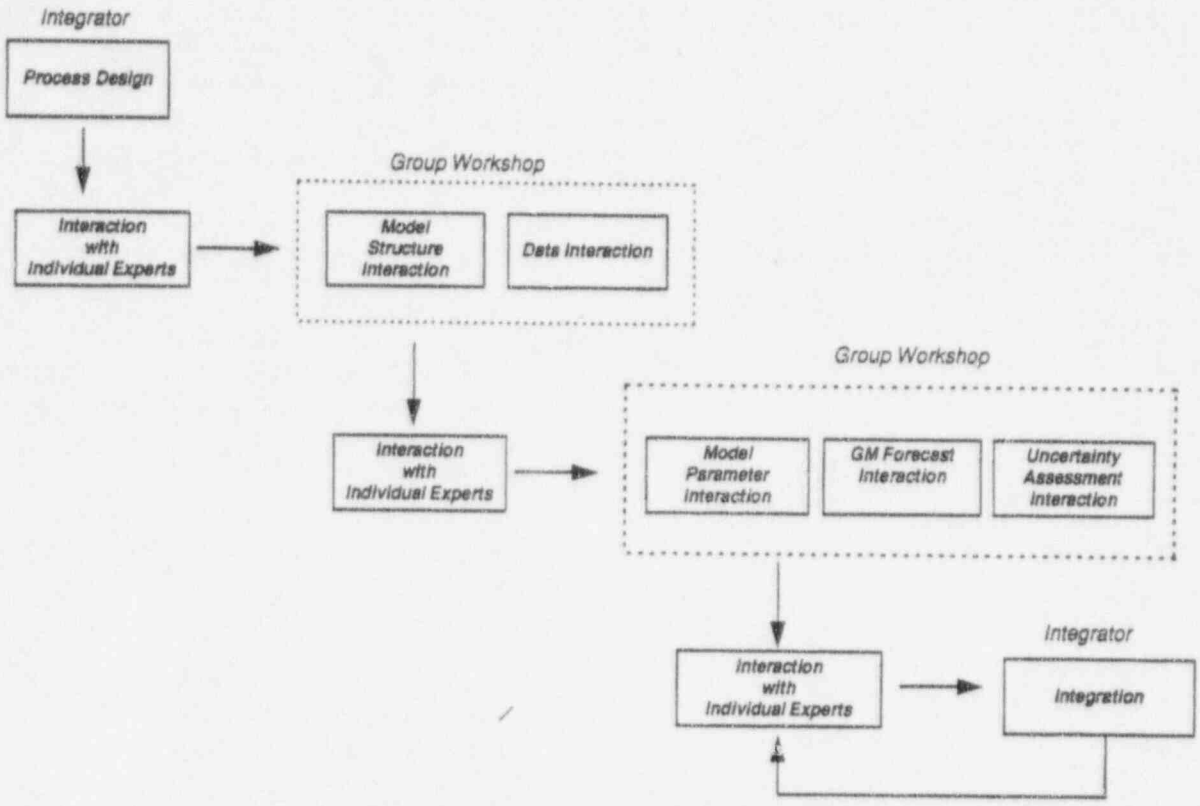

Figure 5-5. Roadmap of GM Elicitation Process

The process begins with a design stage and ends with an integration stage. The group interactions are naturally organized into two group workshops (illustrated by the dotted line boxes in the Figure) but the number of workshops is not as important as ensuring that every type of interaction occurs. TFI interaction with individual experts is also essential at every stage of the process.

The different stages of the TFI process are described below. To help organize the discussion, consideration has been given in each stage of the elicitation process to the purpose and goals of the stage, the process involved in accomplishing the goals, and the products that will result from the stage.

Stage 1: Process Design

In the first stage of the process, the TFI works with the sponsor to lay out the objectives, workplan, and time schedule. It is crucial early on to select and line up the Implementers to help with the logistics of the process.

It is also very important early on to identify potential expert participants and formulate a preliminary schedule of group meetings to ensure that the experts selected will be available roughly when they are needed. Expert selection is described in Chapter 3. The specific types of experts needed are described in the following subsections.

The TFI must fully understand the process laid out in the roadmap (Figure 5-5), as well as the TFI principles in Chapter 3. Moreover, the TFI needs to make sure there are adequate resources in terms of people, time and money to implement the process. Also, the members of the TFI team must be careful to check that they have the necessary expertise; if they do not, the team should be altered or supplemented.

In the design stage, the TFI must work closely with the sponsors of the project. Then the TFI must work with the Implementers to bring the right set of experts onto the project and to make sure they receive all the necessary information concerning their responsibilities and schedules. The TFI and Implementers must work carefully with the sponsors to make sure that necessary

contracts are set up well ahead of any work that needs to be accomplished.

The time frame for the Process Design will typically be on the order of several months because of the large number of people that need to be contacted and assembled and because of the complications and delays that are often encountered in setting up contracts for a large project. It is SSHAC's experience that a number of TFI team meetings are needed to iron out the goals, details, and procedures for each stage. We believe that the details provided in Appendices A and B that describe the SSHAC ground-motion workshops will reduce the necessary preparation time.

The product of the Process Design should be a carefully laid out workplan and time schedule, in which every type of participant has a documented task list, a set of deliverables, a set of milestones, and a budget. It is particularly important to pre-arrange at least approximate dates for large group meetings so that the participants can block out the time on their calendars.

Stage 2: Review of Data

A key lesson from the SSHAC workshops was the importance of early attention to data issues in preparation for reviewing models and expert positions (refer to Chapter 3 for more details). We suggest that a single reviewer (e.g., the resource expert) prepare a comprehensive white-paper discussing the applicable data well in advance of the first group meeting (which we will hereafter call "Workshop #1"). This is a large job and adequate time and resources should be allowed for its completion. This paper should be distributed to meeting participants in advance, with instructions that the participants carefully review the paper and be prepared to critically discuss the paper at the group meeting.

At the workshop, the reviewer would present the paper, and that presentation would be followed by intensive discussion and interactions among the participants. As in all group interactions, the TFI needs to guide the discussion to make sure that it does not stray from the task at hand.

The product would be one or several sets of data against which to compare ground-motion estimates at Stage 4 of the Elicitation Process. The data would have to be compiled into machine-readable form, but this can be done after Workshop #1.

Stage 3: Review of Methods

The main purpose of this stage is for a relatively large group of experts (from which a smaller group will be chosen for Workshop #2) and the TFI to understand the strengths and weaknesses of the basic classes of methods for predicting ground-motion measures. In addition, other goals of Stage 3 are to make sure that all reasonable methods have been considered and to perform an initial screening of the methods. This is accomplished by the following "required" activities:

- Structuring and facilitating complete information and judgment exchange
- Staging presentations by proponents of different modeling approaches
- Ensuring consistent databases and terminology
- Staging debates in critical areas
- Heavy emphasis on structured discussion regarding basic approaches rather than on individual expert opinions.

Each activity above is an essential part of the TFI process. The centerpiece of Stage 3 is a carefully structured group meeting involving intensive interaction among experts and the TFI. Prior to the group meeting, several people with strong knowledge of specific classes of methods for predicting ground-motion measures should be asked to play the role of reviewers and prepare presentations of the methods. At the group meeting these people present the methods and their strengths and weaknesses (without focusing on a particular model), with various types of group discussions following the presentations (for an example of useful types of discussions, see the agenda for the first SSHAC ground-motion workshop, presented in Appendix A).

The meeting itself requires careful facilitation. It is critical for the TFI to set the right tone for the meeting. In doing so, two elements that bear repeating from the general guidance in Chapter 3 are critical:

- The purpose is not to choose the best 'model.' The experts should be made to understand that the TFI concept is founded on the premise that there is no one correct model, and that the meeting will not be focused on trying to identify a single winner or loser. It is very important psychologically to have the participants feel that they are not there to win or lose, but to identify and clearly understand all important scientific and application issues.
- The purpose is not to achieve consensus. Consensus may occur as a serendipitous outcome, but it is important to state explicitly that the meeting will not be a failure if consensus is not achieved. Rather, it should be communicated that disagreement is not only expected, but acceptable.

It is also important that the experts understand that, other than when they are asked to be proponents (which occurs after the first workshop), they are expected to act as independent, informed evaluators (and later as integrators). An important aspect of the TFI process is that the experts are asked to provide input as to how they would integrate all the models, data and information into a composite representation of the overall expert community. If the experts feel involved as evaluators, they will tend to be constructive, and rather than resisting the process, they will assist it.

The focus in this first interaction is on the logic of different modeling approaches, rather than on variations among similar approaches. Initial focus should be on model logic rather than on numbers.

It is essential for the TFI to isolate and play back points of agreement and disagreement. This is accomplished by playing back a clear summary of the conversation frequently during the meeting. A useful facilitation model is the concept of 'active listening' (elaborated in Chapter 3). The TFI should ask experts who are not communicating to try to state each other's positions clearly.

As part of the group meeting itself, the experts should be asked to provide input on specific points in (M, R) space that are most appropriate for constructing the composite ground-motion estimator. There needs to be a clear common understanding of appropriate assumptions concerning magnitude scale and definitions of distance. The TFI should have already considered these issues, but if the experts are not involved in this initial stage, future analyses are likely to contain conflicting assumptions and errors.

Some suggested, but not required, ideas and tools are:

Pre-meeting Contact with Experts - SSHAC strongly suggests that the TFI meet with at least several of the ground-motion experts individually. This greatly aids in anticipating potential confusions and problems, in understanding the subsequent discussion at the group level, and in helping pre-structure discussion topics and define key agenda items.

Pre-meeting Contact with Reviewers - It is essential that the TFI communicate before the meeting with the reviewers to make sure that they understand their role and to promote a standardized format for their presentations. SSHAC's experience has been that without such structure and guidance, some proponents will give excellent presentations, while others will be either ill-prepared or hard to follow.

Influence diagrams - Influence diagrams are an invaluable graphical communication tool for describing basic relationships in a given modeling approach (see Howard and Matheson 1981, for a general description of influence diagrams). Figure 5-6 illustrates an influence diagram that was elicited from a ground-motion expert to describe how one ground-motion model produces an estimate of uncertainty. Such diagrams provide an excellent context for understanding the reasoning underlying a model or scientific argument.

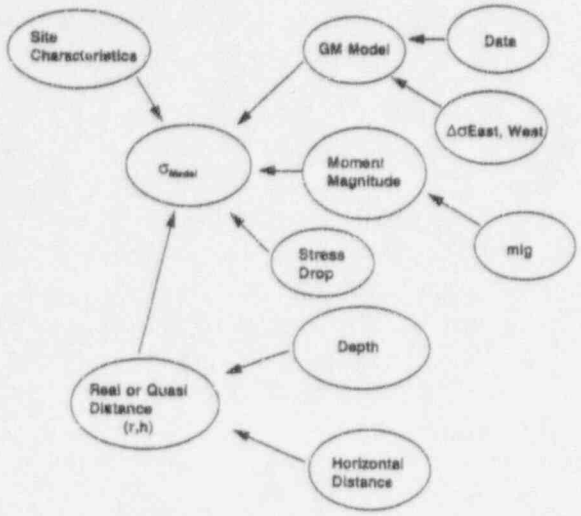

Figure 5-6. Influence Diagram

The key elements in the Stage 3 workshop are an initial session in which the workshop purpose and the various roles of the participants are explained. Second, presentations by experts designated as reviewers (who may or not be members of the expert panel) of representative methods for predicting ground motion form the basis for much of the interchange. Third, the methods should be considered one by one to make sure that all points of confusion and areas of agreement and disagreement are covered in detail. Since the relative efficacy of the different methods depends on magnitude and distance, the discussion needs to be structured (M, R) point by (M, R) point. Fourth, it is important for the experts to document their post-discussion appraisals of the different approaches—this would typically be accomplished through a written survey (the SSHAC Ground-Motion Workshop I survey is described in Appendix A). Finally, it is important that the TFI document, for all participants, a summary of the lessons learned from the meeting, including conclusions about which representative models should be run to provide numerical estimates for the next workshop.

For Workshop #1, it is important to have a large and diverse group of experts, both generalists and specialists who are able to act as reviewers of specific methods. The experts as a group should have a comprehensive understanding of existing data and models and their limitations and should be representative of the overall ground-motion expert community.

The general process of expert selection is described in detail in Chapter 3. Basically, it is important to include enough experts at Workshop #1 such that additional experts would not bring substantially different methods or interpretations to the table. When selecting among specific individuals, it is best to find specialists who are articulate and clear-thinking. It is best to find generalists who are particularly well-respected and who are not perceived to be wedded to one particular approach or interpretation. If possible, it is desirable to include experts who are non - hostile and non-emotional; this contributes to better group dynamics and information exchange. However, if it is necessary to choose between having a diverse set of experts and having a well - behaved group, it is best to go for the diversity. The SSHAC experience, buttressed by extensive decision analysis experience, suggests that if the meeting is structured appropriately and the goals are communicated appropriately, the meeting can proceed without rancor, even if the experts substantially disagree on scientific matters.

The group meeting can be conducted in two or three days. The TFI, however, needs several months to design the meeting and the process, to identify the experts, to solicit participants, and to set up the necessary logistics for the meeting. The reviewers of the methods will require several months to prepare their papers. It is important to take into account the relatively small community of leading ground-motion experts. It is critical to enlist the leading world experts, and this generally necessitates a long lead time.

Basic products of the Review of Methods stage include:

- A set of methods for predicting ground motions clearly understood by the TFI and experts
- A list of specific disagreements, not necessarily resolved, but which are clearly understood and documented

- A set of representative models to be used to forecast ground motion at specific (M, R) points
- The set of (M, R) points to be used in the Stage 4 elicitation
- A set of proponents identified by the TFI to run the models to produce the Stage 4 predictions

Stage 4: Elicitation of Ground Motion at Points in (M, R) Space

Stages 4 and 5 are the heart of the elicitation process. They provide the basic material to be used by the TFI to produce the composite ground-motion estimator. This is done by concentrating ground-motion estimates and discussion on specific (M, R) pairs (determined during and after Workshop #1). The stages are conducted at a group meeting (Workshop #2), smaller than that used for Stages 2 and 3. The workshop must be designed so that information is provided to the TFI in a way that promotes extensive feedback and discussion among experts.

The process starts with a small group of experts in the role of *proponents* who are asked to perform a detailed ground-motion analysis based on a specific model, and then to interpret the results. These proponents should be intimately familiar with particular ground-motion models (generally, these proponents have published these models) providing ground-motion estimates at the (M, R) pairs to the TFI in advance of the meeting, using a specific model. The purpose of the proponent role is not to create advocates, but to create a clear understanding of each model and its results. The proponents also provide a written description of the assumptions and modeling details, as well as an explicit account of the dataset upon which the estimates were based. If possible, proponents should provide numbers and pictures showing what the results would be—based on data only—and compare that to the model estimates. These results may differ. For example, the median ground-motion estimate for one distance may be the result of fitting a curve that applies to a range of distances, whereas the data alone would apply just to that distance.

The TFI needs to make sure that the data used by each proponent are summarized in a form so that easy comparisons can be made. The estimates are displayed by the TFI (and/or resource experts) in a consistent graphical format, along with the data agreed upon in Stage 2. It is useful to include bars representing epistemic uncertainty for each proponent.

The graphs and proponent documentation are distributed, before Workshop 2, to a slightly larger group of experts. This larger group may include the proponents, who would now be asked to wear a different hat, that of *independent evaluator*. The experts-as-evaluators provide estimates for the same (M, R) pairs, using whatever combination of models they wish. Their results are sent to the TFI in advance of Workshop 2. The TFI again prepares graphs showing the various estimates, but they do not have to be distributed before the workshop.

At Workshop 2, the agenda is organized around a discussion of the ground motion estimates. Once again, the TFI attempts to isolate and then focus on areas of strong agreement and disagreement. The purpose is not to achieve consensus (although that is a good outcome if it is a true consensus) but rather a detailed understanding of the rationale for underlying differences so that the experts can each construct an informed composite representation of their own and the overall expert community's state of knowledge. The discussion should illuminate and eliminate any unintended disagreements. Typically, the experts will want to reconsider their estimates after the group discussion. This can be done informally (say overnight) at the workshop, but then needs to be done more carefully immediately after the workshop. Similarly, the TFI may need a round of individual interactions after the group meeting to make sure that the basis for the proponent and expert assessments are fully understand.

Time needs to be allocated for enough iteration so that the TFI can come to a full understanding of the basis for the ground motion estimates. Several months at least are required before Workshop 2 is held for the proponents to (1) perform the model runs, (2) document assumptions and results, (3) receive feedback from the TFI and the other

experts, and then (4) react to the feedback and prepare for the meeting. The meeting itself will typically last two to three days (the SSHAC meeting lasted two days). Roughly a month should be allocated afterwards for individual interactions among the TFI and experts, and for final estimates of the ground motion measures.

The basic product of the model estimate interaction is a well-understood median estimate of ground motion for each specified (M, R) pair, for each model, and for each expert.

Stage 5: Elicitation of Uncertainties

The assessment of uncertainty ranges for both the median and standard deviation of the aleatory distribution on ground motion is naturally, but not necessarily, done in conjunction with the median estimate interaction (Stage 4). Assessments by each proponent and each expert need to be encoded for at least three variables:

- Epistemic uncertainty in median estimate
- Best estimate of the aleatory uncertainty
- Epistemic uncertainty in aleatory uncertainty

If possible, and if there is sufficient time to make sure that all experts are sufficiently well grounded in the concept, it may be appropriate to decompose the assessments of epistemic and aleatory uncertainties into whether they are parametric or modeling uncertainties (see the discussion in Chapter 2). Chapter 2 explains why these different types of uncertainties are needed for a complete specification of the overall uncertainty in ground motion.

The uncertainty-assessment interaction is based on the same paradigm as the median-ground-motion-estimate interaction. At the SSHAC Ground-Motion Workshop 2, the agenda was divided into three major interactions:

- Estimate of median ground motion
- Epistemic uncertainty in the median estimate
- Best estimate and epistemic uncertainty in aleatory uncertainty

Ground-motion experts will not all be familiar with the meanings of these different types of uncertainties, nor will many of them have much experience in or knowledge about probability elicitation. Thus, individual interactions with the experts are required prior to the group meeting to provide education and training in probability elicitation and to elicit probability distributions from the experts (this type of training is described in Chapter 3).

During and after the workshop, the TFI needs to make sure that the experts' probability assessments accurately reflect their true state of information about the uncertainties. At a minimum, two fundamental consistency checks need to be performed:

1. If an expert's uncertainty range (for either the median or the aleatory uncertainty) is narrower than the range of estimates from all experts, the TFI needs to make sure that the expert truly attaches little or no significance to the estimates falling outside the range. In general, it is inconsistent to attach an uncertainty band that is much narrower than the estimates of the set of models and experts that are viewed as credible (this issue is discussed in more detail in Chapter 3). This is especially true when the experts are attempting to represent, not just their own position, but the composite position of the overall community.
2. It is also important to challenge experts whose uncertainty ranges are far greater than the range of model and expert estimates. Such an assessment implies the forecasting error associated with the individual models is quite high—but this would imply that the observed tighter pattern of model estimates is an unlikely coincidence, unless the assessor believes that there is a great deal of correlation among the model forecasts.

The experts need to be comfortable with probability assessment. For this, training in probability assessment and an experienced elicitor is very important. In general, it is important that the variables being elicited are variables with which the experts are intimately familiar.

Generally, experience in other fields suggests that it is far better to assess real observable variables than parameters of complex models or moments of probability distributions (e.g., mean and standard deviation). Due to the way in which data are processed and models are constructed, ground-motion experts seem to find that assessing moments is more natural than specifying a probability distribution directly on ground motion.

The participants and time required for Stage 5 is the same as for Stage 4, since the group meetings for both stages will be held during the same workshop.

The basic product of the uncertainty interaction is a set of probability distributions for each expert on each of the three variables described above. If the elicitation sessions generated influence diagrams and/or conditional distributions, the conditioning and conditional probabilities should be specified explicitly. The TFI should fully understand each expert's rationale underlying the probability assessments. The rationale should be documented by the experts themselves.

Stage 6: Development of Ground-Motion Distributions by the TFI

The basic paradigm for integration is to weight (or weigh) the estimates provided by each model for each (M, R) pair, guided by (1) how the experts as individual evaluators weight the model estimates, and (2) by how the experts integrate all available information into a composite representation of the community. If the experts' estimates as individual evaluators or as integrators are disparate, it is crucial that the TFI understand the sources of the differences before making any final decisions. Once the sources of disagreement are noted, the TFI then has to carefully weigh the strength of the logic, the underlying interpretations, and existing data.

The TFI should carefully consider, step-by-step, each expert-aggregation issue discussed in Chapter 3. It is also useful to apply simplified aggregation models, described in Chapter 3, but these should be viewed as providing rough guidelines only. The value of applying these simple models, especially for TFI team members who are less familiar with probability elicitation principles, is to see how each basic issue can affect the final aggregated probability distribution.

A useful step is for the experts to write down explicitly their judgment about the relative forecasting abilities of the various models and how much overlap or similarity there is between different classes of models. Verbal interaction provides a great deal of information on the rationale for why different experts place different weight on different models, but it is important to quantify these judgments both to ensure that the TFI understands the various positions and to make sure that the experts themselves are thinking consistently about the issues. The survey for the Ground-Motion Workshop 2 provides a starting point for such a quantification (the survey and its results are discussed in Appendix B), although the questions need to be rephrased to be clearer for the experts.

After the TFI is comfortable that the basis for each model and each expert's position is fully understood, it is useful for the TFI to form a preliminary position on the final composite estimates and distributions. (See Chapter 3 on general TFI integration guidance, and Section 5.7.5 below, for a discussion of why equal weights on TI positions is a desirable and likely outcome.) The TFI should document carefully the rationale for the composite representation and present it to the proponents and experts. If resources and time are available, it is best to do this in face-to-face meetings (individually or group); if not, written feedback from the experts is sufficient.

Finally, after receiving feedback from the experts, the TFI team members need to work closely together to construct the final composite representation. Typically, in this step the TFI should continue to interact closely with individual experts to make sure the final representation is based on a complete and accurate state of information.

The final interaction with experts could be done potentially in a several-day meeting, but may require several rounds of individual interaction.

The TFI team members should not rush into a final decision, but should probably iterate with several working meetings. It is essential to make sure that the TFI team completely understands the basis for everything they are integrating.

The product is, for each application: a probability distribution on the median and standard deviation ground motion at each point in (M, R) space. Both the probability distributions and a full and detailed description of the basis for them should be fully documented.

5.7.4 The Technical Integrator (TI) Approach

Several recommendations are appropriate for the case when the Technical Integrator is used to specify ground-motion input to a PSHA. This TI could work as a single entity, using its own expertise (or that of a consultant) to identify ground-motion attenuation equations, or the TI could informally use multiple outside experts to provide input and guidance on the selection and evaluation process. In some cases (for example, in coastal California studies), the equations considered may be entirely taken from the published literature. In other cases, new equations will have to be derived, most likely by modifying similar equations derived from different regions.

Regarding the choice of appropriate ground-motion equations to use for the study, both empirical and analytical equations should be considered. The ultimate choice of equations and how they are used will depend on the region of the study, on available attenuation equations for that region, and on the degree to which attenuation equations from other regions must be adapted to represent characteristics in the study region.

While the TI is not constrained to use explicit numerical weights (i.e., as with the TFI, the TI may choose to "weigh" rather than "weight"), when dealing with multiple models such an approach is encouraged when appropriate. Explicit weights are usually simpler to apply and easier to explain than other aggregation schemes (see the next section for more detail). Compared to the TFI, the TI may have less time and resources, and less control over proponents of different models. Eliciting weights from other experts is usually a simple and straightforward task to perform and the results are easily documented.

In California, it is likely that empirical and analytical equations will be similar for those ranges of magnitude and distance for which numerous data are available with which to calibrate these equations. Thus, the specific weights and credibilities assigned to each equation likely will not be critical. In particular, for distances close to large magnitudes (e.g., distances less than 20 km and magnitudes greater than 7.3), data are lacking and analytical results may differ from empirical results, and different empirical equations may themselves differ because of different extrapolation techniques. In this case, the form of the equation and its consideration of large-magnitude effects (e.g., finite fault rupture) and close-distance effects (e.g., the geometry of the site relative to the fault) must be considered in assessing credibilities of predicted ground motions. For these situations, we recommend that both empirically and analytically based models be considered when selecting the group of attenuation equations used for the PSHA study. Detailed analytical results may provide guidance on appropriate magnitude and distance scaling for large magnitudes and/or short distances, even if the analytical results are not themselves finally used in the hazard calculations.

In many applications, analytical equations have been used with success, but there are differences among available models and among the parameters used as input to those models. An example is the conversion used to relate moment magnitude M to body-wave magnitude mbLg. When such differences exist, they should in general be treated as epistemic uncertainties and both models must be included in the study. Also, the crustal or source parameters for a specific region may be different from generic parameters derived for broader regions. If wave-propagation studies are used to develop attenuation equations, the TI must include all relevant uncertainties in crustal properties. The TI must take care not to

underestimate the epistemic uncertainty in ground-motion prediction for a region with few or no existing ground-motion equations. In such a region, epistemic uncertainty will be high, and the existence of only one or a few (matching) equations is not evidence that epistemic uncertainty is low.

On the other hand, the TI also must not overestimate the uncertainty, just because no empirical observations of strong ground motion are available. Analytical studies conducted in the last decade, calibrated with low-amplitude seismograph records, have gone a long way toward providing an understanding of earthquake ground motions in the central and eastern U.S.

Aleatory uncertainty is relatively easy to estimate for California, where empirical observations are available to quantify scatter about predictions. Here the TI must be careful to incorporate any magnitude dependence of ground-motion scatter into the predictions, as published by the authors of each equation considered.

For other regions of the country, the aleatory uncertainty cannot be estimated from strong-motion observations. The TI may adopt aleatory distributions from California, using similar distributions for the remainder of the country and including, if deemed appropriate, epistemic uncertainty on the magnitude dependence of the aleatory uncertainty. An alternative is to model aleatory uncertainty by estimating aleatory distributions for input variables (such as stress drop) to analytical models, determining the resulting aleatory uncertainty on ground motion as a function of magnitude and distance, and confirming that the derived distribution is not inconsistent with similar distributions from California. There is certainly more epistemic uncertainty in ground motions outside of California, but unless there is a compelling counter-argument, the aleatory uncertainty should be similar to that in California equations.

5.7.5 Step-by-Step Guidance for Ground-Motion Integration

This section contains some summary step-by-step guidance for how ground-motion integration can be performed. The guidance is based on the successful process used to integrate the results of Ground-Motion Workshop 2 (see Appendix B). The guidance is "recommended" rather than "required" because the procedure has been applied in only one workshop.

The process, as performed in GM Workshop 2, works by first comparing model estimates and expert estimates. For simplicity, we will use the word *model* to denote the estimate produced by the model runs of each model proponent and the word *expert* to denote the composite estimate produced by each expert playing the role of integrator. The steps are as follows:

1. TFI posits an intuitive "quick and dirty" median ground-motion estimate.

2. This estimate is compared graphically to the experts' composite median estimates and all the models' median estimates, all overlaid on top of a plot of the available data. This should give the TFI an initial idea as to potential integration problems, if any.

3. The result of equally weighting the experts is then compared to the result of equally weighting the models. This step is for insight only. SSHAC expects equal weights on experts often to be appropriate in forming the final composite distribution. Equal weights on models are almost never appropriate.

4. If unequal weights are still a consideration after steps 1-3 (and after once again interacting with the experts to understand meaningful differences), a range of representative unequal weighting schemes on experts are applied and compared to the equal weighting results. Alternatively, experts are clumped into different subgroups felt to have potentially correlated assessments (Chapter 3), and the subgroups are equally weighted. A sensitivity analysis is performed to see if the different weighting and clumping schemes matter. The point of this step is to determine whether precise unequal weights really matter.

5. If unequal weights on experts are appropriate and the composite estimate is sensitive to the

likely range of unequal weights, then the representativeness of each expert relative to the overall community needs to be evaluated explicitly by the TFI team, again, in consultation with individual experts, as needed, and incorporated into the weights (Chapter 3). Also, although group interaction should have minimized differences in relative interdependence among subgroups of experts and differences in knowledge with respect to the specific application, these issues should be reviewed as well.

6. A final estimate for each (M,R) pair is established.

7. A similar process is used to produce uncertainty ranges.

Recall from Chapter 3 that a well-run facilitation process is expected to result in a defensible and simple equal-weighting process. Thus, SSHAC expects Steps 4 and 5 to be necessary only rarely.

On the more general issue of explicit numerical "weighting" versus non-explicit "weighing" SSHAC's position (discussed also in Chapter 3) is that, while explicit (equal or unequal) numerical weighting is not required, it is a desirable way to arrive at the final TFI estimate for several reasons:

- Explicit weighting provides a decomposition that helps explain how the TFI position was determined
- The TFI position can be explicitly compared to the experts' integrator positions
- Requiring explicit weights on models from experts that they must defend tends to lower the possibility of eliciting extreme and/or non-defensible judgments.
- There are probabilistic expert aggregation models, that, while simplified, provide theoretical underpinnings to the weighting process (these aggregation models can be applied to either experts or ground-motion models).

6. METHODOLOGY FOR CALCULATING SEISMIC HAZARD ESTIMATES AND UNCERTAINTIES

The purpose of a probabilistic seismic hazard assessment (PSHA) is to provide a quantitative assessment of seismic hazard, described by the likelihood that various levels of earthquake-induced ground motions at a site will occur or will be exceeded at a given location in a given future time period. The outputs of a PSHA are estimates of seismic hazard curves, i.e., plots of the estimated probability per unit time of a ground motion variable, e.g., PGA, being equal to or exceeding level a as a function of a. A typical output is a set of curves describing the estimated seismic hazard in terms of curves of the marginal 5th, 50th, and 95th epistemic uncertainty (probability) fractiles of the estimated probability per unit time, P(A>a), as illustrated in Figure 6.1.

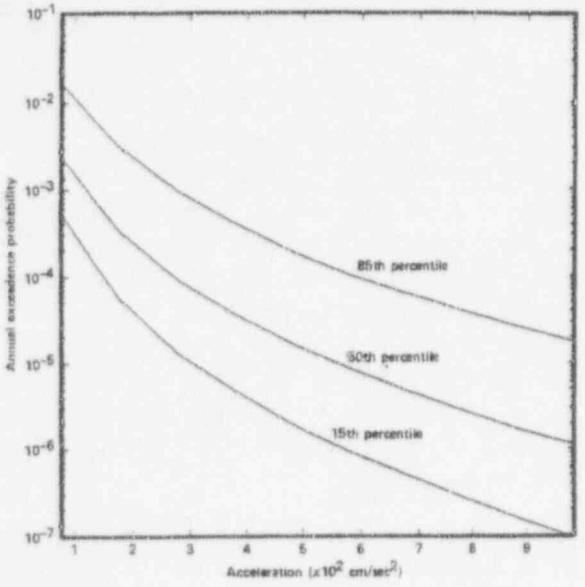

Figure 6-1 Seismic Hazard Curves

Other outputs of a PSHA are described in Chapter 7. Estimation of seismic hazard is based on integrating the seismic source characterizations and ground motion estimates described in Chapters 4 and 5. The mathematical models which form the basis of the seismic hazard calculations are discussed in Chapter 2.

Calculation of the estimated seismic hazard curves involves summing integrals of probability (aleatory uncertainty distributions over multiple seismic areas) is shown in Equation 2.2. The parameters of these distributions are a function of the seismic source characterization and ground motion inputs which are not known exactly (that is, there is epistemic uncertainty); thus, we only have estimates of the distributions. An important part of the calculation involves the quantification of the epistemic uncertainty associated with the estimated seismic hazard due to the epistemic uncertainties in the seismic source and ground motion inputs.

The calculational methodology of assessing the estimated seismic hazard and quantifying its epistemic uncertainty is a nested process. The inner operation is the basic calculation of integrating a specific seismic source characterization with a specific set of ground motion distributions to produce a single estimated seismic hazard curve. This calculation is discussed in Section 6.1.

The outer operation is an uncertainty analysis involving the propagation of the epistemic uncertainties associated with the seismic source characterizations and ground motion distributions to develop the probability (epistemic uncertainty) distribution for the estimated seismic hazard. This epistemic uncertainty is generally assessed in terms of a joint probability distribution of ($PA \geq a$) for a finite number of levels of ground motion, i.e., values of a. The joint probability distribution is sometimes needed for propagating epistemic uncertainty when seismic hazard and fragility are combined in a PRA (ANS/IEEE 1993). One description of epistemic uncertainty is illustrated in Figure 6.1, in which the results of the PSHA are presented as curves of the appropriate fractiles of the marginal (epistemic) probability distributions of the estimated seismic hazard. Calculational methods for the uncertainty analysis are the topic of Section 6.2. Propagation of the epistemic uncertainty in the uncertainty analysis assumes the seismic source characterizations and ground motion inputs are each derived from a

single expert, the TI or TFI. If either or both of these inputs are derived from multiple experts and the multiple inputs are not aggregated prior to doing the calculations, the multiple experts can be thought of as another source of epistemic uncertainty and treated accordingly in the uncertainty analysis. Eliciting and aggregating seismic source and ground motion inputs from multiple experts is discussed in Chapter 3.

6.1 Integration Methods For Producing Seismic Hazard Estimates

The inner loop of the overall seismic hazard estimation process involves the integration of a specific seismic source characterization (SSC) with a specific set of ground motion distributions for all magnitudes and distances to produce a single estimated seismic hazard curve for a given ground motion parameter. Before discussing calculational methods for integrating seismicity and ground motion information to produce an estimate of seismic hazard, it is appropriate to summarize the inputs necessary for seismic hazard evaluation and to reiterate the mathematical identity, discussed in Chapter 2, which is the basis for producing the estimated seismic hazard.

6.1.1 Seismicity and Ground Motion Inputs

6.1.1.1 Seismic Source Characterization

A description of the seismicity throughout the region affecting a site is characterized by

- Seismic source representation of the region, a seismic map

  For purposes of the seismic hazard calculations, the essence of the seismic source representation is that the region be partitioned into areas of homogeneous seismicity, referred to, in this discussion, as "seismic areas." As discussed in Chapter 4, sources are categorized as faults (type 1) or areal (types 2-4) sources. A seismic map is a partition of the region of interest into seismic areas, i.e., into areas of homogenous seismicity.

  Two approaches to developing a seismic map are:

  a. (Bernreuter, D. L., et al. 1989) If the seismic source representation is a partition of the region into "source zones," either faults/fault segments and/or areal sources which are assumed to be areas of homogeneous seismicity, i.e., earthquake expected frequency and magnitude distribution are considered to be homogeneous throughout the zone, a seismic area is equivalent to a source zone. A seismic map is equivalent to a source zonation of the region.

  b. (EPRI 1988) If the seismicity representation is a dual partition of the region into (1) seismic sources to which is associated a maximum magnitude and (2) areas (1°x1° squares in the application by EPRI) in which it is assumed the magnitude-recurrence relation is constant, i.e., earthquake expected frequency and magnitude relative frequency are homogeneous throughout the area—seismic areas are the intersections of seismic sources and the areas in which the magnitude-recurrence relation is constant, i.e., sections of sources of assumed homogeneous seismicity.

- Vector of seismicity information for each seismic area, i.e., area of homogeneous seismicity. The elements of seismicity are:

  - Expected frequency of earthquakes within the area, per time period, of magnitude exceeding a minimum magnitude m_o (e.g., m_{blgo} =5.0); ν
  - Maximum magnitude; m_u

- Magnitude distribution and its parameter(s). Two alternative models are the truncated exponential model and the characteristic earthquake model. Both models involve a parameter β relevant to the exponential portion of the model [$\beta = b\ln 10 = 2.3b$, where b is the slope in the familiar Gutenberg-Richter relation]; the characteristic earthquake model also requires the magnitude/range of magnitudes and frequencies of the characteristic earthquake.

6.1.1.2 Ground Motion Distribution

Assuming that, conditional on magnitude and distance, the distribution of the ground motion parameter is a lognormal distribution, the ground motion is characterized by:

1. The natural logarithm of the median of the ground motion parameter as a function of magnitude and distance, generally given in terms of the value of a ground motion model: $g(m, r)$
2. The (aleatory) standard deviation of the natural logarithm of the ground motion parameter, possibly a function of magnitude and distance: $\sigma(m, r)$

6.1.2 Basic Seismic Hazard Identity

Equation 6.1 shows that, based on the "usual" assumptions of occurrences of earthquakes, seismic hazard $P(A \geq a)$, as a function of ground motion level a, is [see equations 2.2 and 2.3]:

Equation 6.1

$$P(A \geq a \text{ in time } t) = 1 - \exp\left[-\sum_{i=1}^{S} \nu_i t \iint \Phi'\left(\frac{\ln a - g(m,r)}{\sigma} \right) f_{R_i}(r|m) f_{M_i}(m)\, dr\, dm \right]$$

$$\equiv \sum_{i=1}^{S} \nu_i t \iint \Phi'\left(\frac{\ln a - g(m,r)}{\sigma} \right) f_{R_i}(r|m) f_{M_i}(m)\, dr\, dm \qquad (6.1)$$

where:

- S is the number of seismic areas
- ν is the expected frequency, per time period per seismic area, of earthquakes of magnitude at least m_O.
- $\Phi'(.)$ denotes the standard normal complementary cumulative distribution function (CCDF) which is based on the usual assumption that the ground motion parameter is a lognormal aleatory variable. The ground motion distribution is possibly truncated.
- $f_M(.)$ denotes the probability density function of the magnitude distribution.
- $f_R(.)$ denotes the probability density function of distances, from the site, of the locations of earthquakes, given an earthquake occurs in the seismic area. A commonly accepted model is based on assuming earthquakes occur spatially "at random" within a seismic area, thus, f(r)dr represents the proportion of the seismic area at distance r from the site. For some cases, when the seismic area represent a fault/fault segment, the earthquake rupture is represented as a plane instead of a point and if the rupture length depends on magnitude, the distance also depends on magnitude. Thus, the density function is written as $f_R(r|m)$.

6.1.3 Integration Methodology

As shown in the seismic hazard identity in Equation 6.1, the calculation of the estimated seismic hazard involves a sum of two-dimensional integrals. The standard normal CCDF, $\Phi'(.)$, is also an integral so the calculation could be written as a sum of three-dimensional integrals. However, $\Phi'(.)$ is readily evaluated by packaged subroutines. The standard procedure for integrating functions is numerical integration, also called quadrature. This involves partitioning the magnitude and distance ranges into a finite number of subintervals, evaluating the integrand at a selected point(s) within each interval, and summing weighted products of the integrand and the subinterval width. Several one-dimensional quadrature methods exist, including rectangular, trapezoidal, and Simpson's rules, spline quadrature and Gaussian quadrature, corresponding to various orders of accuracy. For the multidimensional estimated seismic hazard calculation, because of the polygonal geometry of the seismic areas and the frequent numerical input of the ground motion models, simple numerical integration methods are generally used.

Considering the levels of uncertainties associated with the inputs into a PSHA, the choice of numerical integration method does not seem to be critical to the analysis. Two important elements of the calculation are developing and keeping track of the geometry of the seismic areas and evaluating the probability density function, $f_R(r)$, of the distance, which, for the commonly accepted model, involves assessing the proportion of a seismic area corresponding to distance r from the site. The calculation also requires a specification of the ground motion distributions for all (m, r). If the medians of the distributions are expressed in terms of a single ground motion model or a set of weighted models, there is no problem. However, as discussed in Chapter 5, the median ground motion may be based on deriving the medians for a finite subset of the (m, r)s. In that case, it is necessary to interpolate between the inputted values to evaluate the median at the non-inputted (m, r)s. One approach mentioned in Chapter 5 is to fit a model, similar to an accepted ground motion model, to interpolate. This is quite reasonable when interpolating between "most likely" or "best estimate" values of the median ground motion. A more difficult issue is how to represent epistemic uncertainty in the median at the interpolated (m, r)s, given epistemic uncertainty in the medians at a selected subset of (m, r)s. The issue is how to represent epistemic uncertainty which accounts for the epistemic correlation associated with the inputted medians. (Note: this is also an issue if a single model uses uncertainty bands to represent epistemic uncertainty.) This is discussed in Section 6.3.

Details on some of the practical aspects of the integration calculation are included in the reports of past PSHAs (e.g., EPRI 1988; Bernreuter, D. L., et al. 1989).

The output of the integration is a single estimated hazard curve, $P(A \geq a)$ as a function of a[1] which represents the estimate of seismic hazard given the specific values of the uncertain inputs. All other products of the PSHA can be derived from the basic seismic hazard calculations or can be evaluated following the same basic concept. Some of the most important such products are deaggregated hazard results and sensitivities (see Chapter 7). Numerical methods for these products are discussed in EPRI 1988 and Bernreuter, D. L., et al. 1995.

6.2 Propagation of Epistemic Uncertainty

Calculation of the estimated seismic hazard curve, as outlined in Section 6.1, is based on a specific seismic source characterization and set of ground motion inputs, i.e., on a specific seismic map, specific values of the seismic parameters for all seismic areas, and specified ground motion distributions for all magnitudes and distances. Since all these inputs are uncertain, i.e., subject to epistemic uncertainty, it is necessary to reflect this epistemic uncertainty in the estimation of seismic hazard. Quantifying the epistemic uncertainty associated with the estimated seismic hazard due to the epistemic uncertainties associated with the seismic source

[1] Note: computationally, the results are vectors of values of $P(A \geq a)$ for a finite set of values of a.

characterization and ground motion distributions is based on propagating these input uncertainties through the seismic hazard calculations to establish the epistemic uncertainty in the estimated seismic hazard.

6.2.1 Descriptions of the Epistemic Uncertainties of the Inputs

Recommended and/or alternative methods of describing the epistemic uncertainties associated with the inputs into a seismic hazard calculation are discussed in the chapters on seismic source characterization and ground motion estimation, Chapters 4 and 5 respectively. Epistemic uncertainty descriptions, as they relate to the elements of the seismic hazard calculations described in Equation (6.1) are summarized here. It is recognized that the basic input information may not always be expressed in terms of the distributions included in Equation (6.1). For example, magnitude distribution inputs may be derived in terms of the parameters (a, b) of the Gutenberg-Richter relation instead of the expected frequency ν and parameter β in Equation(6.1). This is not a problem since the propagation of uncertainties can be based on the uncertainties of the original parameters or on the uncertainties transformed to uncertainties in the parameters of the distributions given in Equation (6.1). The only requirement is that the epistemic uncertainties be completely quantified and that any potential correlations in epistemic uncertainties be recognized and properly handled in the propagation of the epistemic uncertainty.

6.2.1.1 Seismic Source Characterization

The basic epistemic uncertain inputs are the seismic map and the vectors of seismicity parameters, maximum magnitude, earthquake expected frequency and magnitude distribution parameter(s), for each seismic area in the map.

1. Seismic maps: epistemic uncertainty is described in terms of alternative maps based on alternative representations of faults and/or areal sources, accounting for the probability of activity/existence, and alternative source geometries with weights. The number of alternative seismic maps can be significant, especially for regional studies, and can involve considerable computational effort to develop, maintain, and track depending on the number of seismic sources, alternative source boundaries, and the probabilities of activity/existence provided. Ways to reduce the number of alternatives should be considered, e.g., eliminating alternatives with low weights and combining maps which only differ in areas with an insignificant impact on the hazard value.
2. Seismicity parameters for each seismic area:
   - maximum magnitude, m_u: epistemic uncertainty is described in terms of a discrete or continuous probability distribution for m_u for each seismic area. Epistemic correlations between m_us for different seismic areas must be recognized. Such correlations could arise due to seismological considerations or as a result of the seismic sources and seismicity representation, e.g., in the EPRI representation of seismicity, several seismic areas are sections of the same source, hence have the same maximum magnitude, thus their epistemic uncertainty is perfectly correlated.
   - expected earthquake frequency (of magnitudes at least m_o), ν: epistemic uncertainty is described in terms of a discrete or continuous probability distribution.
   - magnitude distribution parameter(s), β (also the characteristic magnitude and frequency, if appropriate): epistemic uncertainty is described in terms of a discrete or continuous probability distribution.

It should be recognized that, for a given seismic area, the latter two parameters, ν and β, are likely to be epistemically correlated. This correlation has the most significant effect on the estimated seismic hazard since it affects the mean hazard as well as the epistemic uncertainty in the estimated seismic hazard. Thus, it is important that this correlation be recognized and accommodated in the uncertainty analysis. This is discussed further in Section 6.3.

For the exponential model (truncated exponential model or the exponential portion of the characteristic earthquake model), alternative parameters to (ν, β) are (a, b), the parameters of a magnitude-recurrence equation. These pairs of parameters are related and it is possible to transform the epistemic uncertainty in one pair to epistemic uncertainty in the other. Again, the parameters (a, b) are epistemically correlated; this correlation must be propagated to correlation between ν and β.

As with the maximum magnitude, if there is epistemic correlation in the seismicity parameters between seismic areas, this must be recognized and accounted for in the specification of epistemic uncertainty. For example, in the EPRI representation of uncertainty and the introduction of "smoothing" between 1°x1° areas, a potential correlation of ν and β between seismic areas is implied.

6.2.1.2 Ground Motion Distributions

The basic ground motion input parameters are the median ground motion and the aleatory standard deviation of the logarithm of the ground motion for all (m, r).

1. Median: epistemic uncertainty is described in terms of a discrete or continuous probability distribution for g(m, r) for all (m, r). Again, it is likely that the estimated median ground motion for different (m, r)s are correlated. This must be recognized and accommodated in the uncertainty calculations.

2. Standard deviation: epistemic uncertainty is described in terms of a discrete or continuous probability distribution for all (m, r), and, if appropriate, recognition and accommodation of epistemic correlation.

If the ground motion inputs are based on ground motion models containing aleatory variables, the inputs include the aleatory variable distributions and their uncertain parameters. These uncertainties must be properly analyzed to assess uncertainty in terms of uncertainties in the median and aleatory standard deviation.

Epistemic correlation in the median and standard deviation between different (m, r)s is a second order effect with regard to the propagation of epistemic uncertainties in PSHAs since it does not affect the mean seismic hazard. Recognition of epistemic correlation in the median and aleatory standard deviation is important only for quantifying the epistemic standard deviation of the estimated seismic hazard. This is discussed further in Section 6.3.

6.2.2 Epistemic Uncertainty Propagation Methods

Conceptually, estimated seismic hazard, i.e., $P(A \geq a)$, is a function of all of the inputs. Since the inputs are epistemically uncertain, $P(A \geq a)$ is a function of a set of probabilistically distributed parameters. There are several methods for propagating this probability through the seismic hazard calculations to derive the (epistemic) probability distribution of $P(A \geq a)$. Two classes of methods are:

- Analytic methods

  Since $P(A \geq a)$ is a function of probabilistic inputs, conceptually, one method is by "transformation of variables." This is not practical for PSHA because of the complexity of the functional relationship. A second type of analytic approach is based on evaluating the moments of the probability distribution of $P(A \geq a)$ in terms of the moments of the probability distributions of the inputs. The classical procedures include the method of moments, Taylor series expansion and response-surface methods.

- Sampling methods

  1. Complete enumeration: This approach is usable if the probability distributions of all inputs are expressed in the format of discrete distributions. Conceptually, the method involves evaluating $P(A \geq a)$ for all combinations of values of the epistemic uncertain parameters/inputs. The probability associated with the resulting value of $P(A \geq a)$ is the product of the probabilities (properly combined to account for epistemic correlations)

associated with the inputs. The resulting set of values and associated probabilities of $P(A \geq a)$ is the (epistemic) probability distribution of $P(A \geq a)$. The number of combinations may become very large if there are a large number of inputs and/or if there are a large number of alternative values for each input. This is recognized in a version of this approach, referred to in the PRA literature as the Direct Probability Distribution (DPD) method (S. Kaplan 1981). To reduce the number of calculations, this version of the method also involves aggregating values and probabilities of intermediate calculations.

If continuous probability distributions are discretized to apply the complete enumeration method, the proper choice of representative values and probabilities is important to derive an adequate estimate of the probability distribution for seismic hazard.

A convenient graphical tool for enumerating the combinations is a logic tree consisting of nodes identifying the uncertain inputs and branches, representing the alternative "values" of the inputs, emanating from each node. A "limb" or continuous connection of the branches for each of the parameters represents a combination. (See EPRI 1988 for an application of logic trees.)

2. Random sampling: Applicable if the probability distributions of the inputs are either continuous or discrete. Conceptually, the methodology is based on sampling values from the probability distributions for each of the inputs and assessing the corresponding value of $P(A \geq a)$. The most straightforward approach is "Monte Carlo sampling," which is based on simple random sampling of each of the inputs. In order to represent satisfactorily the 5th and 95th fractiles of estimated seismic hazard, a minimum sample size of 200 is recommended. If, however, preliminary results show that the mean hazard lies above the 95th fractile, the sample size must be increased appropriately.

   Alternative approaches, based on restricted random sampling, are designed to be more efficient, i.e., for the same sample size to provide an estimate with a lower sampling variability. One method is "importance sampling." This approach is based on sampling the inputs and values of the inputs which are most important, i.e., the inputs to which the seismic hazard is most sensitive. Thus, it requires some knowledge of the sensitivities of the seismic hazard with respect to the uncertain inputs. Another method is "Latin hypercube sampling" (LHS). The basic concept of this approach is to partition the range of each input into the same finite number of equiprobable subintervals and sampling subintervals at random such that each subinterval (of each parameter) is sampled only once. Within a subinterval the sampled value is selected at random. This approach assures that the entire range of each parameter is represented in the sample (R. L. Iman and M. J. Shortencarier 1984). These alternative approaches have not been used in past PSHAs.

Both the complete enumeration and Monte Carlo sampling methods have been used in past PSHAs. They both can, with proper care, be effective computationally and are acceptable for developing the joint (marginal) probability distribution(s) of $P(A \geq a)$ for a finite number of values of a. This distribution is the basis for the fractile curves for the seismic hazard (Fig. 6.1).

The complete enumeration method and the use of accompanying logic trees provides a more transparent presentation of alternative hypotheses and values as well as displaying sensitivities to alternative inputs. If discrete probabilities are used to represent "continuous" epistemic uncertainties, care must be taken to assure that the discretization, and accompanying loss of information, is not too coarse. Monte Carlo sampling is, in general, more efficient and is most useful for large regional studies.

Uncertainty analysis, as discussed, is based on deriving input information from a single pair (one seismicity and one ground motion) of resources. If the overall methodology involves combining inputs from multiple resources, e.g., multiple SSC experts, using some kind of "weighting" algorithm, the relative weights can be treated as a probability to include the epistemic uncertainty associated with the multiple resources in developing the overall probability distribution associated with the estimated seismic hazard.

6.3 Discussion and Recommendations

PSHAs involve extensive computer calculations requiring considerable bookkeeping to handle the multiple summations associated with numerical integration, the potentially large number of seismic areas, and, in the uncertainty analysis, the alternative SSC characterizations and ground-motion uncertainties. It is important to develop the proper combinations of inputs to assure that the models are correctly represented in the calculations. This is particularly important for the uncertainty analysis.

A potentially difficult issue is the representation of epistemic correlation. One way of measuring the significance of potential epistemic correlations is to evaluate their effects on the estimated mean seismic hazard and the epistemic standard deviation of seismic hazard. Since the mean hazard is an important input into a Probabilistic Risk Analysis (PRA) and in design ground motion criteria, for purposes of this discussion, an epistemic correlation is considered to have (1) a first-order effect if it affects the epistemic expected value (i.e., the estimated mean) as well as the epistemic standard deviation and (2) a second-order effect if it only affects the epistemic standard deviation. Three epistemic correlations are important:

1. The epistemic correlation in the median (and standard deviation) of the ground motion parameter between different (m, r)s.
2. The epistemic correlation between the earthquake occurrence rate and magnitude distribution parameter, i.e., (ν, β), within a seismic area. [Equivalently, the epistemic correlation between (a, b) or between the expected frequencies at two (or more) magnitudes within a seismic area].
3. The epistemic correlation in a seismicity parameter, e.g., m_u , ν, or β, between different seismic areas.

If one considers approximating the integrals in Equation (6.1) by sums, it is clear that the seismic hazard calculation involves summing over seismic areas, magnitudes and distances. Since the expected value of a sum is the sum of the expected values of the summands, the correlation between summands does not affect the mean hazard, whereas correlation affecting the means of the summands does. Therefore, correlations such as the first and third type above do not affect the mean hazard, while the epistemic correlation between (ν, β) within a seismic area (i.e., type 2 above) has a first order (mean) hazard effect, because it affects the expected value of the magnitude density function $f_M(m)$ at each m. The other two epistemic correlations only affect the epistemic standard deviation of the seismic hazard. Thus, it is most important to recognize, model, and propagate the epistemic correlations between ν and β within a seismic area.

If (ν, β) are derived from estimates of (a, b) in the Gutenburg-Richter relation based on historical data, the sampling correlation between (a, b) may provide a basis for epistemic correlation between (ν, β). If other sources of epistemic uncertainty are also considered, these also must be included when determining the epistemic uncertainties in (ν, β). If (ν, β) are derived from elicited values of (a, b) or of the expected frequencies at two (or more) magnitudes, potential epistemic correlations in the elicited parameters must be recognized and accounted for in the seismic hazard calculations. One's ability to describe the epistemic correlation should be considered in selecting the inputs. For example, when using historically based estimates of (a, b) careful redefinition of a to a reference magnitude larger than zero can eliminate the sample correlation between a and b. Similar considerations are also advisable if the (ν, β)s are derived from inputs elicited in terms of other parameters, e.g.,

expected frequencies at selected magnitudes. Several approaches to quantifying correlation have been attempted. No one recommended way of modeling this has been identified.

With regard to epistemic correlation, when modeling the epistemic uncertainty to estimate the median of the ground motion distributions, assigning weights to classes of ground motion models to represent epistemic uncertainty will inherently imply certain epistemic correlations between the median ground motions at multiple (m, r)s depending on the relationships of the values of the several models at the different (m, r)s. If a single ground motion model with uncertainty or the FTI approach of developing uncertainty distributions at a finite number of (m, r)s is used to represent uncertainty, a common practice is to use "ground motion models" based on fitting fractiles of the uncertainty distributions at multiple (m, r)s. It should be recognized that this procedure implies perfect epistemic correlation in the medians between the (m, r)s. Assuming perfect correlation does not affect the mean hazard but will inflate the epistemic variance of the seismic hazard estimates, thus producing reduced median estimates and inflated estimates of higher, e.g., 85th, fractiles.

With regard the epistemic correlation in seismicity parameters between seismic areas, this is an issue primarily when the seismicity is variable between seismic areas within a source. Introducing smoothing of the seismicity parameters in the data analysis induces correlation of the seismicity parameters between seismic areas within a source. This needs to be recognized and properly combined with the epistemic uncertainties in the parameters to assure that the full range of potential values of the epistemically uncertain parameters is represented in the uncertainty analysis. (See Appendix I.)

Quantification of epistemic correlation can be difficult but it should be considered in the propagation of the epistemic uncertainty. Given the difficulty in quantifying correlation, it is recommended that the correlation not be represented by a single value of a correlation coefficient. Rather, it is recommended that the sources of the correlation be investigated, e.g., for ground motion median estimates, investigate magnitude scaling and different classes of ground motion models, etc. as sources of epistemic uncertainties. This information should be used in the context of the sampling procedure to assure that the full ranges of potential values of the epistemic uncertain parameters are represented in the sample.

7. GUIDANCE ON DOCUMENTING THE PSHA PROCESS AND RESULTS

7.1 Introduction

One of the major lessons that the PSHA community has learned in recent years is that documenting the PSHA process and results is as important as carrying out the project in a technically competent way. There are at least three reasons why excellent documentation is crucial:

- Only through adequate documentation can others in the technical community understand or review the analysis and the results.
- Only through adequate documentation can a later analysis team with new information or improved models utilize a PSHA to update it, revise it, or validate that it does not need an update or revision.
- Only through adequate documentation can the sponsoring organization retain an adequate record of the process it supported.

With these issues in mind, the SSHAC has developed guidance for PSHA documentation. A thorough documentation effort is required; however, the specific guidance herein cannot be considered as being required in detail because the specific manner in which a given analysis is documented depends on the objectives of the study, and both the details of the site hazard and the factors that affect it.

Note that much of the guidance below is given using the word must, which is intended to convey that the SSHAC strongly feels this particular aspect is crucial. At the same time, it is recognized that methods of presenting data and results change. With this in mind, guidance with respect to format is recommended, leaving the flexibility for modification to meet specific project needs or permit improved methods of presentation. Other documentation guidance is given using the word should, to convey the intent, although slightly weaker than the first category that SSHAC strongly recommends.

7.2 Process Aspects

The PSHA process is a multi-disciplinary evaluation that requires comprehensive documentation. This chapter provides the PSHA analyst with guidance on:

1 Levels of documentation that must be provided or maintained.

2 Elements of the PSHA process that must be documented.

3 Variations in documentation requirements as they pertain to the applications for which the PSHA was performed.

The following subsections describe these features of the PSHA process documentation.

7.2.1 Documentation—Two Tiers

Documentation of the PSHA should be prepared using a two-tiered approach:

- Tier 1 - consists of the documentation that must be reported, either in the main report or in appendices that are published with the main report, and widely accessible.
- Tier 2 - consists of the much larger body of background material that comprises the analysis documentation. This second-tier material should be maintained by the analysis team in an appropriately accessible, usable, and (if appropriate) auditable form. Of course, readily available documentation or references can be cited were appropriate.

Tier 1 and Tier 2 documentation is provided for each element of the PSHA process. These elements of the process are described in the next subsection.

It is strongly recommended that the authors of the PSHA use the two-tiered approach. If however, an alternative format is adopted, the documentation guidance described here must be satisfied.

7.2.2 Documenting the Elements of the PSHA Process

The following list shows the various elements of the PSHA process for which documentation is required. A summary of each is provided in the subsequent paragraphs.

- Roles and Resposibilities of the Project Participants and Consultants
- Comparisons With Other PSHA Studies
- Internal Quality Control and Review
- PSHA Methodology
- PSHA Results
- External Peer Review
- Documenting Citations

Documenting the Roles and Responsibilities of the Participants and Consultants

Even the simplest PSHA will involve a number of participants, and often a number of consultants. The Tier 1 documentation must identify the participants and provide a thorough discussion of the roles of each with care to differentiate the central roles from the supporting roles. Of particular importance is documenting the names of the author or authors who are professionally responsible for the overall performance of the study, and whose reputations support the findings.

Comparison With Other PSHA Studies

For many parts of the PSHA study, a very useful exercise is the comparison of the methods, data, or results with those from other PSHAs that have examined identical or similar geographical areas. If comparable PSHAs have been widely distributed and extensively reviewed, such comparisons can be valuable in demonstrating how different approaches or different data affect the conclusions. While such comparisons are very helpful, they are not required. However, where feasible the documentation of such comparisons should be done in a way that allows review, especially by the analysts who performed the earlier work or who provided its supporting data or models.

Documenting Internal Quality Control and Review

As part of the PSHA's internal quality-control procedures, it is necessary that there be a review of the ongoing work within the project. The process of such review, including the reviewers and their important findings must be documented in Tier 1 to assist both the sponsors and other users of the results in understanding what internal reviews were performed and what was found. The Tier 2 documentation should retain the detailed records of these reviews, as appropriate.

Documenting the PSHA Methodology

As part of the PSHA, there are a number of methodological aspects of the process that must be described. This includes the choice of the stochastic model to describe earthquake occurrences, the magnitude-frequency model, the elicitation methodology, etc. The Tier 1 documentation of the PSHA must provide a comprehensive description of all phases of the methodology that were used. If new models or approaches are used that differ substantially from the recommendations provided in this document (e.g., Chapters 2, 3, 4, 5, 6), a complete description and supporting basis for the alternative approach must be provided in the Tier 2 documentation. In addition, the Tier 2 documentation should describe the implementation of the methodology, such as identifying/describing the software that was used to compute the hazard, the elicitation processes that were conducted, etc. Section 7.4 discusses in more detail the documentation of the PSHA methodology and process.

Documenting PSHA Results

The results of a PSHA are typically presented in terms of fractile seismic hazard curves that define the probability that levels of ground motion may be exceeded at a site. These hazard curves are the composite aggregation of the epistemic uncertainties in the hazard evaluation. In addition, each hazard curve is computed in a composite aggregation of the aleatory uncertainties modeled in the hazard assessment (e.g., earthquake occurrences, ground motion). Recent experience, the requirements of engineering applications and regulatory processes require that comprehensive

documentation of the hazard results be provided. In addition to providing a broad characterization of the hazard, a comprehensive documentation makes the assessment tractable and transparent. Final results must be provided in the Tier 1 documentation with input to the PSHA and intermediate results and evaluations retained as part of the Tier 2 records.

Documenting External Peer Review

If an external peer review has been undertaken, both the principal review findings and the names of the reviewers must be documented in Tier 1. The details of the peer review should be retained as appropriate in the Tier 2 records.

Documenting Citations

It is important to provide proper citation for all data, methods, etc., that are utilized in the PSHA. To avoid this potential confusion or ambiguity, especially where primary earth-sciences data are used that are not readily available or are published in obscure or poorly circulated journals, the documentation should carefully explain where to find the important citations that may be difficult to obtain. Reliance on unpublished data is, of course, acceptable but such data should be considered part of the project files, to be either documented if necessary in Tier 1 (including Tier 1 appendices) or saved as Tier 2 but in an appropriately accessible and usable form.

7.3 Overview: Objective of the Documentation Process

To satisfy the range of PSHA applications (see Chapter 1), guidance for the presentation of results is provided. The objective is to satisfy:

- needs of those involved in the use of the PSHA results (e.g., provide results that satisfy the applications for which the analysis was performed), and
- requirements that the PSHA be tractable and transparent to the general practitioner, analyst, and technical reviewer.

The requirement to make the PSHA tractable is a critical part of the documentation process and, as experience would suggest, one that can be difficult. With these objectives in mind, this chapter provides guidance for the documentation of PSHA results, including the presentation of sensitivity analyses. Guidance is given so that the documentation:

1. Provides results that are required by the application for which the hazard assessment was performed (e.g., for use in developing a seismic zonation map, input to a seismic probabilistic risk assessment).
2. Provides information that permits the analyst or reviewer to understand the constituent parts of the analysis that dominate the seismic hazard (e.g., dominant seismic source, ground motion attenuation model).
3. Demonstrates the sensitivity/insensitivity of hazard results to the uncertainty in key parameters, and variation in the hazard due to the changes in parameter values considered in the hazard assessment.
4. Includes computer-readable (friendly) data files that facilitate the ability to examine specific parameter assessments or scientific interpretations. These data files would contain information that provides the analyst with the opportunity to understand the sensitivity of the results to specific parameters without the added of effort of recomputing the hazard.

The assessment of the seismic hazard at a site entails extensive computations to generate many thousands to tens of thousands of hazard curves, each corresponding to a specified set of parameters (e.g., seismic source geometries, maximum magnitude values, ground motion model). The role of any one or small group of hazard curves (and therefore the parameters that are their basis) is often difficult to determine.

7.4 Documenting the PSHA Process Methodology, Models, and Data Used

7.4.1 Introduction

The basic methodology for performing a PSHA is discussed earlier in this report. Here we provide

guidance to document the methodology, models, and data used.

The analysis team must document in Tier 1 how the overall PSHA has been structured, including the interrelationships among its several parts. The way in which "results" of one part are coupled to subsequent analyses must be discussed in sufficient detail to allow for review of the logic models, the data, completeness, approximations, and any assumptions. All critical aspects, such as the rationale for the binning of certain types of information, the melding together of different models or data, and the structure of the sensitivity and uncertainty analyses, must be documented in detail in Tier 1 in a form that allows for technical review.

This requirement includes documenting the following aspects (the following list can be considered a check list, but of course it is not nor could it be all-inclusive):

1. The basic formulation of the seismic hazard analysis—specifically, the mathematical formulation that describes how the bottom-line results" are derived from the inputs. This usually takes the form of one or more equations like those cited in Chapters 2 to 6, relating such quantities as the probability of exceedance of certain ground motion quantities to other more basic inputs or derived quantities in the analysis.

2. The definitions of all mathematically defined inputs, process variables, and output "results." This should include both mathematical definitions and word definitions for all quantities.

3. Where such limits exist, definitions of the limits of validity of any mathematical formulas or equations used.

4. Careful definitions of the physical units of all quantities (preferably in SI units, but if other units are used, then an explanation is needed of the relationship to SI units).

5. Careful definition and explanation of any mathematical models used, including their ranges of validity, the approximations used in their derivation or use, and any comparisons with other models, similar or different, that are used elsewhere to describe the same or a similar technical issue. Especially helpful here are discussions of any previous uses of the chosen model, including comparisons to other applications between the model and data, direct observations, or inferred analytical results.

6. A discussion of the sources of all experimental data and field observations, including the methods used to obtain these data and observations, the uncertainties (both aleatory and epistemic) as reported by the experimenters/field observers, and any interpretive discussion necessary to understand the ranges of validity of the data and observations.

7.4.2 Documenting the Methods, Models, and Data Used For Seismic Source Characterization

Tier 1

The PSHA report must provide a complete description of the SSC methodology and its implementation to develop the SSC. The Tier 1 documentation must describe:

1. The steps taken to gather the geologic, seismologic and geophysical data used in the PSHA

2. Data resources available to the earth science experts

3. Methodology and steps taken to elicit expert input, including a description of what was elicited from the experts

4. Methods used to define seismic source geometries, magnitude recurrence relationship and maximum magnitude

The SSC is a critical part of the PSHA. It is a multi-disciplinary effort that requires an integrated scientific evaluation and interpretation of a wide range of earth-sciences data. Documentation of the SSC must provide a comprehensive presentation of the process that

was used to develop a model of the seismicity in the vicinity of a site.

Tier 2

In the Tier 2 documentation, a complete cataloging of the data used by the earth-science experts should be retained. Supporting documents, calculation packages, etc. generated by the earth-science experts, PSHA analysts, TFI, and others involved in SSC process should be catalogued and retained.

7.4.3 Documenting the Methods, Models, and Data Used for Analyzing Ground Motion Attenuation

Tier 1

The PSHA report must provide a complete description of the methodology to determine the ground motion models used in the PSHA and the steps taken to implement it. The Tier 1 documentation should describe:

1. The steps taken to gather strong motion data, attenuation models and geophysical data considered in the analysis.
2. Data resources available to the ground motion experts.
3. Methodology and steps taken to elicit expert input, including a description of what was elicited from the experts.
4. Methods used to define the ground motion models.

Tier 2

In the Tier 2 documentation a complete cataloging of the data used by the ground motion experts should be retained, including models and data that were considered. As noted above, it is important that information used in selecting ground motion models be adequately and completely cited and retained. Similarly, supporting documents, calculation packages, etc. that were generated by the ground motion experts, PSHA analyst, TFI and others involved in the analysis or the selection of ground motion models should be catalogued and retained.

7.4.4 Documenting the Methods Used to Produce the PSHA Results

Tier 1

The Tier 1 report must provide a complete description of the mathematical model used to determine the seismic hazard at a site, the approach used to propagate the epistemic uncertainties and the method used to aggregate expert interpretations. The description of the seismic hazard model should fully describe the method used to compute deaggregated hazard results, including the identification of magnitude distance bins, the calculation of the fraction contribution of seismic sources to the total hazard, and the contribution of parameters to the hazard epistemic variance.

Tier 2

For seismic hazard models that differ from the approach described in Chapters 2 and 6, a detailed description of the basis for the alternative model must be presented. In cases where a conventional approach is used, no Tier 2 documentation of the seismic hazard methodology is required.

A description of the software used to compute the seismic hazard should be provided. As a minimum, the software routines, the flow of information, and the software output must be described.

7.5 Documenting the Seismic Hazard Results—Scope

This section provides guidance for documenting the results of the PSHA, including the set of numeric results that quantify the hazard at a site(s), the seismic source characterizations developed to model the active tectonic features in a region (e.g., identification of active seismic sources, the estimate of earthquake rates and maximum magnitudes), and the ground motion attenuation models that are used (or possibly developed in site-specific studies). Guidance is provided for Tier 1 and Tier 2 documentation.

Documentation of the seismic hazard results for a site(s) is divided into three parts:

1. Basic PSHA Results
2. PSHA Deaggregation
3. Sensitivity Analysis

These parts define a hierarchy of results that proceed from basic user-required results to increasing levels of detail that provide insights into the dominant contributors to the site hazard and the sensitivity of the results to parameter variations.

Basic PSHA Results

These are the results that must be generated by the PSHA to satisfy the needs of the specific applications for which the study was performed. Examples of results that document the hazard at a site are the fractile and mean hazard curves for each ground motion measure and the uniform hazard response spectra (UHS). Table 7-1 presents a list and description of Basic PSHA Results.

PSHA Deaggregation

The assessment of the seismic hazard at a site is the result of an aggregation of the aleatory and epistemic uncertainties in the analysis. A deaggregated presentation of the seismic hazard provides a measure of the relative contribution of the constituent parts of the seismic hazard model to the total hazard and an indication of the sensitivity to different parameter assumptions. Table 7-2 provides a list and description of deaggregated seismic hazard results. Note, in some applications deaggregated hazard results may be required as input to certain applications (e.g., studies that require an estimate of the mean magnitude and distance).

Deaggregated seismic hazard results present the hazard in terms of a number of the basic building blocks of the analysis (e.g., the characterization of seismic sources, the prediction of ground motion). By themselves, these results provide a measure of the sensitivity of the hazard to specific inputs and the impact of potential changes to the hazard. For example, a deaggregation of the hazard in terms of seismic sources provides insight to the source(s) that dominate the site hazard. At the same time, this result also demonstrates the insensitivity of the hazard results to parameter variations for sources that make a small contribution. Consequently, deaggregation of seismic hazard results provides valuable insights to the PSHA and the inputs that contribute to the results.

Sensitivity Evaluations

At different stages of the PSHA, sensitivity evaluations are performed. For example, early in the study, sensitivity evaluations may be conducted to identify the dominant factors in the analysis to guide the collection of data, focus the SSC, etc. Similarly, at the conclusion of the study, sensitivity analyses are performed to demonstrate the role of different parameters or their contribution to the epistemic uncertainty. Due to the often complex relationship that may exist between parameters in the analysis (e.g., correlations), sensitivity analyses are a useful means to provide specific insights into their role in the analysis.

7.6 Documentation Guidance: Reporting the Seismic Hazard Results

7.6.1 General Guidance

Table 7-3 provides a summary of general guidance for documenting Basic PSHA Results. The table addresses required fractile levels, use of the mean hazard and presentation of results for rock and soil site conditions. Documentation of the seismic hazard results must include graphic and tabular presentation. In addition, all graphic displays of like results should be provided to a consistent scale for comparison.

Table 7-1 List and Description of Standard PSHA Results

Standard PSHA Result	Required or Optional	Description	Format
Fractile and Mean Hazard Curves	Required	For each ground motion parameter considered in the PSHA, the seismic hazard is expressed in terms of fractile and mean hazard curves. A fractile hazard curve quantifies the probability, p, that the frequency of exceeding each ground motion level is not greater than the value defined by the hazard curve.	The following fractile hazard results are reported: 0.05, 0.15, 0.50, 0.85, and 0.95. In addition, the mean hazard curve is presented.
Uniform Hazard Response Spectra	Required	Response spectrum shapes corresponding to a specified probability of exceedance level. Fractile UHS, and the mean response spectra are determined from the corresponding hazard curves.	The UHS are presented in at least two formats: 1) fractile and mean UHS for a specified probability of exceedance, and 2) the mean or selected fractile level UHS for different probability levels displayed together.
Aggregated Hazard Curves	Optional	Hazard curves that have been generated from an analysis in which a large number of hazard curves have been aggregated to produce a smaller, more manageable set. the combination process should preserve the diversity in hazard curve shapes as well as essential properties of the original set of hazard curves (e.g., mean hazard). This format is used as input to seismic probabilistic risk assessments.	A group of discrete hazard curves is generated, each with a probability weight assigned to it. The hazard curve weights sum to one (see Fig. 7-8).
Ground Motion Contour Map	Optional	To display the hazard in a region, the results can be presented in terms of a ground motion contour map. The ground motion contours define ground shaking levels that have the same probability of exceedance in a specified time period.	A map is produced for each ground motion measure, time period and probability of exceedance (see Fig. 7-9).

Table 7-2 Deaggregated PSHA Results[1]

Deaggregated Result	Description
Magnitude and Distance	A magnitude-distance (M-D) deaggregation entails a presentation of the hazard for selected ranges of magnitude and distances. A M-D aggregation can be presented in terms of the hazard for each M-D pair or in terms of the relative contribution of each M-D pair to the total hazard.
Seismic Sources Individual Sources	Seismic hazard results are presented on a source-by-source basis. The epistemic uncertainty in the activity of the source is not considered. These results can be used to determine the contribution of individual seismic sources to the total hazard. The presentation of source specific results can follow the format for presenting the total hazard (see above).
Magnitude and Distance	For each seismic source, the hazard is deaggregated in terms of magnitude and distance in the same manner as the total hazard (see above). The M-D deaggregation by seismic source provides a breakdown of the difference size earthquakes for the sources that contribute to the site hazard (e.g., importance of the estimate of maximum magnitude).
Ground Motion Attenuation Model	The hazard results are readily deaggregated with respect to the ground motion models that are used in the PSHA. For each attenuation model fractile and mean, hazard results can be presented in a format similar to that used for the total hazard. Each set of fractile results is conditional on the attenuation model considered.
SSC and Ground Motion Experts	If multiple SSC and Ground Motion Experts are used, fractile and mean hazard results can be displayed. The format is similar to the results presented for the total hazard. If a TFI approach is used, results can be presented in terms of particular alternatives or hypotheses that are selected.

Table 7-3 General Guidelines for Documenting PSHAs

PSHA Result/Parameter	Guidance
Fractile Hazard Curves	Results should be presented for the 0.05, 0.15, 0.50, 0.85, and 0.95 fractile levels.
Mean Hazard	The mean seismic hazard curve should be reported for all Basic PSHA Results.
Soil Sites	Basic PSHA results should be presented for both rock and soil site conditions.
Deaggregated Hazard Results	For estimating the relative contribution of a parameter (e.g., seismic source, ground motion model) to the hazard, this should be done using the mean.

[1]Note: Depending on the methods recommended by the SSHAC, there may be other types of deaggregation that could be considered.

7.6.2 Basic PSHA Results

Tier 1 Documentation

Table 7-1 listed the Basic PSHA Results that should be reported. As noted in the table, certain of the results are "Required." whereas others are "Optional." The analyst should note the distinction. The "Required" results are those that must be presented for all applications, whereas results that are listed as "Optional" are those that must be provided for the purpose of satisfying the application for which the PSHA was conducted (see Table 7-1).

For example, most PSHAs are performed for a single site; therefore, a contour map of ground motion is not computed/required. However, in a regional study in which a contour map must be produced, fractile hazard results should also be reported for selected sites.

Fractile and Mean Hazard Curves

The seismic hazard at a site is presented in terms of the annual probability of exceedance of selected ground motion parameters such as PGA (peak ground acceleration), S_a (absolute spectral acceleration), or PSV (pseudo spectral velocity). The latter two are presented as a function of frequency. In some cases, results for displacement are also presented.

Figures 7-1 and 7-2 show the fractile and mean hazard curves for PGA and PSV (1 Hz) for a rock outcrop. Figures 7-2 and 7-3 show the same information for the same site but for soil site conditions (the site in this case is a deep soil site). This information is also presented in tabular form as shown in Table 7-4 for peak ground acceleration. Graphic and tabular results similar to Figures 7-1 to 7-4 and Table 7-4 must be presented for all the ground motion parameters considered in the PSHA.

Note that in presenting the final results, it is necessary to specify certain key parameters, including whether the hazard is for a rock outcrop or for a soil site, and the damping value used when the hazard is presented in terms of response spectrum ordinates.

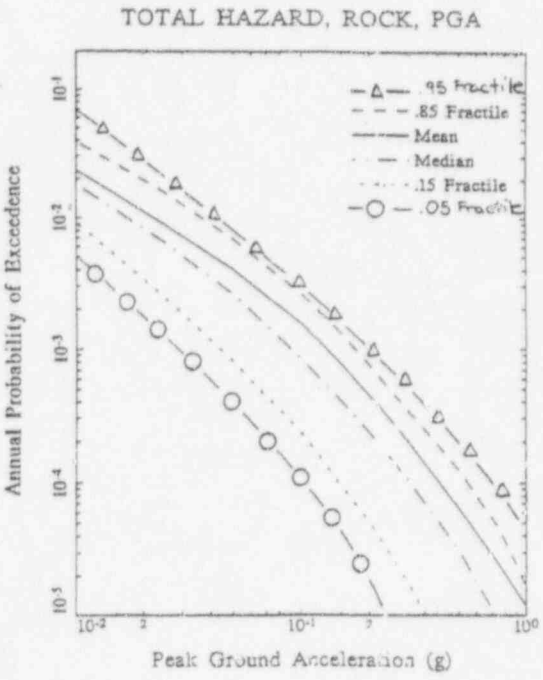

Figure 7-1 Total seismic hazard curves for PGA and rock site conditions.

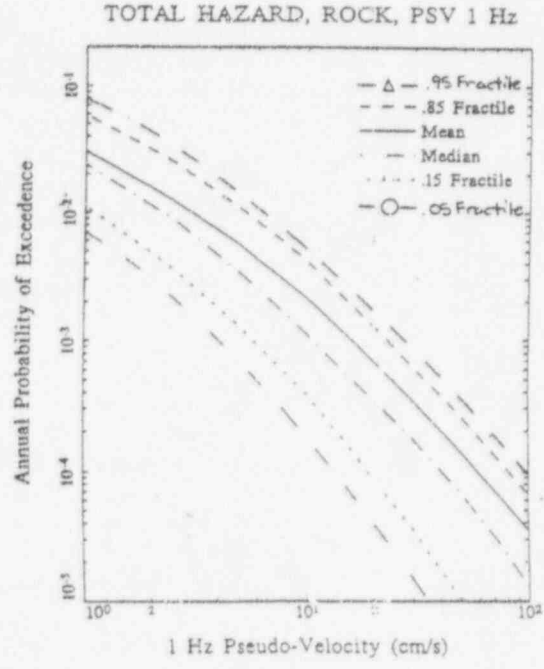

Figure 7-2 Total seismic hazard curves for PSV (1 Hz) and rock site conditions.

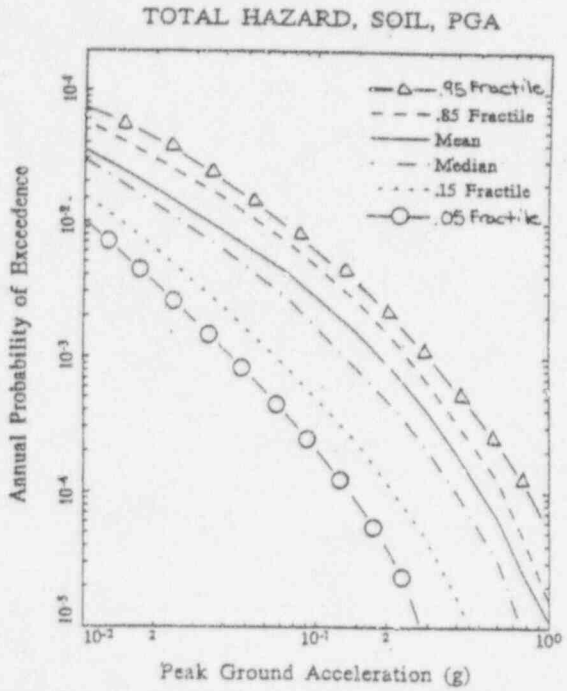

Figure 7-3 Total seismic hazard curves for PGA and soil site conditions.

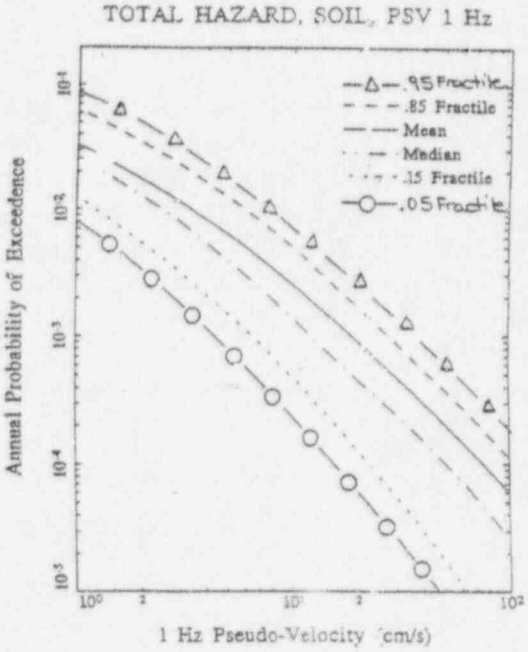

Figure 7-4 Total seismic hazard curves for PSV (1 Hz) and soil site conditions.

Table 7-4 Seismic Hazard Results for Peak Ground Acceleration—Rock (The fractile values were not available.)

Level	Mean	Fractile Levels				
		0.05	0.15	0.50	0.85	0.89
1	2.34E-2					
2.5	8.88E-3					
5	4.07E-3					
10	1.55E-3					
15	7.76E-4					
20	4.49E-4					
30	1.96E-4					
50	6.52E-5					
75	2.48E-5					
100	1.18E-5					

Uniform Hazard Spectra

If the hazard assessment is performed for multiple ground-motion parameters over a range of frequencies, uniform hazard response spectra must be reported. When presenting the UHS results in graphical and tabular format, the following information must be presented:

- annual probability of exceedance level
- frequencies and spectral values
- damping level
- fractile and mean response spectrum values

A number of alternative formats are available to graphically present UHS results. They are:

1. Fractile and mean UHS for a specified annual probability of exceedance and damping level.
2. UHS for a specified damping, multiple annual probability of exceedance levels using the mean or a specified fractile (e.g., 0.50 fractile).
3. UHS for a specified annual probability of exceedance, multiple damping levels using the mean or a specified fractile (e.g., 0.50 fractile) level.

UHS can be displayed in terms of acceleration or velocity. Log-log plots are often used, but other formats can be selected. In the PSHA report, a consistent format should be followed.

When determining the ground motion levels corresponding to a specified annual probability of exceedance, a log-log interpolation scheme is recommended. If extrapolations beyond the computed hazard curve must be made, this should be noted as part of the tabular and graphical presentation.

The UHS should be determined for annual probability levels that satisfy the application for which the study was performed. Typically, a number of probability levels are used, ranging from 10^{-2} to 10^{-5}.

Figure 7-5 shows an example of a UHS for a specified damping level and annual probability of exceedance. For the damping and probability levels shown, the 0.05, 0.15, 0.50, 0.85 and 0.95 fractiles and mean UHS are shown. Figure 7-6 shows the mean UHS for a specified annual probability of exceedance (10^{-3}) and different damping levels. Figure 7-7 shows the mean UHS for a single damping level and different annual probabilities of exceedance.

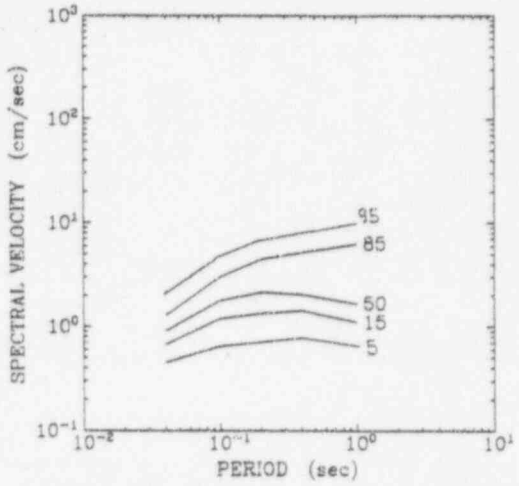

Figure 7-5 Uniform hazard response spectra for soil site conditions for an annual probability of exceedance of 10^{-3} and damping level 5% of critical.

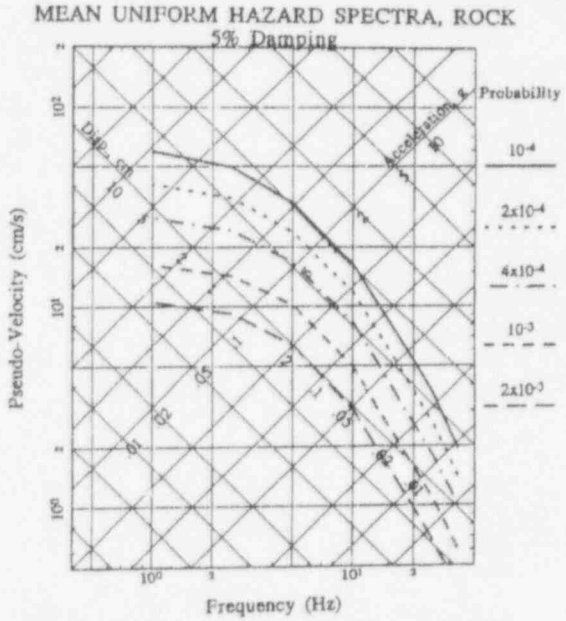

Figure 7-6 Mean uniform hazard response spectra for soil site conditions and annual probabilities of exceedance of $2x10^{-3}$, 10^{-3}, $4x10^{-4}$, $2x10^{-4}$, and 10^{-3} and a damping level 5% of critical.

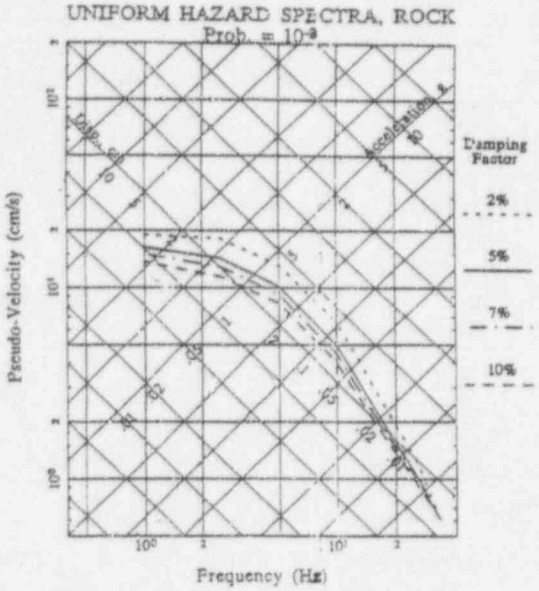

Figure 7-7 Mean uniform hazard response spectra for soil site conditions at four damping levels. Annual probability of exceedance = 10^{-3}.

Table 7-5 Tabulation of Uniform Hazard Response Spectra Results Rock Site Conditions

Frequency	Probability Level	Mean	Fractiles				
			0.05	0.15	0.50	0.85	0.95
1.0	$2x10^{-3}$	0.067					
	10^{-3}	0.104					
	$2x10^{-4}$	0.262					
	10^{-4}	0.379					
2.5	$2x10^{-3}$	0.147					
	10^{-3}	0.230					
	$2x10^{-4}$	0.558					
	10^{-4}	0.791					
5.0	$2x10^{-3}$	0.204					
	10^{-3}	0.326					
	$2x10^{-4}$	0.763					
	10^{-4}	1.06					
10.0	$2x10^{-3}$	0.197					
	10^{-3}	0.319					
	$2x10^{-4}$	1.00					
	10^{-4}	1.06					
25.0	$2x10^{-3}$	0.106					
	10^{-3}	0.171					
	$2x10^{-4}$	0.414					
	10^{-4}	0.582					
PGA	$2x10^{-3}$	0.083					
	10^{-3}	0.129					
	$2x10^{-4}$	0.297					
	10^{-4}	0.410					

For soil sites, UHS results for a rock outcrop and soil-site conditions should be presented (no such examples are provided here). The graphical presentation of the rock and soil UHS should be provided to the same scale.

Table 7-5 shows a typical tabular presentation of UHS results, which lists the mean UHS for different damping values (2%, 5%, 7%, 10%), and presents PSA (pseudo-acceleration) at five frequencies (1, 2.5, 5, 10, and 25 Hz). A table similar to Table 7-5 is presented to accompany each graphical UHS presentation.

When presenting these bottom-line hazard results, it is not usually possible to present "too much information." To the informed reader, the different formats, including both figures and tables, are of great benefit in understanding the results from different perspectives.

Aggregated Seismic Hazard Curves

For input to a probabilistic seismic risk assessment in which an uncertainty assessment is performed, a discrete family of seismic hazard curves must be provided by the PSHA. A discrete family of hazard curves is a group of individual hazard curves with weights assigned to each curve that quantify the epistemic uncertainty in the hazard. The weights sum to one for the group of hazard curves. As noted in Chapter 6, many thousands to many tens-of-thousands of discrete hazard curves may be computed in the hazard assessment. For purposes of performing a seismic risk assessment, it is possible to aggregate this large set of hazard curves into a much smaller, more manageable set for use in a risk assessment.

Cluster analysis techniques (Veneziano, Cornell, and O'Hara 1984) are available to combine hazard curves while retaining the basic probabilistic properties of the original set (i.e., mean and variance). Figure 7-8 presents a schematic illustration of a set of aggregated hazard curves. In most applications, an aggregate set of 10 to 20 hazard curves is adequate for input to a seismic risk assessment. For completeness, the set of aggregated hazard values should also be tabulated. The tabulation lists the hazard curves and their respective weights.

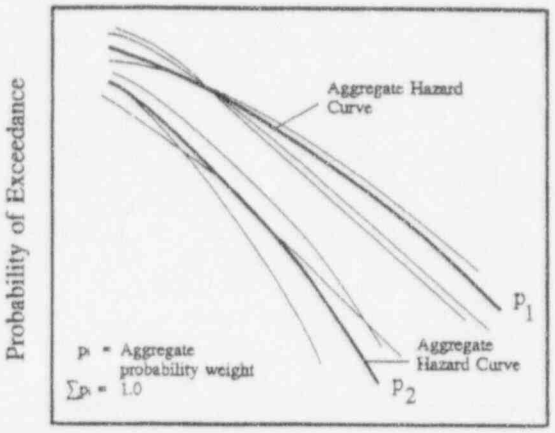

Figure 7-8 Illustration of an aggregate set of seismic hazard curves for PGA for soil site conditions. Each curve has an assigned weight to it.

Ground Motion Contour Maps

For certain applications, contour maps are generated that define the ground motion levels that have the same probability of exceedance for a specified period of time. One such application is in the development of a national seismic design code. Figure 7-9 shows an example of such a map. In defining a ground-motion contour map, the analyst must decide on the following:

- future time period(s) to be considered
- probability of exceedance levels
- ground motion parameters to be mapped (including damping levels)
- fractile levels or mean hazard to be mapped

In addition, the analyst must select a group of sites in a region in order to provide adequate spatial sampling of the ground motion.

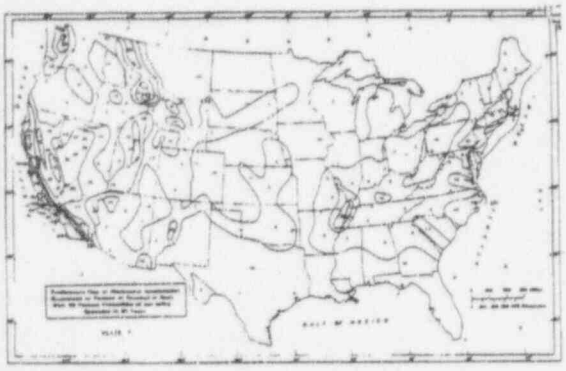

Figure 7-9 Illustration of a ground motion contour map corresponding to specified time period and probability of exceedance level.

Tier 2 Documentation

The Tier 2 documentation of the Basic PSHA Results is contained in the seismic hazard information base maintained in computer-readable data files. The information contained in these files is described in Section 7.12.

7.7 Documenting Deaggregated PSHA Results

7.7.1 Overview

The seismic hazard at a site is an aggregation of:

- the hazard associated with individual seismic sources, and
- earthquakes of different magnitude that occur over a range of distances form the site

In addition, the quantification of the epistemic uncertainty in the hazard is an aggregation of the epistemic uncertainty in individual parameters in the analysis. The form of aggregation is different in each of these examples. In the former case, aggregation is carried out to account for the randomness of earthquake occurrences by location (seismic source to seismic source, and within a seismic source) and earthquake magnitude. In the later case, a very different form of aggregation is performed. Here the epistemic uncertainty in parameter assessments is propagated through the analysis to determine the total epistemic uncertainty in the seismic hazard. To facilitate the understanding of the PSHA, deaggregated results must be presented.

7.7.2 General Guidance

For all PSHA studies, deaggregation results for seismic sources and magnitude and distance must be provided. In addition, if the seismic hazard is estimated for a range of ground-motion parameters, the deaggregation should be performed for at least two spectral frequencies. It is recommended that the deaggregation be performed, as a minimum, for two ground-motion parameters (PSV or S_a) at 1 and 10 Hz. The analyst may consider other frequencies based on the application for which the results will be used. For example, if a hazard assessment is being performed for a long-period structure such as a suspension bridge, results at longer periods (e.g., lower frequencies) will be of interest.

The deaggregation should be performed using mean seismic hazard results. (Note, some applications may require the use of the median hazard.)

7.7.3 Presenting Deaggregated Results

This section describes the deaggregated results that must be provided.

Tier 1 Documentation

Source Deaggregation —The seismic hazard at a site is attributed to the likelihood of ground motions that are generated by multiple sources of seismic activity. An informative PSHA result is to display the mean hazard curve for each seismic source. Figures 7-10 and 7-11 show such a presentation for PGA and PSV at 1 Hz. This is the same site and analysis as for Figures 7-1 and 7-2, except here the mean seismic hazard associated with nine seismic sources is presented. When this information is compared to that in Figures 7-1 and 7-2, one can assimilate clearly the fact that one seismic source dominates the mid-range exceedance probabilities, but a second source dominates at the high end while a third source dominates at the low end of the plot. Table 7-6 presents the results for PGA in tabular form.

A clear picture of the relative contribution of each seismic source to the total mean hazard can be displayed as shown in Figure 7-12. The figure displays, for a selected ground motion parameter, the relative contribution of each seismic source to the total mean hazard at specified ground motion levels. Similar plots can be displayed for each ground-motion parameter. These results are tabulated in Table 7-7.

Magnitude-Distance Deaggregation—Hazard results deaggregated in terms of magnitude and distance must be presented for the total hazard. In addition, magnitude-distance deaggregation should also be presented for the seismic sources that dominate the site hazard. The deaggregation is presented for selected M and D bins that cover the range of magnitudes and distances considered in the analysis and provide adequate representation of the M-D density. The following magnitude and distance have been useful in PSHA applications.

Parameter	Bin
Magnitude	5-5.5, 5.5-6, 6-6.5, 6.5-7, 7-7.5, 7.5-8, >8
Distance (km)	0-10, 10-25, 25-50, 50-100, 100-150, 150-200, >200

Table 7-6 Mean Seismic Hazard Curves for Each Seismic Source Peak Ground Acceleration and Rock Site Conditions

Level (g)	Total	Seismic Source								
		1	2	3	4	5	6	7	8	9

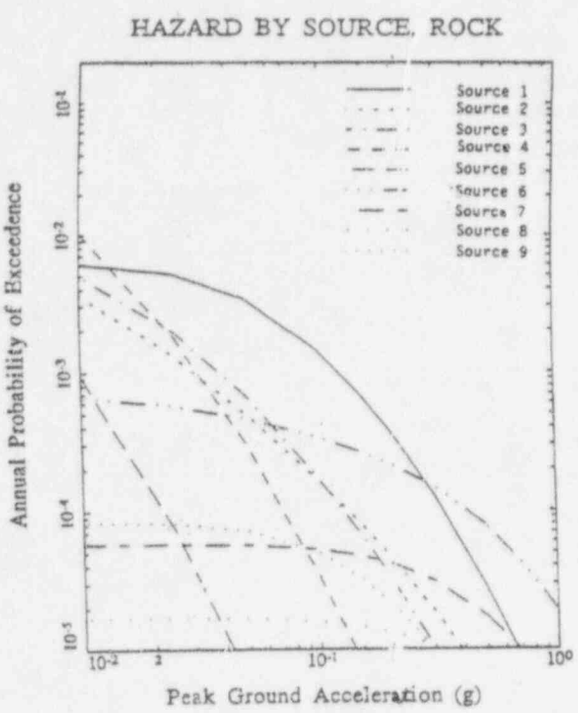

Figure 7-10 Mean seismic hazard curves for individual seismic sources. Results are for PGA and rock site conditions.

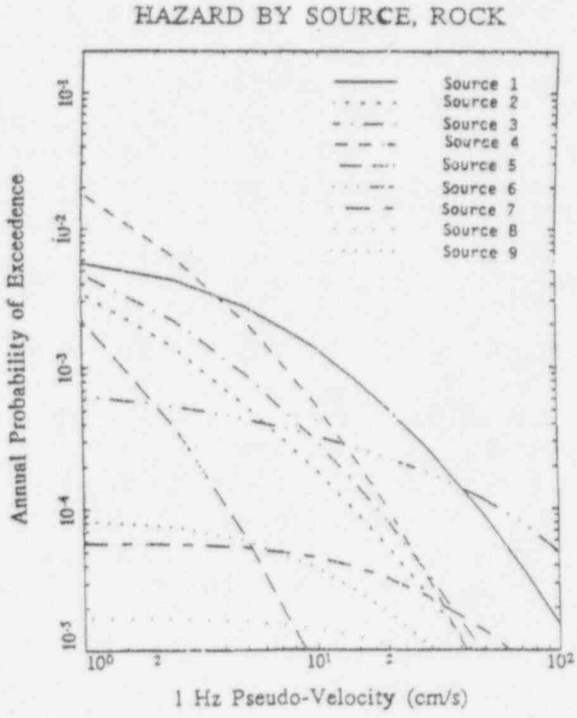

Figure 7-11 Mean seismic hazard curves for individual seismic sources. Results are for PSV(1 Hz) and rock site conditions.

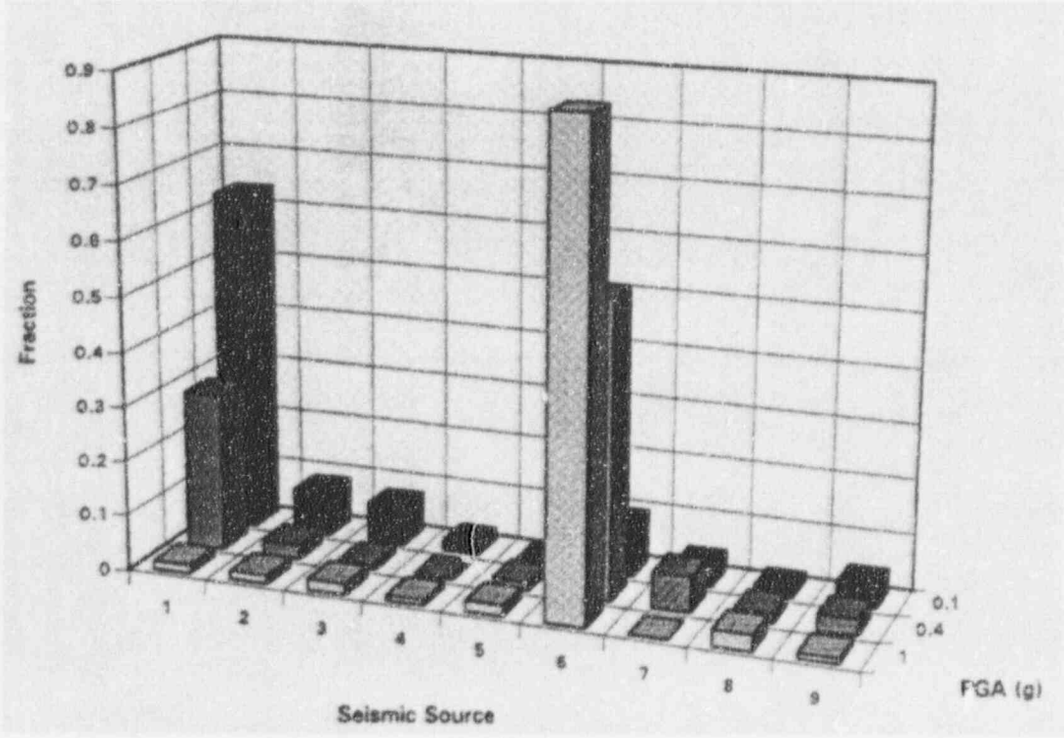

Figure 7-12 Illustration of the relative contribution of individual seismic sources to the total mean hazard. The results are presented for PGA and rock site conditions and three ground motion levels.

Table 7-7 Relative Contribution of Each Seismic Source to the Total Mean Hazard for Peak Ground Acceleration and Rock Site Conditions

Seismic Source	PGA (g)		
	0.10	**0.40**	**1.0**
1	0.63	0.29	0.00
2	0.07	0.00	0.00
3	0.07	0.00	0.00
4	0.02	0.00	0.00
5	0.00	0.00	0.00
6	0.16	0.29	1.00
7	0.03	0.12	0.00
8	0.01	0.00	0.00
9	0.03	0.00	0.00

The PSHA analyst should select the M-D bins appropriate for a given application.

Figure 7-13 shows an example of the M-D deaggregation for the total seismic hazard for PGA at a selected ground motion level. In this case, it is apparent that the hazard is dominated by seismic events with magnitude less than 6.0 that occur within 50 km of the site.

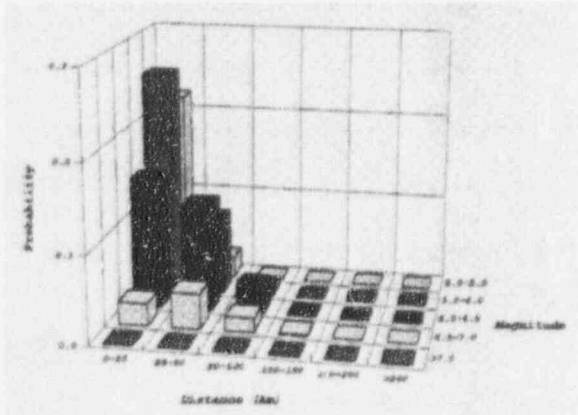

Figure 7-13 Magnitude-distance deaggregation for the total hazard. Results are provided for PGA at 0.25g and rock site conditions.

Table 7-8 tabulates the results presented in Figure 7-13. Similar results should be presented for the seismic sources that dominate the hazard.

Tier 2 Documentation

The Tier 2 documentation of the Deaggregated PSHA Results is contained in the seismic hazard information base that is maintained in computer-readable data files. The information contained in these files is described in Section 7.12.

7.8 Documenting PSHA Sensitivity Analyses

As part of a PSHA, sensitivity analyses are performed to move the study forward and to demonstrate the role of parameters in the analysis. For example, during the course of the study, sensitivity calculations may be performed to provide insights to the factors that will be important/unimportant to the assessment of the hazard at a site. The results of these calculations often become a guide for the PSHA.

At the conclusion of the PSHA, the results of sensitivity analyses provide a means to demonstrate the role that the variation in

individual parameters has on the results. For example, sensitivity calculations can be used to show the variation in the mean hazard due to the different estimates of the maximum magnitude for the seismic source that dominates the hazard. Sensitivity calculations can also be used to determine the contribution of different parameters in the analysis to the total epistemic uncertainty (e.g., total variance).

A number of alternative methods are available to perform sensitivity evaluations. Experimental design and response surface techniques (to name just a few) are examples of sophisticated methods that can be used. In PSHA, relatively simple methods can be used to show the sensitivity of the hazard to parameter variations. As noted above, a deaggregation of the hazard on the basis of seismic sources, magnitude and distance and ground motion attenuation models provides valuable insights to the factors that do/do not contribute to the hazard.

As part of the PSHA documentation, the results of sensitivity calculations should:

1. provide insights to the site hazard that guided the scope and depth of the analysis that was performed
2. demonstrate the sensitivity of the hazard results to the variation in critical parameters in the PSHA (e.g., parameters for the dominant seismic source)

Two types of sensitivity evaluations are recommended. In the first approach, a base case is assumed and the parameter of interest is varied. The results are displayed to demonstrate the variation in the hazard. As an example, consider a site where the hazard is, dominanted by a single seismic source. For this source, three estimates of the maximum magnitude (m_{max}) are defined by the experts. A series of hazard calculations are performed in which the maximum magnitude is set to each of the three m_{max} values.

Table 7-8 Magnitude-Distance Contribution to the Total PGA Hazard at 0.25g—Rock Site Conditions

Distance (km)	Magnitude				
	5.0-5.5	5.5-6.0	6.0-6.5	6.5-7.0	>7.0
0-25	.011	0.26	0.15	0.03	0
25-50	0.02	0.10	0.11	0.05	0.01
50-100	0	0.02	0.06	0.04	0.01
100-150	0	0	0	0.02	0
150-200	0	0	0	0.01	0
>200	0	0	0	0	0

Table 7-9 Contributions to the Epistemic Variance in the Hazard

S_a (g)	Attenuation Model	Hayward Tran.	Recurrence Model	Fault Segmentation	Total Length	Magnitude Distribution	San Andreas Depth	M_{max}	L.P. Assc.	Activity Rate	b-Value
Peak Acceleration											
0.050	0.255	0.008	0.306	0.010	0.018	0.087	0.001	0.106	0.001	0.155	0.054
0.100	0.249	0.005	0.394	0.005	0.014	0.050	0.000	0.077	0.001	0.142	0.064
0.200	0.232	0.001	0.488	0.008	0.013	0.014	0.000	0.066	0.002	0.110	0.066
0.300	0.315	0.000	0.428	0.014	0.013	0.006	0.000	0.060	0.002	0.104	0.057
0.400	0.427	0.000	0.314	0.020	0.013	0.004	0.000	0.057	0.001	0.113	0.051
0.500	0.509	0.000	0.205	0.027	0.011	0.004	0.000	0.057	0.001	0.135	0.052
0.600	0.538	0.000	0.117	0.036	0.009	0.003	0.000	0.063	0.000	0.170	0.062
0.650	0.531	0.000	0.083	0.043	0.008	0.003	0.000	0.069	0.000	0.192	0.070
0.700	0.511	0.000	0.056	0.051	0.006	0.003	0.000	0.077	0.000	0.217	0.080
0.800	0.436	0.000	0.021	0.069	0.004	0.002	0.000	0.098	0.000	0.268	0.103
1.000	0.255	0.000	0.004	0.104	0.001	0.001	0.000	0.145	0.000	0.344	0.144
1.250	0.183	0.000	0.004	0.117	0.000	0.001	0.000	0.183	0.000	0.352	0.161

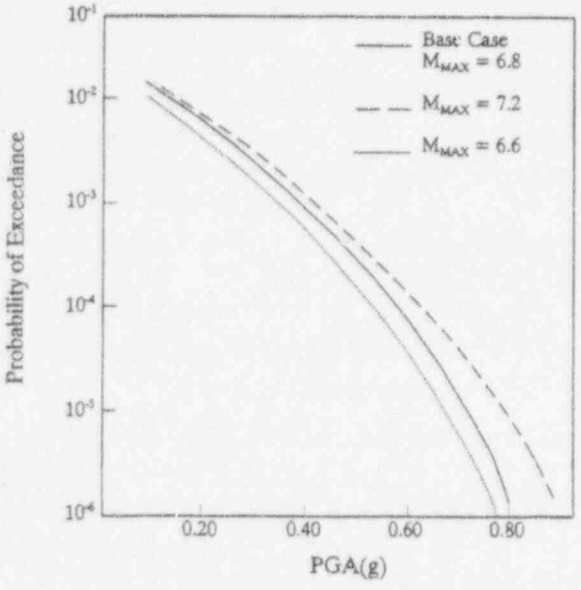

Figure 7-14 Illustration of the sensitivity of hazard results to the variation in maximum magnitude.

Figure 7-14 shows an example of this type of result. A comparison of the mean hazard curves for the three sets of calculations demonstrates the variation in the results based on each of the three maximum magnitude values. Similar sets of hazard calculations can be used to demonstrate the sensitivity of the PSHA results to other parameters. The advantage of this type of sensitivity evaluation is that it provides the analyst, reviewer or user of the PSHA with an understanding of the specific variation in the hazard that is attributed to the change in a particular parameter. At the same time however, these types of sensitivity results may not provide a complete picture of the relative contribution of one or multiple parameters to the hazard.

An alternative format for presenting the sensitivity of the hazard to different parameters is to compute the contribution of the epistemic uncertainty in individual parameters to the total epistemic uncertainty. Table 7-9 presents an example of this type result. This type of sensitivity analysis corresponds to an analysis of the epistemic variance.

Because there are so many different types of sensitivity analyses that could be performed, it is not feasible to present examples here of how to display the results from them. The following is a partial list of the types of analyses for which it may be desirable to present graphical or tabular information:

- sensitivity of hazard to M_{max}
- sensitivity of hazard to various b-values
- sensitivity of hazard to various a-values
- sensitivity of hazard to attenuation parameters, of which there are various that one could display
- sensitivity of hazard to geometry of the dominant fault(s) or other seismic sources
- sensitivity of hazard to slip rate for the dominant fault(s)
- sensitivity to M_L conversion
- sensitivity to local site conditions, such as soil-amplification factors

Tier 2 Documentation

For the sensitivity calculations that are reported in the Tier 1 documentation, backup computer data files similar to those that document the Basic PSHA Results should be provided in the Tier 2 documentation.

7.9 Comparisons with Other Studies

As part of the PSHA documentation, it is useful to the study sponsor to provide a comparison with the results of prior studies. Often, questions about how the current study results compare to previous analyses will arise. It is therefore beneficial to provide in the Tier 1 documentation the results of a comparative evaluation.

If detailed comparative assessments are performed, the supporting documentation should be retained in Tier 2 records.

7.10 Documentation Guidance: Results of the Seismic Source Characterization

The basic purpose of the SSC is to estimate the rate of future earthquake occurrences in the vicinity of a site. This rate has a spatial as well as a temporal component. The spatial variation of earthquake occurrences is modeled through the determination of the temporal component of earthquake occurrence rates within individual seismic sources. The earth science expert's model of earthquake occurrences near a site is defined by the map of seismic sources that are defined and the individual source earthquake recurrence models. Combined, the source map and source recurrence models fully define the spatial and temporal rate of earthquake occurrences in the vicinity of a site.

Tier 1

In the Tier 1 documentation of the SSC, the following must be presented:

1. Seismic source maps that present the earth science expert model for a region. Multiple maps or figures may be required to provide adequate detail legible to the reader. For example, alternative maps and figures may be required to present faults that are modeled as three-dimensional structures in the hazard analysis. Figure 7-15 shows an example for an expert seismic source map. If multiple experts provide input to the PSHA, the seismic source maps for each must be presented.

2. For the seismic sources that dominate the site hazard, the magnitude recurrence model, including the epistemic uncertainty, must be presented. The epistemic uncertainty must be presented in terms of the fractile recurrence curves at the 0.05, 0.15, 0.50, 0.85 and 0.95 fractile levels. The mean recurrence is also presented. Figure 7-16 shows this type of result.

3. The PSHA model for seismicity in the vicinity of the site is presented for the rate of earthquake occurrences at several (at least 3) magnitudes. Figure 7-17 shows an example of a map that displays the mean rate of earthquake occurrences above magnitude 5.

4. For the region around a site (approximately 200 km), a comparison of the historic rate of earthquake occurrences and the PSHA estimate must be presented. Figure 7-18 shows this type of comparison.

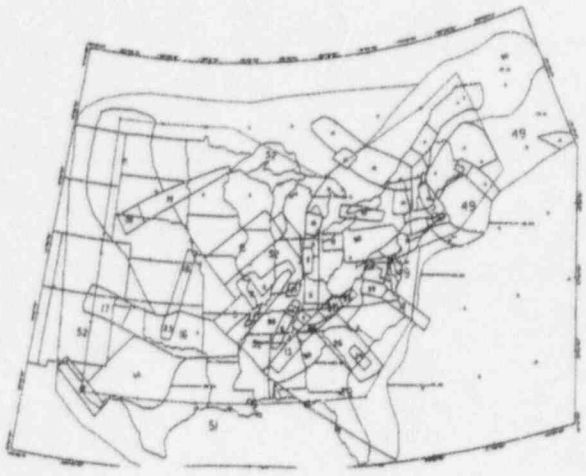

Figure 7-15 Example of an expert seismic source map.

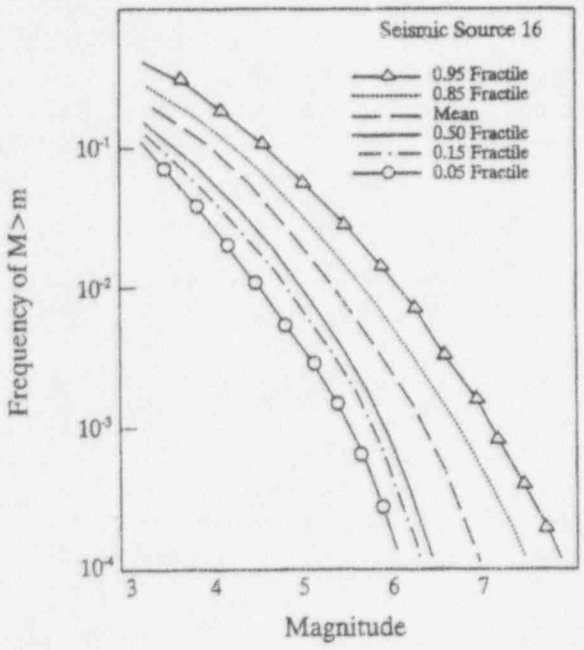

Figure 7-16 Illustration of magnitude recurrence for a seismic source.

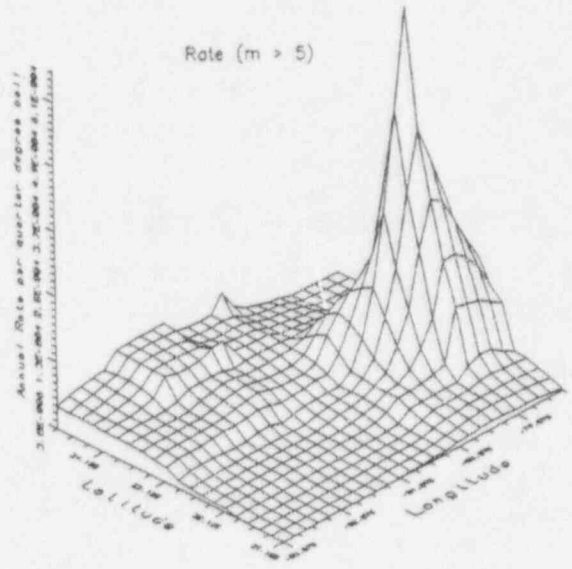

Figure 7-17 Map showing the mean rate of earthquake occurrences above magnitude 5.0.

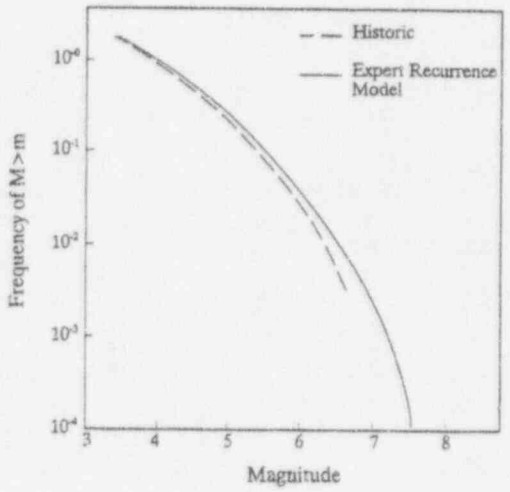

Figure 7-18 Comparison of historic seismicity and the expert recurrence model for a region.

Tier 2

The Tier 2 documentation of the SSC results should contain supporting documentation of the seismic source maps and calculation results of seismic source recurrence models.

7.11 Documentation Guidance: Results of Characterization of Ground Motion Attenuation

The graphical presentation of ground-motion-attenuation results can be done in a number of different ways. Perhaps the most difficult challenge in presenting this type of information is when the results from several different models are being displayed on the same figure. The problem here is to differentiate among many different types of curves, or types of results, on the same figure. The guidance, although obvious, is well worth repeating:

- keep everything legible
- use similar scales for graphs that the reader might want to compare

7.12 Computer-Readable Data Files

Tier 1 Documentation

As part of the PSHA documentation, it is recommended that computer-readable (friendly) data files be generated and retained by the analyst and the study sponsor, if desired. This section provides the guidance for documentation of these data files.

As part of the Tier 1 documentation of the PSHA, the analyst should describe the information that is retained in computer readable data files. The description should include:

1. The content of each type of data file that is retained
2. Potential uses of these data files
3. How the files can be used/accessed (i.e., Can a spreadsheet be used to read them?)

The purpose is to make the sponsor of the study and the technical reviewer aware of the content of the data files and their potential applications. This write-up need not contain a detailed description of each computer data file; rather, it should briefly describe their content and availability.

Tier 2 Documentation

As part of the PSHA Tier 2 documentation computer data files should be created that contain all of the intermediate and final calculations that form the basis for the Basic PSHA Results, the PSHA Deaggregated Results and the Sensitivity Analyses reported in the Tier 1 documentation. These guidelines identify the information that should be contained in these files; however, specific file formats, etc. are not specified.

The information contained in the Tier 2 documentation of the PSHA results should be comprehensive enough to permit the analyst or technical reviewer to conduct sensitivity evaluations, examine the impact of individual parameters to the results without having to rerun the hazard analysis, which may not be an option for the sponsor or the technical reviewer. Table 7-10 contains a summary of the information that should be provided in the Tier 2 documentation of the PSHA results. Figure 7-19 shows a tabular summary of a computer data file.

Table 7-10 PSHA Computer Data Files—Tier 2 Documentation

Item/Data File	Description
Computer Data File Documentation	Complete documentation of the computer data files generated as part of a PSHA project should be created and maintained. This should include a description of the file content and format.
Seismic Hazard Input Files	Computer data files that are the input to the PSHA calculation software. There may be one or multiple files that contain information for each seismic source (e.g., geometry data, seismicity parameters, fault rupture models, etc.) and the rules for combining seismic sources.
Seismic Source Hazard File	For each seismic source, data files should be provided that contain the hazard results for the alternative parameter values used to quantify the hazard and the probability values assigned to each parameter. Examples of alternative parameters include different estimates of the maximum magnitude, seismic activity rate, b-value, ground motion attenuation models, source geometry, etc. For each hazard curve, information should be provided that makes it possible to identify exactly the parameters used to produce that particular hazard result.
Total Hazard File	In a format similar to the hazard results for individual seismic sources, a data file should be provided that contains each of the individual hazard curves generated in the analysis. In addition to the information provided for each seismic source, this data file must identify the seismic sources that were included in the hazard curve determination and the probability weight assigned to the hazard result.
Deaggregated Hazard File	Data files should be provided for the deaggregated hazard results in the same way that the total hazard is documented. This should be done on a source-by-source basis and for all sources combined (see above).
Sensitivity Analyses	Data files should be provided for sensitivity analyses that are performed in the same way that the total hazard is documented (see above).

```
SAMPLE PROBLEM  TEAM 1, SITE 1                    89.80     35.81
A      1 1.00 1 6.5 0.85 1 0.50  3.120 3.723 4.768 5.770 7.049
A      1 1.00 2 6.8 0.15 1 0.50  3.087 3.658 4.628 5.531 6.659
A      1 1.00 1 6.5 0.85 2 0.50  2.617 3.128 4 090 5.042 6.251
A      1 1.00 2 6.8 0.15 2 0.50  2.602 3.093 3.998 4.872 5.961
B      3 0.75 1 6.0 0.33 1 0.50* 3.347 3.710 4.317 4.967 5.877
B      3 0.75 2 6.5 0.34 1 0.50* 3.285 3.638 4.226 4.820 5.598
B      3 0.75 3 6.9 0.33 1 0.50* 3.266 3.612 4.190 4.764 5.483
B      3 0.75 1 6.0 0.33 2 0.50* 3.028 3.391 4.031 4.679 5.575
B      3 0.75 2 6.5 0.34 2 0.50* 2.982 3.325 3.939 4.548 5.345
B      3 0.75 3 6.9 0.33 2 0.50* 2.971 3.304 3.902 4.494 5.245
B      4 0.25 1 6.0 0.33 1 0.50* 4.193 4.555 5.154 5.792 6.689
B      4 0.25 2 6.5 0.34 1 0.50* 4.121 4.472 5.045 5.619 6.373
B      4 0.25 3 6.9 0.33 1 0.50* 4.097 4.440 5.000 5.548 6.234
B      4 0.25 1 6.0 0.33 2 0.50* 3.852 4.232 4.871 5.509 6.393
B      4 0.25 2 6.5 0.34 2 0.50* 3.798 4.154 4.764 5.356 6.128
B      4 0.25 3 6.9 0.33 2 0.50* 3.783 4.129 4.719 5.289 6.008
C      4 1.00 1 5.5 0.85 1 0.50* 3.975 4.464 5.262 6.072 7.201
C      4 1.00 2 6.2 0.15 1 0.50* 3.702 4.161 4.886 5.567 6.442
C      4 1.00 1 5.5 0.85 2 0.50* 3.498 4.020 4.669 5.688 6.789
C      4 1.00 2 6.2 0.15 2 0.50* 3.250 3.724 4.503 5.223 6.129
```

Problem title and the site coordinates
Results for Source A

For each line the following is given:
- *Source name*
- *Seismicity option and its probability*
- *Maximum magnitude value and its probability*
- *Attenuation function and its probability*
- *Flag; '*' if this is a host source, else blank*
- *-log (base 10) probability of exceedance for each ground motion level*

Figure 7-19 Schematic illustration of the content of a Source Data File.

8. SEISMIC HAZARD GLOSSARY

TERM	DEFINITION
Acceleration (ground)	Acceleration at the ground surface produced by seismic waves. Typically expressed in g, the vertical acceleration of gravity at the earth's surface (9.80665 m/s^2).
Acceleration, Spectral	Pseudo-absolute response spectral acceleration, given as a function of period or frequency and damping ratio (typically 5%). It is equal to the peak relative displacement of a linear oscillator of frequency f attached to the ground, times the quantity $(2\pi f)^2$. It is expressed in g or cm/s^2.
Active Fault, Active Source	A fault or area source that on the basis of historical, seismological, or geological evidence is considered to have a non-zero probability of producing an earthquake in the present tectonic environment.
Activity Rate	See "Recurrence."
Aleatory Uncertainty	The uncertainty inherent in a non-deterministic (stochastic, random) phenomenon. Aleatory uncertainty is reflected by modeling the phenomenon in terms of a probability model. In principle, aleatory uncertainty cannot be reduced by the accumulation of more data or additional information. Sometimes called randomness.
Area Source	A region of the earth's crust that is assumed for PSHA to have relatively uniform seismic source characteristics. (See also "Seismic Source Zone").
Attenuation, Ground Motion	Decrease in severity (or amplitude) of ground shaking with increasing distance from the earthquake source.
Background Source	A regional scale area source. (Type IV source in text).
b-value	A parameter describing the decrease in the relative frequency of occurrence of earthquakes of increasing sizes. It is the slope of a straight line relating absolute or relative frequency (plotted logarithmically) to earthquake magnitude or intensity, the Gutenberg-Richter recurrence relationship.
Bandwidth	A range of frequencies or periods.

Complementary Zone	See "Background Zone."
Control Point	The location in the soil profile where the control motion is specified.
Control Motion	The input time history to a seismic site response analysis.
Convolution	Complex multiplication in the frequency domain. Used in site response analysis to take the ground motion at a given depth and "propagate" it upward through the soil column and in probability calculations.
Design Earthquake	The magnitude, distance, and other parameters representing the design ground motion.
Design Ground Motion	A specification of the seismic ground motion at a site used for the earthquake-resistant design of a structure.
Design Spectrum	A set of curves for design purposes that gives spectral acceleration, velocity, or displacement (usually absolute acceleration, pseudo-relative velocity, and relative displacement) of a single degree of freedom oscillator as a function of natural period of vibration and damping. (Alternate: The spectral representation of design ground motion).
Distance, Epicentral	Distance from the epicenter to a specific location (site).
Distance, Fault	Shortest distance from the fault to a specific location (site).
Distance, Hypocentral	Distance from the hypocenter to a specific location (site).
Distance, JB	Shortest distance from a point immediately above the ruptured portion of the fault to a specific location (site) (after Joyner and Boore, 1981).
Duration (of ground motion or earthquake rupture)	The length of time during which ground motion at a site shows certain characteristics (e.g., perceptibility, large amplitudes). (See "Corner Frequency").
Earthquake	A sudden motion or trembling of the earth caused by the abrupt release of slowly accumulated strain. The ground motion may range from violent at some locations to imperceptible at others. (Alternate: Naturally occurring shear failure of rock masses within the earth that gives rise to propagating seismic waves).

Epistemic Uncertainty	Uncertainty attributable to incomplete knowledge about a phenomenon which affects our ability to model it. Epistemic uncertainty is reflected in a range of viable models, multiple expert interpretations, and statistical confidence. In principle, epistemic uncertainty can be reduced by the accumulation of additional information. (See "Modeling Uncertainty").
Exceedance Probability	The probability that a specified level of ground motion for at least one earthquake will be exceeded at a site or in a region during a specified exposure time.
Expected Value	The average value, taken with respect to its probability distribution, of an aleatory (random) variable.
Expected Occurrence Rate	The expected value of the number of occurrences of an event (e.g., earthquakes) per unit area per unit time; generally denoted as ν.
Family of Hazard Curves	A set of hazard curves used to reflect the epistemic uncertainties associated with estimating seismic hazard. A common family of hazard curves used in describing the results of a PSHA are curves of fractiles of the probability distributions of estimated seismic hazard as a function of the level of ground motion parameter.
Fault	A planar or gently curved fracture surface or zone in the earth across which there has been relative displacement.
Fault, Dip-Slip	A fault in which the relative displacement is along the direction of the dip of the fault plane; either down-dip (normal fault) or up-dip (reverse fault).
Fault, Normal	A dip-slip fault in which the block above the fault has moved downward relative to the block below. This type of fault represents crustal extension.
Fault, Reverse	A dip-slip fault in which the block above the fault has moved upward relative to the block below, and the fault dip $> 45^{\circ}$.
Fault, Strike-Slip	A fault in which the relative displacement is along the strike of the fault plane, either right- or left-lateral.
Fault, Thrust	A dip-slip fault in which the block above the fault has moved upward relative to the block below, and the fault dip $< 45^{\circ}$. This type of fault represents crustal compression

Fault Zone	The zone of deformation comprising a fault.
Focal Mechanism	The combination of the dip angle of the fault and the direction of slip across the fault; faults are classified as strike-slip, normal or reverse. (See "Fault"). (Alternate: Geometrical representation of earthquake faulting expressed in terms of the strike and dip of the fault plane and the rake angle of the slip vector with respect to the fault plane).
Frequency, Corner	Frequency at which the amplitude spectrum of an earthquake transitions from a low-frequency level controlled by the seismic moment, to a high-frequency level controlled by the stress drop. $1/f_c$ is approximately the duration of the earthquake rupture.
Ground Motion Attenuation Model	An analytic model used to relate some measure of ground motion (peak ground acceleration, spectral acceleration, etc.) to distance, magnitude, source and path parameters. A variety of such models exist. A simple, commonly used form is $g(m,r) = C_1 + C_2 M + C_3 \log R + C_4 R$. The ground motion model is part of a model for observed ground motion measures, e.g., $\log A = g(m,r) + E$ where E denotes aleatory uncertainty. Inherent in the model of the observed ground motion measure is a model of the aleatory uncertainty, often taken to be a normal (Gaussian) probability distribution, i.e., $E \sim$ Normal $(0,\sigma)$ where σ, the standard deviation of E, quantifies the aleatory variability of the ground motion measure. If more complex models are considered, including source and path parameters, e.g., stress drop, and if any of these parameters are aleatory uncertain parameters, the model should include their (aleatory) probability distribution similar to that given for E above.
Gutenberg-Richter Relation	A model of the relationship between frequency and magnitude of earthquakes (in some specified region) expressed as $\log N = a - bM$ where N is the number of earthquakes with magnitude greater than M.
Hypocenter, focus	The point in the earth at which an earthquake is initiated.
Intensity	A measure of the effects of an earthquake at a particular place. Commonly used scales to specify intensity are the Rossi-Forel, Mercalli, and Modified Mercalli.

Lower Bound Magnitude

The lowest earthquake magnitude considered in deriving the seismic hazard curve for a site. (The choice of the lower bound magnitude is based on arguments that smaller earthquakes will not structurally damage well-engineered structures).

Magnitude

A measure of earthquake size, determined by taking the common logarithm (base 10) of the largest ground motion observed during the arrival of the P-wave or seismic surface wave and applying a standard correction for distance to the epicenter. (Alternatively: A measure of earthquake size).

Magnitude, Body-Wave

Magnitude derived from the largest displacement amplitude of body waves (P or S).

Magnitude, Coda-Wave

Magnitude derived from the amplitude and duration of the seismic coda.

Magnitude Distribution

The (conditional) aleatory probability distribution of earthquake magnitude, given the occurrence of an earthquake, assumed to be homogeneous at all locations throughout a source/subsource/seismic area.The probability distribution (given a sufficient number of earthquake events) is estimated as

$$f_M(m)\Delta m = \frac{\text{number of earthquakes with magnitude in } \Delta m}{\text{total number of observed earthquakes}}$$

where $f_M(m)$ is the probability density function. Due to lack of sufficient historical data, this distribution is often taken to be the Gutenberg-Richter relation.

Magnitude, Lg

Magnitude derived from the displacement amplitude of Lg waves; often used in Eastern North America because it can be accurately measured from typical low-gain seismographs at long distances from the source.

Magnitude, Moment

Earthquake magnitude derived from the seismic moment. Approximately equal to local magnitude for moderate earthquakes, and to surface-wave magnitude for large earthquakes.

Magnitude, Richter or Local (1935)

Common logarithm of the trace amplitude (in microns) of a standard Wood-Anderson seismograph located on firm ground 100 km from the epicenter. Correction tables are used to account for other distances and ground conditions.

Magnitude, Surface-Wave	Earthquake magnitude determined from the maximum amplitude of 20 s period surface waves.
Maximum Magnitude	The largest earthquake that a seismic source is capable of generating. The maximum magnitude is the upper-bound to recurrence curves.
Maximum Credible	The phrase used to specify the largest value of a variable, e.g., the magnitude of an earthquake, which might reasonably be expected to occur. A confusing term with no quantifiable definition. Not recommended for use in PSHA.
Mean	Average value of a set of data.
Mean Occurrence Rate Estimate	An estimate of the expected occurrence rate, usually taken as the total number of occurrences of an event (e.g., earthquakes) observed in a specified area and time interval divided by the area times length of time. (See "Rate of Seismicity").
Median (sample median)	Fiftieth fractile of the probability distribution of a variable. (Middle value of an ordered list of a set of data).
Modeling Uncertainty	The variability of a model predicted value from the value of the quantity being predicted. In principle, it can be reduced or eliminated by further testing, data accumulation, or more detailed modeling. It is one source of epistemic uncertainty. (Often called systematic uncertainty).
Outcrop Motion	Motion specified at the free surface of either a real or hypothetical bedrock outcrop at the ground surface. This motion thus represents the earthquake motion unaltered by surface soft soil layers.
Peak Acceleration	Maximum value of acceleration displayed on an accelerogram.
Peak Displacement	Maximum value of displacement obtained or calculated from a record of ground motion.
Peak Velocity	Maximum value of velocity obtained or calculated from a record of ground motion.
Randomness	See "Aleatory Uncertainty."

Rate of Seismicity	Rate of occurrence of earthquakes above some specified magnitude for a specified region.
Recurrence, Recurrence Rate, Recurrence Curve	The frequency of earthquake occurrence of various magnitudes often expressed by the Gutenberg-Richter relation.
Recurrence Interval	The mean time period between earthquakes of a given magnitude.
Recurrence Model	A model to express the relative number or frequency of earthquakes having different magnitudes. A common recurrence model is the exponential magnitude distribution.
Repeat Time	See "Recurrence Interval."
Response Spectrum	A set of curves that gives spectral acceleration, velocity, or displacement as a function of period of vibration and damping.
Return period	Commonly used to express the mean time period between ground motions of a particular amplitude (increase of annual frequency).
Seismic Hazard Curve	A plot of an estimate of the expected frequency of exceedence (over some specified time interval) of various levels of some characteristic measure of an earthquake (often peak ground acceleration). The time period of interest is often taken as one year, in which case the curve is called the annual frequency of exceedence.
Seismic Moment	A measure of the size of an earthquake based on interpretations of how much stress was relieved over the area of the fault or rupture surface. It is defined by the product of the rupture area, the average slip, and the crustal shear modulus.
Seismicity	Denotes the propensity for earthquakes to occur in a region and the possible magnitudes, locations and depths of these earthquakes.
Seismic Source	General term to define faults or area sources. (Types 1-4 in text).

Seismic Source Characteristics	The parameters that characterize a seismic source for PSHA, including source geometry, probability of activity, maximum magnitude, and earthquake recurrence.
Seismic Source Zone	See "Area Source."
Seismic Zone	A region showing relatively elevated levels of observed seismicity.
Seismogenic	Capable of generating tectonically significant earthquakes.
Seismotectonic Province	A region of the earth's crust having similar seismicity and tectonic characteristics.
Site Response (amplification)	The amplification (increase or decrease) of earthquake ground motion by rock and soil near the earth's surface in the vicinity of the site of interest. Topographic effects, the effect of the water table, and basin edge wave-propagation effects are sometimes included under site response.
Source Zone	See "Area Source."
Spatial Clustering	Observed or inferred proximity of earthquake occurrences.
Stationary Poisson Process	A probabilistic model of the occurrence of an event over time (space) characterized by the following properties: (1) the occurrence of the event in a small interval is constant over time (space), (2) the occurrence of two (or more) events in a small interval is "negligible," and (3) the occurrence of the event in non-overlapping intervals is independent. This model is often used to model the temporal and spatial occurrence of earthquakes within a source zone/seismic area.
Stress Drop	The average shear stress released across a rupture surface during an earthquake. (1 bar = 1.013×10^6 dyne/cm^2).
Tectonic Province	See "Seismotectonic Province."
Temporal Clustering	Occurrences of multiple closely timed earthquakes separated by longer periods of quiescence. Events that tend to cluster represent a deviation from a stationary Poisson process.

Upper Bound Magnitude	See "Maximum Magnitude."
Uncertainty	See "Epistemic Uncertainty" and "Aleatory Uncertainty."
Variance	The expected value, taken with respect to its probability distribution, of the squared deviation of an aleatory variable from its expected value.
Zonation	The process of developing seismic source maps (or a set of seismic zones).

REFERENCES

Abrahamson, N. A. and D. Sykora. 1994. "Variation of Ground Motions Across Individual Sites." In *Proceedings of the Fourth DOE Natural Phenomena Hazards Mitigation Conference*. U.S. Department of Energy.

Abrahamson, N. A. and J. J. Litehiser. 1989. "Attenuation of Vertical Peak Acceleration." *Bull. Seism. Soc. Am.*, 79:549-580.

Abrahamson, N. A., P. G. Somerville, and C. A. Cornell. 1990. "Uncertainty in Numerical Strong Motion Predictions." In *Proceedings of the Fourth U.S. National Conference on Earthquake Engineering*, Palm Springs, California, I (May 20-24):407-416. .

Algermissen, S. T., E. V. Leyendecker, G. A. Bollinger, N. C. Donovan, J. E. Ebel, W. B Joyner, R. W. Luft, and J. P. Singh. 1991. "Probabilistic Ground-Motion Hazard Maps of Response Spectral Ordinates for the United States." In *Proceedings of the Fourth International Conference on Seismic Zonation*, Earthquake Engineering Research Institute, Oakland, California, II:687-694.

Amendola, A. 1986. "Uncertainties in Systems Reliability Modeling: Insights Gained Through the European Benchmark Exercises." *Nuclear Engineering and Design*, 93:215-225.

Amos, C., et al. 1987. *Evaluation of Severe Accident Risks and the Potential for Risk Reduction: Peach Bottom, Unit 2*. Report NUREG/CR-4551, vol. 3. Nuclear Regulatory Commission, Washington, D.C.

Anderson, J. G. 1991. "Strong Motion Seismology." In U.S. National Report to International Union of Geology and Geophysics 1987-1990. *Reviews of Geophysics*, Supplement, 700-720.

Anderson, J. G. and S. E. Hough. 1984. "A Model for the Shape of the Fourier Amplitude Spectrum of Acceleration at High Frequencies." *Bull. Seism. Soc. Am.*, 74:1969-1993.

ANS/IEEE. 1983. *PRA Procedures Guide*. Report NUREG/CR-2300. Nuclear Regulatory Commission, Washington, D.C.

Apostolakis, G. 1990. "The Concept of Probability in Safety Assessments of Technological Systems." *Science*, 250:1359-1364.

Apostolakis, G. 1993. *A Commentary on Model Uncertainty*. In *Proceedings of the Workshop on Model Uncertainty: Its Characterization and Quantification*, Annapolis, MD. A. Mosleh, N.O. Siu, C. Smidts, and C. Lui, eds.

Atkinson, G. M. 1993a. "Notes on Ground Motion Parameters for Eastern North America: Duration and H/V Ratio." *Bull. Seism. Soc. Am.*, 83:587-596.

Atkinson, G. M. 1993b. "Earthquake Source Spectra in Eastern North America." *Bull. Seism. Soc. Am.*, 83:1778-1798.

Atkinson, G. M. 1995. "Optimal Choice of Magnitude Scales for Seismic Hazard Estimates in Eastern North America." *Seism. Res. Lett.*, 66:51-55.

Atkinson, G. M. and D. M. Boore. 1990. "Recent Trends in Ground Motion and Spectral Response Relations for North America." *Earthquake Spectra*, 6:15-35.

Atkinson, G. M. and D. M. Boore. 1995. "Ground Motion Relations for Eastern North America." *Bull. Seism. Soc. Am.*, 85:17-30.

Atkinson, G. M. and R. R. Mereu. 1992. "The Shape of Ground Motion Attenuation Curves in Southeastern Canada." *Bull. Seism. Soc. Am.*, 82:2014-2031.

Bender, B. 1986. "Modeling Source Zone Boundary Uncertainty in Seismic Hazard Analysis." *Bull. Seism. Soc. Am.*, 76:329-342.

Benioff, H. 1955. "Mechanism and Strain Characteristics of the White Wolf Fault as Indicated by the Aftershock Sequence." In *Earthquakes in Kern County, California During 1952*. Gordon B. Oakeshott, ed. *California Division of Mines Bulletin*, 171:199-202.

Benjamin, J. R. and C. A. Cornell. 1970. *Probability, Statistics, and Decision for Civil Engineers*. New York: McGraw-Hill.

Bernreuter, D. L., et al. 1981. *Seismic Hazard Analysis*. Report NUREG/CR-1582, vols. 1-5, prepared by Lawrence Livermore National Laboratory for the U.S. Nuclear Regulatory Commission.

Bernreuter, D. L., J. B. Savy, R. W. Mensing, and J. C. Chen. 1989. *Seismic Hazard Characterization of 69 Nuclear Plant Sites East of the Rocky Mountains*. Report NUREG/CR-5250, vols. 1-8, prepared by Lawrence Livermore National Laboratory for the U.S. Nuclear Regulatory Commission.

Boatwright, J. and D. M. Boore. 1982. "Analysis of the Ground Accelerations Radiated by the 1980 Livermore Valley Earthquakes for Directivity and Dynamic Source Characteristics." *Bull. Seism. Soc. Am.*, 72:1843-1865.

Bollinger, G. A., M. C. Chapman, and M. S. Sibol. 1993. "A Comparison of Earthquake Damage Areas as a Function of Magnitude Across the United States." *Bull. Seism. Soc. Am.*, 83:1064-1080.

Bonduelle, Y. 1987. *Aggregating Expert Opinions by Resolving Sources of Disagreement*. Ph.D. Dissertation, Stanford University.

Bonilla, M. G., R. K. Mark, and J. J. Lienkaemper. 1984. "Statistical Relations Among Earthquake Magnitude, Surface Rupture Length, and Surface Fault Displacement." *Bull. Seism. Soc. Am.*, 74:2379-241.

Bonano, E. J. and G. Apostolakis. 1991. "Theoretical foundations and practical issues for using expert judgments in uncertainty analysis of high-level radioactive waste disposal." *Radioactive Waste Management and the Fuel Cycle*, 16: 137-159.

Boore, D. M. 1983. "Stochastic Simulation of High-Frequency Ground Motions Based on Seismological Models of the Radiated Spectra." *Bull. Seism. Soc. Am.*, 73:1865-1894.

Boore, D. M. 1989. "The Richter Scale: its Development and Use for Determining Earthquake Source Parameters." *Technophysics*, 166:1-14.

Boore, D. M. and G. M. Atkinson. 1987. "Stochastic Prediction of Ground Motion and Spectral Response Parameters at Hard-Rock Sites in Eastern North America." *Bull. Seism. Soc. Am.*, 77:440-467.

Boore, D. M. and J. Boatwright. 1984. "Average Body-Wave Radiation Coefficients." *Bull. Seism. Soc. Am.*, 74:1615-1621.

Boore, D. M. and W. B. Joyner. 1988. "Measurement, Characterization, and Prediction of Strong Ground Motion." In Earthquake Engineering and Soil Dynamics II—Recent Advances in Ground Motion Evaluation, J.L. von Thun, ed. *ASCE Geotech.*, Pub. No. 20.

Boore, D. M. and W. B. Joyner 1989. "The Effect of Directivity on the Stress Parameter Determined from Ground Motion Observations." *Bull. Seism. Soc. Am.*, 79:1984-1988.

Boore, D. M., W. B. Joyner, and T. E. Fumal. 1993. "Estimation of Response Spectra and Peak Accelerations from Western North America Earthquakes: An Interim Report." *U.S. Geol. Surv. Open-File Rept.*

Boore, D. M., W. B. Joyner, and T. E. Fumal. 1994a. "Estimation of Response Spectra and Peak Accelerations from Western North American Earthquakes: An Interim Report: Part 2." *U.S. Geol. Surv. Open-File Rept.*

Brune, J. N. 1970. "Tectonic Stress and the Spectra of Seismic Shear Waves from Earthquakes." *J. Geophys. Res.*, 75:4997-5009.

Call, H. J. and W. A. Miller. 1990. "A Comparison of Approaches and Implementations for Automating Decision Analysis." *Reliability Engineering and System Safety*, 30:115-162.

Campbell, K. W. 1987. "Predicting Strong Ground Motion in Northern Utah." In *Evaluation of Regional and Urban Earthquake Hazards and Risk in Utah*, W.W. Hays and P.L. Gori, eds. *U.S. Geol. Surv. Open-File Rept.*

Campbell, K. W. and Y. Bozorgnia. 1994. "Near-Source Attenuation of Peak Horizontal Acceleration from Worldwide Accelerograms Recorded from 1957 to 1993." In *Proceedings from the Fifth U.S. National Conference on Earthquake Engineering*, Chicago, Illinois, 3 (July 10-14):283-292.

Chibber, S. and G. Apostolakis. 1993. "Some Approximations Useful to the Use of Dependent Information Sources." *Reliability Engineering and System Safety*, 42:67-86.

Chibber, S., G. Apostolakis, and D. Okrent. 1994. "On the Use of Expert Judgments to Estimate The Pressure Increment in the Sequoyah Containment at Vessel Breach." *Nuclear Technology*, 105:87-103.

Clemen, R. T. 1985. "Extraneous Expert Information." *Journal of Forecasting*, 4:329.

Clemen, R. T. 1987. "Combining Overlapping Information." *Management Science*, 33:332-340.

Clemen, R. T. 1989. "Combining Forecasts: A Review and an Annotated Bibliography." *International Journal of Forecasting*, 5:559-583.

Clemen, R. T., and R. L. Winkler. 1992. "Sensitivity of Weights in Combining Forecasts." *Operations Research*, 40, no. 3:609-614.

Clemen, R. T., and R. L. Winkler. 1985. "Limits for the Precision and Value of Information from Dependent Sources." *Operations Research*, 33:427-42.

Cooke, R.M. 1991. *Experts in Uncertainty: Expert Opinion and Subjective Probability in Science*. New York: Oxford University Press.

Coppersmith, K. J. 1991. "Seismic Source Characterization for Engineering Seismic Hazard Analysis." In *Proceedings of the Fourth International Seismic Zonation Conference*, 1:1-60.

Coppersmith, K. J. and R. R. Youngs 1986. "Capturing Uncertainty in Probabilistic Seismic Hazard Assessment Within Intraplate Tectonic Environments." In *Proceedings of the Third U.S. National Conference on Earthquake Engineering*, Earthquake Engineering Research Institute, 1:301-312.

Cornell, C. A. 1993. "Which 84th Percentile Do You Mean?" In *Proceedings of the Fourth DOE Natural Phenomena Hazards Mitigation Conference*, 152-159.

Cornell, C. A. 1968. "Engineering Seismic Risk Analysis." *Bull. Seism. Soc. Am.*, 58:1583-1606.

Cornell, C.A., and S. R. Winterstein. 1988. "Temporal and Magnitude Dependence in Earthquake Recurrence Models." *Bull. Seism. Soc. Am.*, 27, no. 4.

Cornell, C.A., H. Banon and A.F. Shakal. 1977. *Seismic Motion and Response Prediction Alternatives*. Massachusetts Institute of Technology, MIT CE R77-34 no. 579, Cambridge, MA.

De Bejar, L. A. and K. Ganapathi. 1992. "A Note on Inelastic Response Spectra for Systems with Bilinear Spring Force and Kinematic Strain Hardening." *Earthquake Spectra*, 8:181-200.

De Finetti, B. 1974. .*Theory of Probability,* vols. 1 and 2, New York :Wiley.

De Groot, M. 1974. "Reaching a Consensus." *J. Amer. Statist. Assoc.*, 69:118-121.

De Groot, M.H. 1988. "Modern Aspects of Probability and Utility." In *Accelerated Life Testing and Expert Opinions in Reliability*, C.A. Clarotti and D.V. Lindley, eds. North Holland, Amsterdam.

DeWispelare, A.R., L.T. Herren, M.P. Miklas, and R.T. Clemen. 1993. *Expert Elicitation of Future Climate in the Yucca Mountain Vicinity*. Report NRC-02-88-005, Center for Nuclear Waste Regulatory Analyses, San Antonio, Texas.

Electric Power Research Institute (EPRI). 1989. *Probabilistic Seismic Hazard Evaluations at Nuclear Power Plant Sites in the Central and Eastern United States: Resolution of the Charleston Earthquake Issue*. Prepared by Risk Engineering Inc., Yankee Atomic Power Company and Woodward Clyde Consultants, EPRI Report NP-6395-D.

Electric Power Research Institute (EPRI). 1985. *Seismic Source Method of Seismic Hazard Analysis and EQZONE Computer Code Manual*. EPRI NP-4349-CCM.

Electric Power Research Institute (EPRI). 1988. *Seismic Hazard Methodology for Central and Eastern United States*. EPRI NR-4726, vols. 1-10.

Electric Power Research Institute (EPRI). 1989. *Probabilistic Seismic Hazard Evaluations at Nuclear Power Plant Sites in the Central and Eastern United States: Resolution of the Charleston Earthquake Issue*. EPRI Report NP-6395-D.

Electric Power Research Institute (EPRI). 1993. *Earthquakes of Stable Continental Regions, Evaluation of Large Earthquake Potential*. EPRI Research Project 2556-12, vols. 1-6.

Electric Power Research Institute (EPRI). 1993. *Guidelines for Determining Design Basis Ground Motions*. Rept. No. EPRI TR-102293, vols. 1-5.

Electric Power Research Institute (EPRI). 1986-1991. *Seismic Hazard Methodology for the Central and Eastern United States*. Rept. No. NP-4726, vols. 1-10.

Genest, C. and J. Zidek. 1986. "Combining Probability Distributions: A Critique and an Annotated Bibliography." *Statistical Science*, 1:114-148.

Global Seismic Hazard Analysis Program. 1993. "Special Issue: Technical Planning Volume." *Annali di Geofiscia*.

Gokhale, D. V. and S. J. Press. 1982. "Assessment of a Prior Distribution for the Correlation Coefficient in a Bivariate Normal Distribution." *Journal of the Royal Statistical Society*, Series A, 145:237-249.

Gutenberg, B. 1955. "Magnitude Determination for Larger Kern County Shocks, 1952; Effects of Station Azimuth and Calculation Methods." In Earthquakes in Kern County, California, during 1952, Gordon B. Oakeshott, ed. *California Division of Mines Bulletin*, 171:171-175.

Hanks, T. C and H. Kanamori. 1979. "A Moment Magnitude Scale." *J. Geophys. Res.* 84:2348-2350.

Hanks, T. C. 1982. "f_{max}." *Bull. Seism. Soc. Am.*, 72:1867-1879.

Hanks, T. C. and A. C. Johnston. 1992. "Common Features of the Excitation and Propagation of Strong Ground Motion for North American Earthquakes." *Bull. Seism. Soc. Am.*, 82:1-23.

Hanks, T. C. and D. M. Boore. 1984. Moment-Magnitude Relations in Theory and Practice. *J. Geophys. Res.*, 89:6229-6235.

Hanks, T. C. and R. K. McGuire. 1981. "The Character of High-Frequency Strong Ground Motion." *Bull. Seism. Soc. Am.*, 71:2971-2095.

Helmberger, D., D. Dreger, R. Stead and H. Kanamori. 1993. "Impact of Broadband Seismology on the Understanding of Strong Motions." *Bull. Seism. Soc. Am.*, 83:830-850.

Herrero, A. and P. Bernard. 1994. "A Kinematic Self-Similar Rupture Process for Earthquakes." *Bull. Seism. Soc. Am.*, 84:1216-1228.

Herrmann, R. B. 1985. "An Extension of Random Vibration Theory Estimates of Strong Ground Motion to Large Distances." *Bull. Seism. Soc. Am.*, 75:1447-1453.

Hora, S.C. and R. L. Iman. 1989. "Expert Opinion in Risk Analysis: The NUREG-1150 Methodology." *Nuclear Science and Engineering*, 102:323-331.

Horton, S. P. 1994. "Simulation of Strong Ground Motion in Eastern North America." In *Proceedings of the Fifth U.S. National Conference on Earthquake Engineering*, Chicago, Illinois, 3 (July 10-14):251-260.

Howard, R. and J. Matheson. 1981. "Influence Diagrams." In *The Principles and Applications of Decision Analysis*, R. Howard and J. Matheson, eds., Strategic Decisions Group, Menlo Park, California.

Iman, R.L., and M. J. Shortencarier. 1984. *A FORTRAN 77 Program and User's Guide for the Generation of Latin Hypercube and Random Samples for Use with Computer Models*. Report NUREG/CR-3624, SAND83-2365, Sandia National Laboratories, Albuquerque, NM.

Johnston, A. C., and L. R. Kanter. 1990. "Earthquakes in Stable Continental Crust." *Scientific American*, 262:68-75.

Johnston, A. C., K. J. Coppersmith, L. R. Kanter, and C. A. Cornell. 1993. "The Earthquake of Stable Continental Regions." Rept. No. EPRI TR-102261, vols. 1-5, Palo Alto, California.

Joyner, W. B. 1984. "A Scaling Law for the Spectra of Large Earthquakes." *Bull. Seism. Soc. Am.*, 74:1167-1188.

Joyner, W. B. 1987. "Strong-Motion Seismology." In U.S. National Report to International Union of Geology and Geophysics 1983-1986. *Reviews of Geophysics*, 25:1149-1160.

Joyner, W. B. and D. M. Boore. 1981. "Peak Horizontal Acceleration and Velocity From Strong-Motion Records Including Records from the 1979 Imperial Valley, California, Earthquake." *Bull. Seism. Soc. Am.*, 71:2011-2038.

Joyner, W. B. and D. M. Boore. 1988. "Measurement, Characterization, and Prediction of Strong Ground Motion." In *Earthquake Engineering and Soil Dynamics II*, Proceedings of the Am. Soc. Civil Eng. Geotech. Eng. Div. Specialty Conf., June 27-30, 1988, Park City, Utah, 43-102.

Kanamori, H., H. K. Thio, D. Dreger, E. Hauksson, and T. Heaton. 1992. "Initial Investigation of the Landers, California, Earthquake of 28 June 1992 Using TERRAscope." *Geophys. Res. Lett.*, 19:2267-2270.

Kaplan, S. and B. J. Garrick. 1981. "On the Quantitative Definition of Risk." *Risk Analysis*, 1:11-27.

Kaplan, S. 1981. "On the Method of Discrete Probability Distributions in Risk and Reliability Calculation-Application to Seismic Risk Assessment." *Risk Analysis*, 1:189-196.

Kaplan, S. 1992. "Expert Information vs. Expert Opinions—Another Approach to the Problem of Eliciting/Combining/Using Expert Opinion in PRA." *Reliability Engineering and System Safety*, 35:61-72.

Keeney, R. L. and D. von Winterfeldt. 1991. "Eliciting Probabilities from Experts in Complex Technical Problems." *IEEE Transactions on Engineering Management*, 38:191-201.

Kennedy, R., S. Short, T. Kipp, H. Banon, F. Tokarz, and K. Merz. 1984. *Engineering Characterization of Ground Motion*. Report NUREG/CR-3805, U.S. Nuclear Regulatory Commission.

Krawinkler, H., A. Nassar, and M. Rahnama. 1991. "Damage Potential of Loma Prieta Ground Motions." *Bull. Seism. Soc. Am.*, 81:2048-2069.

Lee, P. M. 1989. *Bayesian Statistics*. New York: Oxford University Press.

Lindley, D. V. 1988. "The Use of Probability Statements." In *Accelerated Life Testing and Experts' Opinions in Reliability*, C. A. Clarottl and D. V. Lindley, eds., North Holland, Amsterdam.

Linstone, H. A. and M. Turoff. 1975. *The Delphi Method, Techniques and Applications*. Reading Massachusetts: Addison-Wesley.

Martin, G., ed. 1994. Proceedings. NCEER, SEAOC, BSSC Workshop on Site Response During Earthquakes and Seismic Code Provisions. University of Southern California.

McGuire, R. K. 1978. "Seismic Ground Motion Parameter Relations." In *Proceedings of the Proc. Am. Soc. Civil Eng., J. Geotech. Eng. Div.*, 104:481-490.

Meyer, M., and J. Booker. 1991. *Eliciting and Analyzing Expert Judgment: A Practical Guide*. New York: Academic Press.

Milne, W. G. and A.G. Davenport. 1969. "Distribution of Earthquake Risk in Canada." *Bull. Seism. Soc. Am.*, 59, no. 2:729-754.

Morris, P. A. 1971. *Bayesian Expert Resolution*. Ph.D. Dissertation, Dept. of Engineering-Economic Systems, Stanford University.

Morris, P. A. 1974. "Decision Analysis Expert Use." *Management Science*, 20, no. 9:1233.

Morris, P. A. 1977. "Combining Expert Judgments: A Bayesian Approach." *Management Science*, 23, no. 7:679.

Morris, P. A. 1983. "An Axiomatic Approach to Expert Resolution." *Management Science*, 29, no. 1:24.

Morris, P. A. 1986. "Observations on Expert Aggregation." *Management Science*. 32, no. 3.

National Research Council. 1988. *Probabilistic Seismic Hazard Analysis*. Washington, D.C.: National Academy Press.

Newmark, N. M. and W. J. Hall. 1978. *Development of Criteria for Seismic Review of Selected Nuclear Power Plants*. Report NUREG/CR-0098, U.S. Nuclear Regulatory Commission, May.

Nuttli, O. W. 1973. "Seismic Wave Attenuation and Magnitude Relations for Eastern North America." *J. Geophys. Res.*, 78:876-885.

Nuttli, O.W., and R.B. Herrmann. 1984. "Ground Motions of Mississippi Valley Earthquakes." *J. of Technical Topics in Civil Engineering*, ASCE, 110 (May):54-69.

Oliver, R. M. and J. Q. Smith. 1990. *Influence Diagrams, Belief Nets and Decision Analysis*. New York: Wiley.

Otway, H. and D. von Winterfeldt. 1992. "Expert Judgment in Risk Analysis and Management: Process, Context, and Pitfalls." *Risk Analysis*, 12:83-93.

Ou, G. B. and R. B. Herrmann. 1990. "A Statistical Model for Ground Motion Produced by Earthquakes at Local and Regional Distances." *Bull. Seism. Soc. Am.*, 80:1397-1417.

Pacific Gas and Electric Company. 1988. *Final Report of the Diablo Canyon Long-Term Seismic Program*. U.S. Nuclear Regulatory Commission docket nos., 50-275 and 50-323.

Papageorgiou, A. S. 1988. On Two Characteristic Frequencies of Acceleration Spectra: Patch Corner Frequency And f_{max}. *Bull. Seism Soc. Am.*, 78:509-529.

Paté-Cornell, M. E. and P. S. Fischbeck. 1992. "Aversion to Epistemic Uncertainties in Rational Decision Making: Effects on Engineering Risk Management." In *Risk-Based Decision Making in Water Resources V*, Y.Y. Haimes, D.A. Moser, and E.Z. Stakhiv, eds. New York: American Society of Civil Engineers.

Press, S.J. 1989. *Bayesian Statistics: Principles, Models, and Applications*. New York: Wiley.

Reiter, L. 1990. *Earthquake Hazard Analysis*. New York: Columbia University Press.

Richter, C. F. 1935. An Instrumental Earthquake Magnitude Scale. *Bull. Seism. Soc. Am.*, 25:1-31.

Ripley, B. D. 1981. *Spatial Statistics*. New York: Wiley.

Sadigh, K. R. 1993. "A Review of Attenuation Relationships for Rock Site Conditions from Shallow Crustal Earthquakes in an Interplate Environment." In *Proceedings of the International Workshop on Strong Motion Data*, Menlo Park, California, 2 (December 13-17):179-236.

Saikia, C. K. 1994. Modified Frequency-Wavenumber Algorithm for Regional Seismograms Using Filon's Quadrature—Modeling Lg Waves in Eastern North America. *Geophys. J. International*, 118:142-158.

Saikia, C. K. and P.G. Somerville. 1995. "Simulated hard-rock motions in Saint Louis, Missouri, from Large New Madrid Earthquakes (abstract)." *Seismological Research Letters*, March.

Savage, L. J. 1972. *The Foundations of Statistics*. New York: Dover Publications.

Savy, J. B., A. C. Boissonnade, R. W. Mensing, and C. M. Short. 1993. *Eastern U.S. Seismic Hazard Characterization Update*. Report UCRL-ID-115111, Lawrence Livermore National Laboratory.

Schwartz, D. P. and K. J. Coppersmith. 1984. "Fault Behavior and Characteristic Earthquakes: Examples from the Wasatch and San Andreas Faults." *Journal of Geophysical Research*, 89:5681-5698.

Schwartz, D. P. and K. J. Coppersmith. 1986. "Seismic Hazards: New Trends in Analysis Using Geologic Data." In *Active Tectonics*, R. E. Wallace, ed., pp. 215-230, Washington, D.C.: Academic Press.

Schwartz, D. P. 1988. "Geology and Seismic Hazards: Moving into the 1990's." In *Earthquake Engineering Soil Dynamics II-Recent Advances in Ground Motion Evaluation*, J.L. Von Thun, ed., ASCE Geotechnical Special Publication 20:1-42.

Sewell, R. T. 1988. "Damage Effectiveness of Earthquake Ground Motion: Characterizations Based on the Performance of Structures and Equipment." Ph. D. Thesis, Stanford University, vol. 1, 340 pages.

Shachter, R. D. 1988. "Probabilistic Influence and Influence Diagrams." *Operations Research*, 36:589.

Shakal, A. F. and D. L. Bernreuter. 1981. *Empirical Analyses of Near-Source Ground Motion*. Report NUREG/CR-2095, U.S. Nuclear Regulatory Commission.

Silva, W. J. and R. K. Green. 1989. "Magnitude and Distance Scaling of Response Spectral Shapes for Rock Sites with Applications to North American Tectonic Environment." *Earthquake Spectra*, 5:591-624.

Siu, N. and G. Apostolakis. 1981. "Probabilistic Models of Cable Tray Fires." *Reliability Engineering*, 3:213-227.

Siu, N. and G. Apostolakis. 1985. "On the Quantification of Modeling Uncertainties." From the *8th International Conference on Structural Mechanics in Reactor Technology (SMiRT)*, Brussels, Belgium.

Slemmons, D. B. 1977. *Faults and Earthquake Magnitude*. U.S. Army Corps of Engineers, Waterways Experiment Station, Misc. Papers S-73-1, Report 6.

Somerville, P. G. 1992. "Engineering Applications of Strong Ground Motion Simulation." *Tectonophysics*, 218:195-219.

Somerville, P. G. and R. Graves. 1993. "Conditions That Give Rise to Unusually Large Long-Period Ground Motions." *The Structural Design of Tall Buildings*, 2:211-232.

Somerville, P., M. Sen, and B. Cohee. 1981. "Simulations of Strong Ground Motions Recorded During the 1985 Michoacan, Mexico and Valparaiso, Chile Earthquakes." *Bull. Seism. Soc. Am.*, 81:1-27.

Toro, G. R. and R. K. McGuire. 1987. "An Investigation into Earthquake Ground Motion Characteristics in Eastern North America." *Bull. Seism. Soc. Am.*, 77:468-489.

Toro, G. R. and R. K. McGuire 1991). Assessment of the 1988 Saguenay Earthquake-Implications on Attenuation Functions, Seismic Hazard Analysis, Risk Engineering, Inc., Golden, CO, Published by Oak Ridge National Laboratory as ORNL/TM-11897.

Toro, G. R., W. J. Silva, R. K. McGuire, and R. B. Herrmann. 1992. "Probabilistic Seismic Hazard Mapping of the Mississippi Embayment." *Seism. Res.* Let., 63:449-475.

Trauth, K. M., S. C. Hora and R. V. Guzowski. 1993. *Expert Judgment on Markers to Deter Inadvertent Human Intrusion into the Waste Isolation Pilot Plant*. Report SAND92-1382, Sandia National Laboratories, Albuquerque, NM.

Trifunac, M. D. and A. G. Brady. 1975. "On the Correlation of Seismic Intensity Scales with the Peaks of Recorded Strong Ground Motion." *Bull. Seism. Soc. Am.*, 65:139-162.

Trifunac, M. D. and V. W. Lee. 1979. *Dependence of Pseudo Relative Velocity Spectra of Strong Motion Acceleration on the Depth Of Sedimentary Deposits*. Report No. CE 79-02, Dept. Civil Eng., Univ. of Southern California, Los Angeles.

Tversky, A. and D. Kahneman. 1974. "Judgment Under Uncertainty: Heuristics and Biases." *Science*, 185:1124-1141.

U.S. Atomic Energy Commission. 1973. Design Response Spectra for Seismic Design of Nuclear Power Plants. *Regulatory Guide* 1.60, 4p.

Van Dyck, J. F. M. 1985. *Statistical Analysis of Earthquake Catalogs*. Ph.D. Thesis, Department of Civil Engineering, Massachusetts Institute of Technology, December.

Velasco, A. A., C. J. Ammon, and T. Lay. 1994. "Empirical Green Function Deconvolution of Broadband Surface Waves: Rupture Directivity of the 1992 Landers, California (M_W = 7.3), Earthquake." *Bull. Seism. Soc. Am.*, 84:735-750,

Veneziano, D. 1988. "The Use of Intensity Data in Ground Motion Estimation." In *Proceedings: Earthquake Ground-Motion Estimation in Eastern North America*, Electric Power Research Institute Report EPRI NP-5875, 5–1 to 5–57.

Veneziano, D. 1994. "Uncertainty and Expert Opinion in Geologic Hazards." Presented at *The Earth, Engineers, and Education Symposium in Honor of R.V. Whitman*, Massachusetts Institute of Technology, October 7-8.

Veneziano, D. and A. L. Pais. 1986. "Automatic Source Identification Based on Historical Seismicity." In *Proceedings of the 8th European Conference on Earthquake Engineering*, Lab. Nacional de Engenharia Civil, Lisbon, 1, no. 2.1:63-70.

Veneziano, D. and Luc Chouinard. 1987. "Local Models of Seismicity and Their Estimation." In *Proc., Symposium on Seismic Hazards, Ground Motions, Soil Liquefaction and Engineering Practice in Eastern North America*, K. H. Jacob, ed., National Center for Earthquake Engineering Research Research, Rept. No. NCEER-87-0025.

Veneziano, D., C. A. Cornell, and T. O'Hara, 1984. *Historic Methods of Seismic Hazard Analysis*. EPRI NP 3438, Electric Power Research Institute, Palo Alto, California.

Webster's, 1988. *New World Dictionary*. New York: Simon & Schuster.

Weichert, D. H. 1980. "Estimation of the Earthquake Recurrence Parameters for Unequal Observation Periods for Different Magnitudes." *Bull. Seism. Soc. Am.*, 70:1337-1346.

Wells, D. L. and K. J. Coppersmith. 1994. "New Empirical Relationships Among Magnitude, Rupture Length, Rupture Width, Rupture Area and Surface Displacement." *Bull. Seism. Soc. Am.*, 84, no. 4:974-1002.

Winkler, R. L. and W. L. Hays. 1975. *Statistics*. New York: Holt, Reinehart and Winston.

Winkler, R. L. 1968. "The Consensus of Subjective Probability Distributions." *Management Science*, 15, no. 2:B-61.A

Winkler, R. L. 1981. "Combining Probability Distributions from Dependent Information Sources." *Management Science*, 27, no. 4:479.

Winkler, R. L. 1986. "Expert Resolution." *Management Science*, 32, no. 3.

Winkler, R. L. 1993. "Model Uncertainty: Probabilities for Models?" In *Proceedings of the Workshop on Model Uncertainty: Its Characterization and Quantification*, Annapolis, MD, October 20-22, A. Mosleh, N. Siu, C. Smidts, and C. Lui, eds.

Wyss, M. 1979. "Estimating Maximum Expectable Magnitude of Earthquakes from Fault Dimensions." *Geology*, 7:336-340.

Youngs, R. R. and K. J. Coppersmith. 1985. "Implications of Fault Slip Rates and Earthquake Recurrence Models to Probabilistic Seismic Hazard Estimates." *Bull. Seism. Soc. Am.*, 75:939-964.

Youngs, R. R., K. J. Coppersmith, C. L. Taylor, M. S. Power, L. A. DiSilvestro, M. L. Angell, N. T. Hall, J. R. Wesling, and L. Mualchin. 1993. "A Comprehensive Seismic Hazard Model for the San Francisco Bay Region." In *Proceedings of the Conference on Seismic Hazards in the Eastern San Francisco Bay Region.*

Zeng, Y., J. G. Anderson, and G. Yu. 1993. "A Composite Source Model for Computing Realistic Synthetic Strong Ground Motions." In *Proceedings of the International Workshop on Strong Motion Data*, Menlo Park, California, December 13-17, 2:439-451.

Appendix: Review Report by the National Academy of Sciences/ National Research Council

Review of

Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts

Panel on Seismic Hazard Evaluation
Committee on Seismology
Board on Earth Sciences and Resources
Commission on Geosciences, Environment, and Resources
National Research Council

NATIONAL ACADEMY PRESS
Washington, DC 1997

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

This report has been reviewed by a group other than the authors according to procedures approved by a Report Review Committee consisting of members of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine.

This report was prepared with the support of the U.S. Nuclear Regulatory Commission under Award No. NRC-G-04-92-117. The opinions, findings, conclusions and recommendations expressed herein are those of the author(s) and do not necessarily reflect the views of the U.S. Nuclear Regulatory Commission.

Additional copies of this report are available from:

Committee on Seismology
Board on Earth Sciences and Resources
National Research Council
2101 Constitution Avenue, NW
HA 372
Washington, DC 20418
202-334-2744

The cover photo, courtesy of Walter Hays, U.S. Geological Survey, shows damage to Nimitz Freeway resulting from the Loma Prieta earthquake of 1989.

Printed in the United States of America

PANEL ON SEISMIC HAZARD EVALUATION

CARL KISSLINGER, *Chairman,* University of Colorado, Boulder
KEIITI AKI, University of Southern California, Los Angeles
WALTER J. ARABASZ, University of Utah, Salt Lake City
DENNIS K. BENSON, Appropriate Solutions, Columbus, Ohio
JOHN E. EBEL, Boston College, Massachusetts
THOMAS C. HANKS, U.S. Geological Survey, Menlo Park, California
JAMES S. LANGER, University of California, Santa Barbara
NORMAN C. RASMUSSEN, Massachusetts Institute of Technology, Cambridge
LEON REITER, Nuclear Waste Technical Review Board, Arlington, Virginia
DANIELE VENEZIANO, Massachusetts Institute of Technology, Cambridge

National Research Council Staff

WILLIAM E. BENSON, Study Director
JUDITH L. ESTEP, Administrative Assistant

COMMITTEE ON SEISMOLOGY

THOMAS H. JORDAN, *Chairman*, Massachusetts Institute of Technology, Cambridge
WALTER J. ARABASZ,† University of Utah, Salt Lake City
F. A. DAHLEN,* Princeton University, New Jersey
STEVEN M. DAY, San Diego State University, California
THOMAS C. HANKS, U.S. Geological Survey, Menlo Park, California
CHARLES A.LANGSTON, Pennsylvania State University, University Park
THORNE LAY, University of California, Santa Cruz
STEWART A.LEVIN, Mobil Exploration & Production Technical Center, Dallas, Texas
STEPHEN D. MALONE, University of Washington, Seattle
DONALD L. PAUL,† Chevron Petroleum Technology Company, Houston, Texas
JAMES R. RICE, Harvard University, Cambridge, Massachusetts
PAUL G. SOMERVILLE, Woodward-Clyde Consultants, Pasadena, California
ANNE M. TREHU, Oregon State University, Corvallis
JOHN E. VIDALE, University of California, Los Angeles

National Research Council Staff

CHARLES MEADE, Senior Program Oficer
WILLIAM E. BENSON, Senior Program Oficer
JUDITH L. ESTEP, Administrative Assistant

*Term expired December 31, 1995.
†Term expired December 31, 1994.

BOARD ON EARTH SCIENCES AND RESOURCES

J. FREEMAN GILBERT, *Chairman*, University of California, San Diego
THURE CERLING, University of Utah, Salt Lake City
MARK P. CLOOS, University of Texas at Austin
JOEL DARMSTADTER, Resources for the Future, Washington, DC
KENNETH I. DAUGHERTY, E-Systems, Fairfax, Virginia
WILLIAM R. DICKINSON, University of Arizona, Tucson, *emeritus*
MARCO T. EINAUDI, Stanford University, California
NORMAN H. FOSTER, Independent Petroleum Geologist, Denver, Colorado
CHARLES G. GROAT, University of Texas, El Paso
DONALD C. HANEY, Kentucky Geological Survey, Lexington
SUSAN M. KIDWELL, University of Chicago, Illinois
SUSAN KIEFFER, Kieffer & Woo, Inc., Palgrave, Ontario
PHILIP E. LAMOREAUX, P.E. LaMoreaux and Associates, Inc., Tuscaloosa, Alabama
SUSAN M. LANDON, Thomasson Partner Associates, Denver, Colorado
J. BERNARD MINSTER, Scripps Institution of Oceanography, La Jolla, California
ALEXANDRA NAVROTSKY, Princeton University, New Jersey
JILL D. PASTERIS, Washington University, St. Louis, Missouri
EDWARD C. ROY, JR., Trinity University, San Antonio, Texas

National Research Council Staff

CRAIG M. SCHIFFRIES, Director
THOMAS M. USSELMAN, Associate Director
WILLIAM E. BENSON, Senior Program Officer
ANNE M. LINN, Senior Program Officer
CHARLES MEADE, Senior Program Officer
LALLY A. ANDERSON, Staff Associate
VERNA J. BOWEN, Administrative Assistant
JENNIFER T. ESTEP, Administrative Assistant
JUDITH ESTEP, Administrative Assistant

COMMISSION ON GEOSCIENCES, ENVIRONMENT, AND RESOURCES

GEORGE M. HORNBERGER, *Chairman*, University of Virginia, Charlottesville
PATRICK R. ATKINS, Aluminum Company of America, Pittsburgh, Pennsylvania
JAMES P. BRUCE, Canadian Climate Program Board, Ottawa, Ontario
WILLIAM L. FISHER, University of Texas at Austin
JERRY F. FRANKLIN, University of Washington, Seattle
DEBRA S. KNOPMAN, Progressive Foundation, Washington, DC
PERRY L. McCARTY, Stanford University, California
JUDITH E. McDOWELL, Woods Hole Oceanographic Institution, Massachusetts
S. GEORGE PHILANDER, Princeton University, New Jersey
RAYMOND A. PRICE, Queen's University at Kingston, Ontario
THOMAS C. SCHELLING, University of Maryland, College Park
ELLEN K. SILBERGELD, University of Maryland Medical School, Baltimore
VICTORIA J. TSCHINKEL, Landers and Parsons, Tallahassee, Florida

National Research Council Staff

STEPHEN RATTIEN, Executive Director
STEPHEN D. PARKER, Associate Executive Director
MORGAN GOPNIK, Assistant Executive Director
GREGORY SYMMES, Reports Officer
JAMES MALLORY, Administrative Officer
SANDI FITZPATRICK, Administrative Associate
MARQUITA SMITH, PC Analyst

The National Academy of Sciences is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Bruce Alberts is president of the National Academy of Sciences.

The National Academy of Engineering was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. William A. Wulf is interim president of the National Academy of Engineering.

The Institute of Medicine was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Kenneth I. Shine is president of the Institute of Medicine.

The National Research Council was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Bruce Alberts and Dr. William A. Wulf are chairman and interim vice-chairman, respectively, of the National Research Council.

Preface

In the 1980s two studies produced probabilistic seismic hazard estimates for nuclear power plant sites in the central and eastern United States. The first, sponsored by the U.S. Nuclear Regulatory Commission (USNRC), was conducted by Lawrence Livermore National Laboratory. The second, sponsored by utilities in the Seismicity Owners Group, was conducted by the Electric Power Research Institute (EPRI). The studies produced similar hazard curves and generally similar estimates of relative hazard. But for several sites absolute hazard levels differed by two or more orders of magnitude.

Because absolute hazard levels are important for nuclear power plant design, a new study, sponsored jointly by the USNRC, EPRI, and the U.S. Department of Energy, was undertaken by the newly formed Senior Seismic Hazard Analysis Committee (SSHAC) to determine the source of the major discrepancies in the two hazard estimates and to derive a robust probabilistic seismic hazards analysis methodology that could be used for future estimates.

At the same time, the USNRC asked the National Research Council (NRC) to review the work of the SSHAC study and evaluate the proposed methodology. This review was undertaken by the Panel on Seismic Hazard Evaluation of the NRC's Committee on Seismology which followed the work of the SSHAC study and produced the present critique of the SSHAC report.

Carl Kisslinger
Chairman

Contents

Executive Summary

This review and commentary by the National Research Council's Panel on Seismic Hazard Evaluation presents the panel's evaluation and critique of the report titled *Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts* (U.S. Nuclear Regulatory Commission, NUREG/CR-6372, Washington, DC, 1997). The reviewed report was prepared by the Senior Seismic Hazard Analysis Committee (SSHAC), a committee created and sponsored by the U.S. Nuclear Regulatory Commission (USNRC), the U.S. Department of Energy, and the Electric Power Research Institute. The panel was appointed at the request of the USNRC to provide an independent interactive review of the results of SSHAC's efforts.

SSHAC's charge from its sponsors' perspective was to provide an up-to-date procedure for obtaining reproducible results from the application of probabilistic seismic hazard analysis (PSHA) principles established in past practice, not to advance the foundations of PSHA or develop a new methodology. This focus led to an emphasis on procedures for eliciting and aggregating data and models for performing a hazard analysis, rather than an examination of the earth science foundations of PSHA. SSHAC focused on process because previous PSHA studies have shown that different groups of experts can produce highly discrepant results. A second major theme in the SSHAC report is the treatment of uncertainties in data and models in arriving at stable estimates of seismic hazard at a selected site.

With this in mind, the panel found that the SSHAC report offers substantial contributions to the foundations and practice of PSHA. In particular, the panel commends SSHAC for emphasizing the need for critical evaluation of expert opinion. But the panel also identified some limitations in both the report and the recommended procedures, of which

potential users should be aware. Only certain key points are highlighted here in the summary; the rest are included in later chapters.

MAIN FEATURES OF THE SSHAC REPORT

As stated above, the SSHAC report focuses on procedures for using experts in probabilistic seismic hazard analysis and for determining uncertainties at key stages of the analysis process. In its treatment of the use of expert opinion, SSHAC outlines four possible levels of effort and complexity. But the SSHAC report is strongly flavored by emphasis on hazard analysis for nuclear and other critical facilities, and SSHAC therefore discusses at great length its highest-level (level 4) procedure for evaluating expert opinion. And although SSHAC includes proper disclaimers the unwary reader could gain the incorrect impression that the high-level (level 4) PSHA procedure is needed for every hazard analysis.

The panel agrees that all PSHA projects should share the same basic principles and goals, but that the elaborate level 4 methodology is not required for every PSHA study. SSHAC does indeed recognize that alternate simpler methods are probably adequate for less critical facilities, but the simpler methods are not discussed in detail and the reader is not fully advised about other sources of information. Adequate disclaimers in the SSHAC report should protect the analyst who chooses to use procedures other than those recommended by SSHAC from the need to defend that decision in a regulatory setting.

THE SSHAC METHODOLOGY

SSHAC's contributions to PSHA methodology include the testing and full explication of the technical facilitator/integrator (TFI) entity, which is the essential ingredient in implementing SSHAC's high-level (level 4) analysis.[1] The TFI approach was found to be very effective in two workshops on ground motion estimation and led to an unexpected degree of agreement among the experts consulted, who began with many diverse viewpoints. The panel notes that TFI elicitation procedure is not

[1]For a description of the TFI entity, see Chapter 2.

synonymous with PSHA methodology. Nor is the TFI approach recommended by SSHAC for every PSHA study.

In outlining its four levels of complexity, SSHAC visualizes three distinct roles that experts should play at various stages of the process. First, an expert may start out as the proponent of a particular position (data or model). Then the expert is asked to become an objective evaluator of the positions of the other experts in the group. Finally, the expert becomes an integrator and aggregates all the positions to arrive at a putative position of the whole informed scientific community. This estimation of the position of the whole informed community by integration of the positions of a sample of well-qualified experts is the primary goal of the more complex SSHAC procedure. **The panel questions whether any group of experts can truly assess the view of the whole informed scientific community on the entire range of relevant issues.**

BACKGROUND WORKSHOPS

SSHAC sponsored workshops on seismic source characterization, ground motion estimation, and earthquake magnitudes. These workshops are documented in detail in Appendixes A, B, C, and H of the SSHAC report. The workshops contributed both to the development of the procedures SSHAC recommends and to advancement of our knowledge of the earth science elements of PSHA for the eastern United States. Because SSHAC focused on procedures for PSHA rather than technical issues, some of these valuable results are presented but not highlighted. They deserve more attention.

THE TREATMENT OF UNCERTAINTY

The SSHAC report emphasizes the importance of how uncertainty is treated because the results of a PSHA can be influenced heavily by uncertainties in the data, the models, or both. SSHAC's treatment distinguishes and emphasizes the difference between two types of uncertainty: aleatory (i.e., uncertainty due to variability inherent in the phenomenon under consideration) and epistemic (uncertainty due to our limited knowledge of the phenomenon). After separation, these two

components must be quantified for the model or parameter under consideration. The panel has more trouble with this element than any other in the SSHAC report.

Recognition of the two kinds of uncertainty is useful initially when eliciting and combining expert inputs. Experts need to be aware of the sources of uncertainties (e.g., limitations of available data) so that they can make informed assessments of the validity of alternative hypotheses, the accuracy of alternative models, and the value of data and then transmit those uncertainties to the TFI. However, as detailed in Chapter 3 of this report, the panel believes that the statistical analysis and uncertainty separation procedures recommended by SSHAC may in some cases be more sophisticated than is warranted by the data or the purposes for which the results are to be used.

During the planning of a PSHA, a detailed analysis of uncertainty would be helpful but typically is not available. It may be sufficient for planning purposes to conduct limited sensitivity analyses, using bounding hypotheses, and to consider the level of effort that would be required to reduce the associated uncertainty.

In addition, the value of an epistemic/aleatory separation to the ultimate user of a PSHA is doubtful. In particular, it is not clear that such a separation would be more helpful than the display of expert-to-expert variability of a mean hazard at the time of an analysis, with an explanation of the source of the differences.

The panel also notes that the SSHAC report's discussions and recommendations on uncertainty and the use of experts are quite independent of PSHA and can be applied to other types of risk analysis. The panel believes that the SSHAC report makes a solid contribution to the methodology of hazard analysis, especially in the use of expert opinion.

1
Introduction

> "The future utility of PSHA in decision making depends to a large degree on our ability to implement the process in a meaningful and cost-effective way. Development of the SSHAC guidelines was planned with this goal in mind."
>
> *—from Sponsors' Perspective, SSHAC Report*

This review and commentary by the National Research Council's Panel on Seismic Hazard Evaluation presents the panel's evaluation of the report *Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts* (U.S. Nuclear Regulatory Commission, NUREG/CR-6372, Washington, DC, 1997). That report was prepared by the Senior Seismic Hazard Analysis Committee (SSHAC) (not a committee of the National Research Council) with sponsorship and oversight by the U.S. Nuclear Regulatory Commission (USNRC), the U.S. Department of Energy (DOE), and the Electric Power Research Institute (EPRI).

WHAT IS SEISMIC HAZARD ANALYSIS?

Earthquakes present a threat to people and the facilities they design and build. Seismic hazard analysis (SHA) is the evaluation of potentially damaging earthquake-related phenomena to which a facility may be subjected during its useful lifetime. An SHA is done for some practical purpose, typically seismic-resistant design or retrofitting. Although strong vibratory ground motion is not the only hazardous effect of earthquakes (landslides, fault offsets, and liquefaction are others), it is the cause of much widespread damage and is the measure of earthquake hazard that has been accepted as most significant for hazard resistance planning.

The level of effort put into an SHA depends on the investment in the facility that might be lost and the consequences to society should it fail. Critical facilities are those that are deemed so important to the functioning of society or whose catastrophic failure will have such disastrous consequences that a maximum (and necessarily costly) effort to assess seismic and all other natural hazards is justified. The SSHAC project was born in the context of SHA for such critical facilities, nuclear power plants in particular. Even though SSHAC broadened its concept of the applicability of its recommended approach to SHA, its report is strongly influenced by this orientation toward very large, costly facilities for which the end goal is to prevent catastrophic failure, even at great expense.

Two general approaches to SHA have been developed and applied. The first approach uses discrete, single-valued events to arrive at scenario-like descriptions of the hazard. Typically, a seismic source location, a maximum earthquake associated with that source, and a ground motion attenuation relationship are specified. The ground motion at the site of interest implied by the chosen inputs is then calculated. The frequency of earthquake occurrence is usually not taken into account, and there is no formal and open way of treating uncertainties. This approach has been labeled deterministic seismic hazard analysis (DSHA) and has been used for many years in the design of power plants, large dams, and other critical facilities.

The other approach is probabilistic seismic hazard analysis (PSHA) and is the subject of the SSHAC effort. PSHA allows the use of multivalued or continuous events and models incorporating the effects and frequencies of all earthquakes that could impact a site. PSHA can easily incorporate model and parameter uncertainties. The results of a PSHA, including the uncertainties, can be represented as a series of curves (mean, median, or selected fractiles), showing the annual frequency of exceeding different levels of the chosen measure of ground motion. The intent of high-level PSHA is to capture and display as much as possible of the knowledge provided by existing data, theory, and computational simulations.

It should be noted that the procedures recommended by SSHAC for the elicitation and aggregation of expert opinion as input to PSHA are equally applicable for compiling the input for DSHA. The only essential difference between DSHA and PSHA is that the latter carries units of time while the former usually does not (Hanks and Cornell, 1994). In the case of a specific design situation, both DSHA and PSHA result in estimates of

ground motion values or time histories that provide the basis for earthquake-resistant design. PSHA yields, in addition, the annual frequency of exceedance of that ground motion level together with attendant uncertainties. SSHAC's responsibilities did not extend to a discussion of the steps by which project engineers and sponsors use the output of a hazard assessment. One approach to this issue is presented in a recent paper by McGuire (1995).

Projection of the location, severity, and frequency of occurrence of future extreme natural events inherently involves a variety of uncertainties. Yet decisions on the siting and design of needed facilities must be made in the face of these uncertainties. No amount of statistical analysis, no matter how rigorously based and carefully done, can totally compensate for the incompleteness of available data and the defects of our evolving scientific knowledge. **A primary objective of SSHAC was to acknowledge and document uncertainties explicitly so that users of PSHA will be able to make better-informed decisions.**

BACKGROUND AND CONTEXT OF THIS REPORT

The Panel on Seismic Hazard Evaluation was created under the Committee on Seismology of the National Research Council in October 1992. The panel was formed in response to a request from the USNRC to provide an independent review and evaluation of a report on PSHA to be produced by SSHAC.

The work of the panel was influenced by several factors. First, the USNRC asked the panel to provide an "interactive review," that is, to submit feedback to SSHAC as it worked in order to avoid the production by SSHAC of a report in which the panel might find serious flaws after it was completed. This request raised serious questions as to how the panel could meet its requirement and not become so involved in the production of the SSHAC report that the objectivity of the panel's own review would be compromised. The panel agreed with the USNRC to provide "arms-length" interaction with SSHAC and developed methods of operation to achieve that goal.

Another factor affecting the work of the panel was a change in the charge to SSHAC after it began its work. The original task assigned by the sponsors concentrated on the reconciliation of two studies done in the mid-1980s by Lawrence Livermore National Laboratory (LLNL) and EPRI of the earthquake hazard at nuclear power plant sites in the United

States east of the Rocky Mountains. These studies were prompted by advice to the USNRC from the U.S. Geological Survey, based on its reconsideration of the likelihood that a major earthquake, such as the Charleston, S.C. earthquake of 1886, could occur again in Charleston or elsewhere along the eastern seaboard. The possibility of such an earthquake could have implications for the safety of nuclear power plants in the eastern United States. A brief history of the LLNL and EPRI studies is given in the SSHAC report.

Although the two studies ranked the many sites approximately the same (from most hazardous to least hazardous in terms of the mean hazard estimates), the absolute hazard values for specific sites, in terms of the mean value of the annual probability of exceeding a specified level of ground motion, differed greatly, with the LLNL results consistently greater.

The problem is illustrated in Figure 1.1, which displays the hazard at three widely separated sites as the annual frequency of occurrence of peak ground acceleration (PGA), the ground motion parameter chosen for this evaluation. The median hazard curve from each study is shown, as well as the 85th and 15th percentile curves. In two of the three cases shown, the median hazard calculated by LLNL is well above that derived by EPRI, and the "uncertainty," measured by the spread of the 15th and 85th percentile curves, is much greater for LLNL than EPRI. Also, the uncertainty is large, a factor of 5 or more at potentially damaging levels of ground motion (PGA greater than 200 cm/sec^2).

The mean hazard curves, not shown in the figure, differ by even greater factors in many cases. This is because the LLNL median and 85th percentile curves are above the EPRI results, and arithmetic averages spanning several orders of magnitude give greatest weight to the largest numbers. This explains the relatively high values of the mean hazard derived by LLNL but it does not get at the fundamental cause for the differences in the estimates.

The desirability of discovering the cause(s) of the discrepancies was obvious, not only for intellectual reasons (why did competent scientists working from the same or similar knowledge and data bases get vastly different answers?), but also for the practical reason that the quantitative estimate of seismic hazard is important in judging whether earthquakes represent a substantial threat, as well as the weight of earthquakes relative to other natural hazards in making design and retrofitting decisions. The USNRC funded LLNL to investigate the

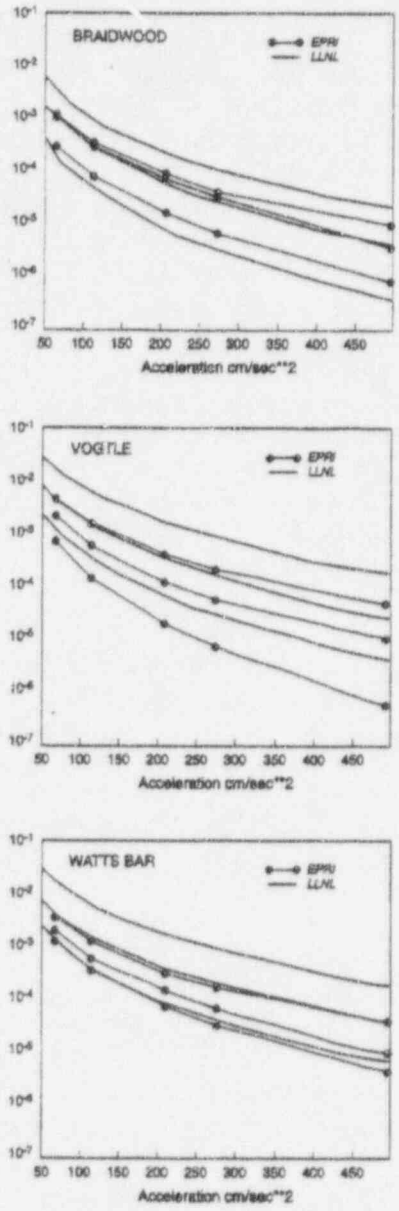

FIGURE 1.1 Median, 15th, and 85th percentile hazard curves for three representative separated sites in the eastern United States, illustrating the differences in results of the LLNL and EPRI studies. The ordinate is the estimated annual frequency of exceedance of the peak ground acceleration shown as the abscissae (adapted from Figures 2.3.1, 2.3.7, 2.3.8 in Bernreuter et al., 1987).

problem. LLNL's study (Bernreuter et al., 1987) concluded that the factors involved in the discrepancy were: (1) different values were chosen for the lower-bound earthquake when the groups were integrated over seismicity to calculate the hazard, (2) different ground motion models were used, and (3) LLNL included a correction for local site effects and EPRI did not. This explained why the two studies obtained different answers but does not explain why competent analysts arrived at significantly different inputs to the hazard calculations.

As SSHAC was being assembled, the underlying cause of the discrepancies between the two studies was identified by further study at LLNL. Researchers there concluded that the differences were due to the ways in which the inputs provided by experts had been elicited. Once this was recognized and taken into account, the differences in the outputs (mean hazard curves) were reduced from orders of magnitude to small factors that represented satisfactory agreement, given the many uncertainties in every step of the analysis. This resolution of the original problem led to changes in the SSHAC charter (1994), from which the following items are selectively cited to provide the context within which the SSHAC report was developed:

> *Objective:* To develop implementation guidelines, including recommended methodology, suitable for the performance of PSHA for seismic regulation of nuclear power plants and other critical facilities.
>
> *Requirements and Guidelines* (for the implementation guidelines and methodology):
>
> - Be able to provide probabilistic seismic hazard results in the form of fractile probabilities and mean values over a range of ground motion levels suitable for use in probabilistic seismic risk assessments for nuclear facilities.
> - Be defined in sufficient detail that, when independently applied by different organizations, no ambiguity exists on how the PSHA is to be performed and comparable results are obtained.
> - It is specifically not the objective of this program to advance PSHA methodology or to develop a new PSHA methodology. Rather, an important step in reaching the

objective of this program is expected to be the completion of evaluations of independent PSHA applications by LLNL and EPRI as well as other relevant applications.

- The outcome of this process will be the recommended methodology and implementation guidelines for PSHA in nuclear power plant licensing.

The emphasis on methodology for doing PSHA as the central theme is reflected in the title of the SSHAC report. The focus on siting nuclear facilities, though not emphasized explicitly in the report, strongly influenced its concentration on high-level PSHA.

It should be recognized that the charges to SSHAC and to the panel did not call for the defense or promotion of PSHA as a method for evaluating earthquake hazards. **SSHAC has produced a document that sets forth its conclusions and recommendations on the proper way to do a PSHA if that is the approach chosen by project developers and their analysts.** Neither the SSHAC report nor the panel evaluates the efficacy of PSHA relative to other methods, DSHA in particular. The SSHAC report does provide criteria that can be used to decide the appropriate level of effort for a specific study. Some of the issues related to alternatives to a full-blown PSHA and alternatives to SSHAC's recommended procedures are discussed elsewhere in this report.

The panel offers its appraisal of the SSHAC report, with primary emphasis on the scientific validity of the work and its conclusions, with appropriate attention to the clarity of the presentation, possible sources of misinterpretation, and the report's contributions to PSHA.

INTERACTIONS OF THE PANEL WITH SSHAC

The panel met with SSHAC three times (June 28-29, 1993; May 27-28, 1994; and December 9-10, 1994). Members of SSHAC, representatives of the three sponsoring organizations, and scientific and technical consultants to SSHAC attended the meetings. In addition, Thomas Hanks, a member of the panel, attended a number of SSHAC meetings as liaison observer.

By the nature of its charge, the panel was not able to begin its work until it received a draft product from SSHAC and could not finish its work until it had received the complete final SSHAC report. The June

1993 meeting was devoted primarily to briefings by agency representatives, SSHAC members, and scientific consultants, designed to educate the panel about the goals of SSHAC, the background of the problems being addressed, and the procedures SSHAC would follow. A spokesman for the USNRC explained that the agency wanted two products from SSHAC: (1) a set of guidelines for the process of seismic hazard assessment, and (2) a set of guidelines for the agency, using current data sets and computer codes, to reevaluate the hazards at existing sites. A SSHAC spokesman concluded that the central thrust of the project was to develop, justify, and illustrate methods for capturing both the inherent uncertainties in the parameters that go into an analysis and the disagreement among experts about the values of these parameters. At this time, the panel decided that it needed two additional members, one who could provide expertise in expert opinion analysis and decision science and one with extensive knowledge of both the deterministic and probabilistic approaches to seismic hazard assessment.

By May 1994 the focus of the SSHAC effort had changed, as noted above, from the reconciliation task to the more substantial and significant task of building on the lessons learned from prior experience in hazard assessment to develop scientifically sound procedures for doing PSHA. The SSHAC chairman explained that his committee's goal had been broadened to the development of a methodology that would be applicable not only to nuclear power plants but to other critical facilities as well. SSHAC members presented detailed technical briefings in their areas of expertise, so that the panel gained insight into the flavor of the report that SSHAC would produce. Vigorous discussions of both earth science and decision science issues provided a forum for the panel to explore details of the proposed SSHAC approaches and to convey in broad terms some concerns of the panel. Points raised in these discussions and the panel's evaluation of how SSHAC treated each are addressed elsewhere in this report.

The December 9-10, 1994, panel meeting was based on a detailed review of a draft report submitted by SSHAC. The draft was incomplete; in particular, the extensive appendixes, which on later examination proved to be essential and very valuable contributions of the SSHAC effort, were not available. But, the panel did conduct a detailed review of the main report. SSHAC members, as well as the agency representatives, were present for this review. The results of the review were submitted in the form of a formal letter report to the USNRC on March 16, 1995 (reproduced here as Appendix B). The USNRC forwarded this letter

report to SSHAC as part of its oversight of the final version of the SSHAC report.

The March 1995 letter report was the principal formal feedback from the panel to SSHAC. The letter report offered the panel's general comments on the SSHAC draft, a statement of concerns and problems, with suggestions for improvement, and a summary of specific scientific and technical concerns that the panel thought should be addressed. A draft of the final SSHAC report was sent to the panel on October 6, 1995. The present report is based on the panel's review of the October 6 draft, supplemented by several figures and parts of the appendixes that were submitted later. (Although the October 6 draft needed editing the panel was informed that the work of SSHAC was completed and that no further substantive changes in the SSHAC report would be made.)

The expectations of the sponsoring organizations are expressed succinctly in the last sentence of the *Sponsors' Perspective* that opens the SSHAC report, which is quoted at the beginning of this chapter. The panel has reviewed and evaluated the SSHAC report in light of these expectations and how well the goal has been achieved.

ORGANIZATION OF THE PANEL'S REPORT

The panel determined that the SSHAC report could be reviewed under four main headings: (1) process (elicitation and aggregation) and documentation, (2) the treatment of uncertainty, (3) seismic source characterization, and (4) ground motion estimation. The first two concentrate on the decision science components of PSHA, the latter two on the earth science inputs. Following a chapter on each of these, the panel offers a summary of its findings and recommendations.

2
Process and Documentation for a Probabilistic Seismic Hazard Analysis

By its own definition, the main emphasis of the Senior Seismic Hazard Analysis Committee's (SSHAC) report is on the procedural rather than the technical aspects of probabilistic seismic hazard analysis (PSHA). SSHAC argues that many of the major potential pitfalls of PSHA are procedural and therefore goes to great efforts to outline what it views as an appropriate process. In SSHAC's view the important aspects of "process" have to do primarily with experts, their interaction, and methods for translating their views into useful input for a PSHA. Of particular significance is the role assigned to the facilitation/integration team that organizes and directs a PSHA project and its use of experts. SSHAC lays out two basic principles underlying the PSHA process and its results:

1. Regardless of the scale of a PSHA study, the goal (as stated by SSHAC) is "*to represent the center, the body, and the range of technical interpretations that the larger technical community would have if they were to conduct the study.*"
2. "*It is absolutely necessary that there be a clear definition of ownership of the inputs into the PSHA, and hence ownership of the results of the PSHA.*"

The panel supports these principles as ideological guidelines for planning and executing a PSHA study, at least in the case of critical facilities. The first is, or should be, the goal of a sponsor in initiating a PSHA, the assumption being that using the collective input of the *informed* technical community would be the best, and most defensible, way of defining seismic hazard. That principle also has an enabling effect because, as discussed later, it allows experts to transcend the role of being proponents of models (the usual mode in scientific discourse) into the

roles of objective evaluators and integrators. **The extent to which this goal can reasonably be pursued in a particular case should depend on the scope and importance of the project and the resources available to support the study.**

The second principle is important because it assigns to an identified entity, the "owner," clear intellectual or scientific responsibility for the conduct and results of a PSHA. This does not necessarily mean that the "owner" agrees with every particular input or result but that the owner feels confident that the PSHA has fulfilled the purpose of representing the larger technical community and can be defended in scientific and regulatory arenas, as necessary. These principles underlie the primary recommendations of the SSHAC report that deal with the PSHA process.

LEVEL OF EFFORT IN A PSHA

SSHAC recognizes that a PSHA can be carried out at different levels of effort and emphasizes that the effort expended should match the importance of the facility, the degree of controversy, uncertainty, and complexity associated with the relevant scientific issues, and external decision factors, such as regulatory concerns and the resources available. This is shown in Table 2.1, taken from Chapter 3 of the SSHAC report.

Four levels of study are defined, the first three of which rely on a single entity called the technical integrator (TI), who is responsible for all aspects of the PSHA, including specifying the input. Although experts may be involved on a consulting basis, there is no formal elicitation of their views. The highest level of study (level 4) makes use of formally elicited expert judgment. As such, a new entity called the technical facilitator/integrator (TFI) is needed. The role of the TFI is discussed below. A large part of the SSHAC report is devoted to defining what is necessary to carry out a level 4 study and explaining the function of the TFI because the ideas are new, not because this level of effort is required for every seismic hazard assessment. **It would be inappropriate to infer that all PSHAs require the considerable resources needed to carry out the level 4 PSHA described by SSHAC.**[2]

[2]Nor does SSHAC make such a claim or inference. This statement is more a caveat to users than a criticism of SSHAC.

The Panel endorses the conceptual framework embodied in Table 2.1, recognizing that the application of PSHA to engineering and regulatory problems is varied and that the level of effort needed should also vary.

SSHAC points out that most site-specific studies make use of some type of TI approach. The TI performs analyses, accumulates information relevant to each issue, and develops a representation of the technical community's views on the relevant input models, parameters, and their uncertainties. At the lowest level of effort (level 1) the technical community's views are determined primarily by a literature search. At higher levels the TI makes use of outside technical researchers and proponents to gain insight into different data sets and models.

The panel emphasizes that a TI must still be guided by the principles of representation and ownership described above.

The importance of peer review is discussed below, but the panel stresses its particular significance when the TI mode is used. Reliance on a single entity (TI) to characterize the input of the whole technical community may be a very efficient mode of operation, but additional assurance is needed to provide confidence that the results are a reasonable representation of the community's views.

THE MULTIPLE ROLES OF EXPERTS

The TFI process views experts as acting in different roles—proponents, evaluators, and integrators. The proponent role is one in which the expert explains, and argues for, the choice of a particular model or set of parameters. The aim is to make sure that the different views in the technical community are presented and discussed by the expert panel. If necessary, individuals outside the expert panel may be brought in to argue points of view with which panel members may not be comfortable. The next role the experts are asked to assume is that of independent evaluators representing their own views of the information presented. Mean estimates of model, component, or parameter values are elicited, along with their uncertainties as appropriate. The result should be the group's composite views of the issues at hand. The experts are encouraged to evaluate their own and other models according to their own technical judgment, without regard to who originally proposed the models. In the past, most PSHAs that have relied on formally elicited expert judgment have strived to get experts to think in this manner. The hope was that the experts' composite view also represented the composite view of the technical community as a whole.

TABLE 2.1 Degrees of PSHA Issues and Levels of Study (Table 3-1 of the SSHAC Report)

Issue Degree	Decision Factors	Study Level
A Non-controversial; and/or insignificant to hazard	• Regulatory concern • Resources available • Public perception	1 TI evaluates/weights models based on literature review and experience; estimates community distribution
B Significant uncertainty and diversity; controversial; and complex		2 TI interacts with proponents & resource experts to identify issues and interpretations; estimates community distribution
C Highly contentious; significant to hazard; and highly complex		3 TI brings together proponents & resource experts for debate and interaction; TI focuses debate and evaluates alternative interpretations; estimates community distribution
		4 TFI organizes panel of experts to interpret and evaluate; focuses discussions; avoids inappropriate behavior on part of evaluators; draws picture of evaluators' estimate of the community's composite distribution; has ultimate responsibility for project

To more truly represent the technical community's view, the SSHAC report recommends that the experts be specifically asked to assume the role of integrators and to characterize their perception of how the *technical community as a whole* would view the issues at hand. Thus, although the expert may view his/her assessment as being the most correct, he/she is explicitly thrust into the role of trying to fulfill the first principle of PSHA as outlined above and must be willing to do so. This mode of expert behavior may not be achievable in all issues. Also, the

panel is not aware of any objective way to test the assumption that a whole technical community's views can be accurately determined from the interactions of a small group of experts.

SSHAC introduces some useful concepts in its discussion of the interaction among experts. One is that in the process of eliciting, aggregating, evaluating, and integrating the opinions of experts the TFI (discussed in the next section) should create an atmosphere in which there will not be "winners" and "losers." Another useful idea is the avoidance of unintended dissent or consensus. Apparent disagreement may arise because of lack of communication and understanding among these disagreeing; the process of "active listening," in which a listener is asked to give back what he/she has just heard, is a step toward eliminating disagreement where it really does not exist. At the other extreme is the development of an apparent but false consensus; the TFI should strive for consensus among the experts only if it is really agreed on.

The panel views the role of expert as integrator as important and worthwhile. However, successful implementation of the integrator role of the experts should be viewed more as a goal to strive for than a uniformly and demonstrably achieved measure of success.

The SSHAC report implies four basic criteria for the identification and selection of experts: (1) technical expertise, (2) strong communication skills, (3) willingness to assume the role of independent evaluator, and (4) willingness to commit the time and effort to participate actively in the study. The choice of disciplines to be represented and the breadth of knowledge of each expert depend on the issues to be addressed and whether or not interdisciplinary subgroups of experts will be formed to provide input. SSHAC also strongly recommends a formal nomination process based on consulting the literature and asking technic societies, government organizations, and knowledgeable individuals to submit the names of potential experts. **Whatever the issue or structure of elicitation, the panel believes that the credibility and quality of an elicitation-based PSHA depend very much on the choice of experts. The panel supports the need for careful attention to the selection process and finds the criteria suggested by SSHAC to be reasonable and likely to be effective.**

TECHNICAL FACILITATOR/INTEGRATOR

One of SSHAC's main contributions to PSHA methodology is the introduction of the technical facilitator/integrator (TFI) concept. The SSHAC report describes this new function in Section 3.3.1 as follows:

> The TFI is a single entity who has the responsibility and is empowered to represent the composite state of information regarding a technical issue of the scientific community.... The TFI process is centered on the precept of thorough and well-documented expert interaction as the principal mechanism for integration.

As SSHAC acknowledges, a major stimulus for its charge was the need to resolve the differences in hazard estimates between the Lawrence Livermore National Laboratory and the Electric Power Research Institute studies. SSHAC's investigation revealed that the process of elicitation and the procedures for integration allowed room for considerable misunderstanding and potential misinterpretation. Six areas in which improvements could lead to a better outcome are detailed in Section 3.3.2.2 of the SSHAC report:

1. Overly diffused responsibility
2. Insufficient face-to-face expert interaction
3. Inflexible aggregation schemes
4. Imprecise or overly narrow objectives
5. Outlier experts
6. Insufficient feedback

The TFI concept was designed to resolve these procedural issues. This approach is described in detail in Chapters 3 through 5 and Appendix J of the SSHAC report. **The panel concurs that, in cases in which decisions about a critical facility of major complexity depend on controversial and uncertain inputs, the TFI approach offers an effective mechanism for capturing the best of what is known about the particular issues.**

The Proposed TFI Process

The seven steps proposed by SSHAC for the TFI approach (Section 3.3.4) were first suggested by Keeney and von Winterfeldt (1991), based on their experience in eliciting expert judgment for probabilistic risk assessment of nuclear power plants. The steps are:

1. Identification and selection of technical issues
2. Identification and selection of experts
3. Discussion and refinement of technical issues
4. Training for elicitation
5. Group interaction and individual elicitation
6. Analysis, aggregation, and resolution of disagreements
7. Documentation and communication

A flow chart of the process as applied to ground motion elicitation by SSHAC is reproduced here as Figure 2.1. Appendix J of the SSHAC report spells out the background, evolution, and details of the TFI process as developed by SSHAC. Appendix J must be read carefully; readers may need to consult additional references in order to fully understand some of the issues discussed, such as the weighting of individual expert inputs.

The TFI process requires careful and time-consuming setup procedures to ensure that all participants are clear on the objectives of the study, their roles in the study, and the intended results. The TFI (an individual or, perhaps, a team of two or three people) must be highly competent in the relevant subject areas, adept at elicitation and group process, and thorough. Because a strong TFI will have a major influence on the outcome of the elicitation/aggregation process, **it is essential that, if more than one TFI is assigned to work on a particular analysis project, they all be equally well qualified.**

The panel concludes that for appropriate issues the TFI process holds significant promise for PSHA. This process was developed by SSHAC as part of its effort to overcome limitations of previous PSHA studies. The panel cautions, however, that this process is expensive, time consuming, and demanding of all participants. SSHAC's criteria for identifying the issues for which the full TFI process is justified (Table 2.1) must be understood by project sponsors and their analysts.

As discussed in the next chapter, each element of a seismic hazard analysis may involve high degrees of uncertainty. Many situations arise in

which competent experts may legitimately disagree in their interpretation of extant data and theory. In view of the complexity of the issues and models involved in PSHA, SSHAC concluded that an improvement in the process of elicitation would help focus attention on the technical issues by reducing previously observed problems in "consensus," unintended agreement, and unintended disagreement.

At each step of the elicitation process, the TFI strives for complete understanding by each expert of all technical issues. The goal is that all experts are "on the same page." The results of two ground motion workshops conducted by SSHAC and documented in Appendixes A and B of its report indicate that investment in the TFI process bore substantial results.

The panel is aware that the TFI process, as implemented in these workshops, has rarely been used in the earth sciences. An example of the application of the process in a related subject field is provided by a probabilistic volcanic hazards analysis (Coppersmith et al., 1995).

TREATMENT OF EXPERT INPUT

Integration of Expert Opinion

SSHAC correctly points out that in theory it is always possible to formulate the expert integration problem as a Bayesian inference problem in which the opinions rendered by the experts are viewed as "noisy observations" of the quantities of interest (e.g., parameter values, distributions). Difficulties lie in the formulation of an "observation model" tailored to each expert combination task and sometimes in implementing the Bayesian analysis to produce *a posteriori* uncertainties. A discussion of combination problems and models is given in Appendix J of the SSHAC report. SSHAC repeatedly warns against blindly using any specific model and stresses that the models described in Appendix J are only examples for illustration. The panel agrees with these warnings and adds the following comments:

- In essence, Appendix J presents two very different types of models: (1) the so-called classical models, which emphasize the "noisy observation" interpretation of expert opinion, and (2) the TFI model, which regards each expert as being potentially correct, with a probability

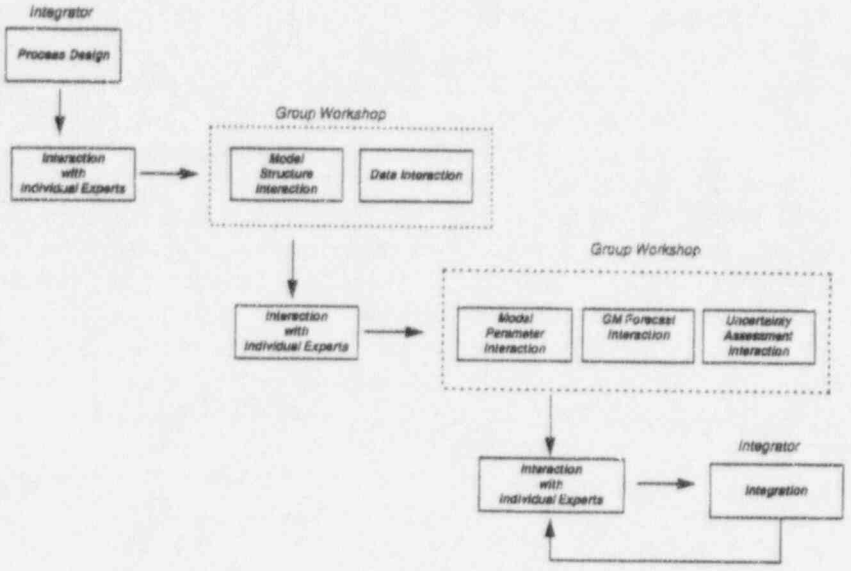

FIGURE 2.1 Roadmap of ground motion elicitation process (Figure 5-5 of the SSHAC report).

proportional to an assigned weight. Although this interpretation of the TFI model is not given in the SSHAC report, the fact that the community distribution is defined as a weighted sum of the expert distributions is equivalent to saying that each expert is correct with a probability equal to his/her assigned weight. At the end of Appendix J, the two approaches are compared numerically and shown to produce very different results. Without an in-depth discussion of when each type of model (or neither) is applicable, Appendix J may leave the reader confused. The classical models combine distribution functions with the meaning of uncertainty on the value of an unknown parameter. Hence, in this case the object of estimation is an unknown *scalar quantity* and the distributions express uncertainty on that quantity according to different experts. The TFI model, on the other hand, combines distribution functions that express the state of uncertainty of the scientific community according to different experts. In this second case the object of estimation is the *distribution function* itself. Therefore, while the inputs to, and results from, both

models are in the form of probability distributions, such distributions have different meanings in the two cases and should not be compared.

- The community distribution, which the TFI model estimates, is defined in Appendix J, Section 5, of the SSHAC report as "the mixture of the distributions of the individual experts if [the decision maker] believed that the experts . . . in this 'perfect community' were effectively equally informed on the issue of interest and equally interdependent. . . ." **As the entire SSHAC procedure revolves around this distribution, the panel believes that its definition should have been given in the main report, with a detailed explanation and justification.**
- SSHAC gives expressions for the mean and variance of the community distribution after stages 1 and 2 of the TFI process. Given the approximate nature of the results for the variance and the fact that distributions, not just mean values and variances, are needed, a much simpler and basically as accurate combination rule would be to take the weighted average of the distributions provided by the experts. The statement in Appendix J that "determination of the predictive (i.e., *a posteriori*) distribution follows a straightforward but cumbersome Bayesian statistical analysis" indicates that SSHAC knows how to perform a fully nonparametric Bayesian estimation of the community distribution function. This panel could think of no straightforward procedure to do so (one would need to consider the expert distribution estimates as random processes given the true community distribution function, with serious practical and conceptual implications). Because determination of community distribution and its uncertainty is at the core of the SSHAC approach, the report should have been more explicit about such a procedure.
- SSHAC favors an equal weighting integration scheme, unless there are clear indications that different weights should be used, for example, to reduce the influence of outliers. Linear combination rules with equal (unequal if necessary) weights are applied to parameter estimates (classical models) as well as to the probability distributions that, according to the panel of experts, quantify uncertainty in the scientific community (TFI model). Conditions for "equal weights" are set forth in the report. The panel believes that there may be some confusion about linear combination with equal weights and symmetrical (but possibly nonlinear) treatment of the expert assessments. The conditions quoted in the SSHAC report apparently lead to symmetrical treatment, not necessarily to averaging. There is a brief reference to nonlinear

combination rules in the section on nonequal weights in Appendix J, with little discussion. Analysts are advised to verify whether the conditions of linearity and normality of the observation model apply before using a linear combination rule. Contrary to what SSHAC states (e.g., Figure J-6), in some cases it would be better to combine the parameters of the distributions provided by the experts rather than the distributions themselves (combining the parameters results in a *nonlinear* combination of the distributions.) For example, if the experts agree on all distribution characteristics except for a location parameter, combining the estimated locations would be the right thing to do.

In view of these limitations and the objective difficulties in properly combining expert opinions, the panel recommends the following:

1. Use the models in Appendix J of the SSHAC report for reference, not as prescriptive or even recommended combination procedures.

2. Do not accept the results of a mechanical combination rule unless they are consistent with judgment.

3. If a mechanical combination rule is used, a general way by which to derive that rule is to view experts as noisy observers of the quantity being estimated. This approach is always the correct one from a Bayesian viewpoint, irrespective of the problem at hand. What differs in different cases is the nature of the observation errors, which need not necessarily be normal, additive, or independent.

4. When combining expert opinions on distribution functions, the correct Bayesian approach requires the use of a random process formalism, unless the problem can be reduced to a discrete one through appropriate parameterization. In all but the simplest cases a formal analysis becomes prohibitive, and the panel recommends primary reliance on judgmental combination procedures.

Weighting

One of the more problematic aspects of PSHA has always been the aggregation of input from different experts, especially when one or more expert opinions are outliers relative to the views of the rest of the

participants. This problem has led to consideration of weighting of different experts' opinions based on quantitative or qualitative assessments of the degree of expertise (typically a highly subjective exercise). The extensive interactive education and elicitation process proposed by SSHAC is intended to bring all expert participants to parity. This process should make it more reasonable to use equal weighting of all the experts. Appendix B of the SSHAC report states that equal weights were used for the combination of expert opinions and concludes that the TFI "integration process is robust."

The panel concurs that equal weighting of experts should be the clearly preferred target in a multiple-expert PSHA. To achieve this, proper choice of experts and group interactions should be emphasized, as outlined in Chapter 4 and Appendix H of the SSHAC report. In the case in which a different weighting scheme is applied, the burden of proof rests with the TFI; nevertheless, every effort should be made to obtain expert concurrence on the weights used or modification applied.

Dependency Among Experts

A related aggregation problem, dependency among experts, is, on the surface, exacerbated by the TFI process. The overall community is composed of a finite number of experts who rely on a finite number of models and methodologies. While one or more of the participating experts may not be thoroughly familiar with the entire range of such models and methodologies at the beginning of the exercise, such familiarity is an objective of the TFI process. As shown in the second SSHAC ground motion workshop, this interactive process narrowed the range of estimates as the experts increased their knowledge and understanding of issues and methods. One goal of a well-executed TFI process is that all participating experts are better able to make informed independent judgments.

Peer Review

SSHAC requires that peer review be an integral part of the PSHA process. The panel concurs. SSHAC defines two types of review: (1) participatory and (2) late stage. Participatory peer review involves "full

and frequent access throughout the entire project" by the reviewers. The advantage of a participatory review is the opportunity to subject interim results and deliberations to independent feedback. This provides the PSHA team with an opportunity for adjustment and limits the possibility that a lengthy and costly effort might be found to have serious flaws in the end. SSHAC recognizes that a limitation of participatory peer review is that "peer reviewers might lose their objectivity as they interact with the project over time." The panel views a participatory peer review as equivalent to a backup group of experts who provide oversight of the work of the primary team. **Safeguards must be established to preserve the objectivity of the review process.** As explained in the introduction to this report, this panel was asked to provide participatory peer review to SSHAC, and the panel insisted on a process by which it would not become so deeply involved in the preparation of its report that its objectivity would be compromised. The panel believes that this is also a necessary precaution for peer review of any PSHA study.

The late-stage review is closer to the traditional academic review in that it occurs near the end of a project. SSHAC strongly recommends participatory peer review on the grounds that a late-stage review can be risky, especially with regard to the process aspects of a PSHA study. Table 3-2 in the SSHAC report summarizes its recommendations on how to structure the peer review process.

The panel concludes that participatory review, as part of a PSHA process, would serve to improve the quality of a study insofar as it is another step toward incorporating the views of the broad informed scientific community. Other considerations—for example, the requirements of regulatory bodies—might call for a late-stage review also.

Documentation

Chapter 7 of the SSHAC report puts much emphasis on the importance of fully documenting every PSHA study. The guidelines on documentation are intended to ensure that each step of the PSHA process is not only completely recorded but also that the records are stored in accessible formats that permit the technical community to review all operations and decisions. This documentation also greatly facilitates later reanalysis and update as new information becomes available, perhaps eliminating the necessity of redoing the entire PSHA.

The panel believes that the calculated seismic hazard derived from each individual expert's input needs to be presented. It is not clear whether this is included in SSHAC's recommendations. Regardless of how the aggregation is carried out, it is important to be able to compare results caused by each expert's input with those of the composite produced by aggregating the individual inputs. This comparison provides users with a good indicator of the diversity of input and its impact on the final calculations, as discussed in Chapter 3.

SSHAC proposes that this documentation follow a two-tiered approach that is to be applied to every element of a PSHA. Tier 1 documentation is defined as all documentation that must be published as part of the main report or its appendixes, so that it is widely accessible. Simply stated, tier 2 is everything else that constitutes background material for the analysis. SSHAC's prescription for what materials should go into the two tiers is spelled out for each of the elements of a PSHA (i.e., seismic source characterization, ground motion attenuation, and the methods used to produce the PSHA results).

The SSHAC report specifically states that the computer software used should be identified and archived. This would include any relevant programs and code that would be necessary for an independent analyst to replicate the study. Should problems be identified later with either the computer code or the input data, reanalysis is greatly facilitated. **The panel recommends that specialized computer programs needed to implement the SSHAC procedures be readily accessible to any group that wants to engage in seismic hazard evaluation as part of a research program or business venture.** The availability of these programs becomes especially important if the procedures recommended by SSHAC are so successful that they become the standard adopted by governmental regulatory bodies and the major engineering concerns of the nation.

To facilitate the accurate and timely documentation of PSHA projects, **the panel recommends that an individual or small team be designated as the Project Archivist and that a documentation plan be in place at the beginning of each project**. The thoroughness and complexity of the SSHAC approach, especially when the TFI is used, require that all participants have ready access at any time to materials generated previously. This implies a documentation process that keeps current with the rest of the project.

The panel concludes that the discussion of the documentation process in Chapter 7 of the SSHAC report provides thorough and useful guidance for numerous other applications in addition to seismic hazard assessment. **Documentation is not one of the more glamorous aspects of the scientific enterprise, but it is essential to the full realization of the benefits of the large investment in data acquisition, analysis, and interpretation that are characteristic of large projects.**

3
Treatment of Uncertainty

A fundamental aspect of the Senior Seismic Hazard Analysis Committee's (SSHAC) methodology is the distinct and separate treatment of aleatory and epistemic uncertainty. Throughout its report, SSHAC emphasizes the need to distinguish between these two types of uncertainty, the quantifications of their contributing sources, and the propagation and full display of the epistemic component to users (see, e.g., Sections 1.8 and 1.9). SSHAC deals with techniques to assess, elicit, combine, propagate, document, and display epistemic uncertainty, and it is clear that much if not most of the effort in any probabilistic seismic hazard analysis (PSHA) conducted according to SSHAC's recommendations would have to be expended in activities related to the handling of uncertainty.

The two fundamental types of uncertainty are defined by SSHAC as:

- Epistemic: the uncertainty attributable to incomplete knowledge about a phenomenon that affects our ability to model it.
- Aleatory: the uncertainty inherent in a nondeterministic (stochastic, random) phenomenon.

Epistemic uncertainty may be reduced with time as more data are collected and more research is completed. Aleatory uncertainty, on the other hand, cannot be reduced by further study, as it expresses the inherent variability of a phenomenon.

Making a rigorous separation between aleatory and epistemic uncertainty, as advocated by SSHAC, requires a level of effort and expertise much greater than that for most PSHA efforts. Therefore, the panel thinks it is appropriate to elaborate as to when and why such classification may be needed and indeed whether it is appropriate (these

issues are not addressed directly by SSHAC). In this regard, it is useful to consider separately two questions:

1. Is the aleatory/epistemic classification unique and clear?
2. Why is a separate treatment of epistemic and aleatory uncertainty needed and to what degree should it be pursued in a PSHA analysis?

Embedded in the second question are issues of utilization of results in which epistemic uncertainty and aleatory uncertainty are separated (i.e., of results stated in a "probability of frequency" format), either in the process of conducting the PSHA study or in the process of decision making by the ultimate user. In this chapter the panel briefly reviews SSHAC's position on these issues and makes some recommendations.

IS THE ALEATORY/EPISTEMIC DISTINCTION UNIQUE AND CLEAR?

SSHAC correctly points out that the classification of uncertainty as epistemic or aleatory depends on the model used to represent seismicity and ground motion. For example, epistemic uncertainty would be much greater if, in the assessment of seismic hazard at an eastern U.S. site, instead of representing random seismicity through homogeneous Poisson sources one used a model with an uncertain number of faults, each with an uncertain location, orientation, extent, state of stress, distribution of asperities, and so forth. As little is known about such faults, the total uncertainty about future seismicity and the calculated mean hazard curves would be about the same, irrespective of which model is used. However, the amount of epistemic uncertainty would be markedly different; it would be much greater for the more detailed, fault-based model. Consequently, the fractile hazard curves that represent epistemic uncertainty would also differ greatly.

A reasonable interpretation of the probabilistic models used in seismic hazard analysis is that they represent not intrinsic randomness but uncertainty on the part of the analyst about the actual states and laws of nature—for example, about the number of earthquakes of magnitude 6 to 7 that will occur in the next 50 years in a given crust volume. According

to this interpretation, all or most of the uncertainty in PSHA is due to ignorance. In certain cases, uncertainty due to ignorance may be expressed numerically by long-term relative frequencies. For example, with a very long record of seismicity, one could extract the long-term relative frequency with which earthquakes of magnitude 6 to 7 occur in a generic 50-year period. In the absence of other relevant information, it is reasonable to use this long-term relative frequency as a measure of epistemic uncertainty about the occurrence of the event in the next 50 years. Note that as interest in PSHA is typically in the occurrence of rare events in the near future and because the occurrence of such events depends to a large extent on the current physical conditions of the earth's crust near the site, ignorance or epistemic interpretation of the occurrence probability is more appropriate than the long-term relative frequency or aleatory interpretation. In certain parts of its report, SSHAC concedes that in reality there may be just one type of uncertainty. For example, Section 2.2.3 reads, in part:

> . . . Even though we have discussed probabilities appearing in the model of the world and the epistemic model, and we have given them different names, leading philosophers of science and uncertainty (e.g. de Finetti 1974; de Groot 1988) believe that, conceptually, there is only one kind of uncertainty; namely, that which stems from lack of knowledge.

Other statements support this position. For example, Section 2.2.6 states that ". . . the different terminology [aleatory versus epistemic] is not intended to imply that these uncertainties are of fundamentally different nature." Similarly, Section 1.8 points out that in the context of seismic hazard analysis, "the division between the two different types of uncertainty, epistemic and aleatory, is somewhat arbitrary." **The panel concludes that, unless one accepts that all uncertainty is fundamentally epistemic, the classification of PSHA uncertainty as aleatory or epistemic is ambiguous.**

Reference to a particular class of seismicity models (e.g., the models described in Sections 2.1 and Chapter 4 of the SSHAC report) produces some stability in the epistemic/aleatory distinction. However, if such distinction is to have any impact on the decisions, the basis for choosing any particular model type should be made clear, as alternative

and equally valid choices would lead to different decisions. In view of this undesirable dependence of epistemic uncertainty on the models selected for PSHA, one may question whether the epistemic/aleatory uncertainty decomposition is actually called for in a PSHA study and the extent to which it is needed for decision making by the users. These questions are addressed in the following section.

IS THE EPISTEMIC/ALEATORY SEPARATION NEEDED?

SSHAC does not provide a clear rationale for the need to separate aleatory uncertainty from epistemic uncertainty, although the report refers to several uses of this separation. Sections 2.2.5 and 2.2.6 of the report cite facilitated communication of results, discipline on the part of the analyst, and completeness of results. A "theoretical foundation" for the aleatory/epistemic distinction is offered in Section 2.2.6 by quoting a result by de Finetti in probability theory that shows how to combine epistemic and aleatory uncertainty to quantify total uncertainty for a particular (the binomial) model. However, the same result indicates neither how to separate the two uncertainties in practice (this is acknowledged by SSHAC) nor how to make decisions considering epistemic uncertainty. Therefore, the panel finds reference to de Finetti's result not relevant to whether or why the aleatory/epistemic distinction is necessary.

Reference to the decision-making implications of the epistemic/aleatory character of the uncertainty is made at the end of SSHAC's Appendix F, where it is stated that: "because epistemic and aleatory uncertainties are treated differently in making design and retrofit decisions, and because the median hazard is sometimes the preferred central measure of hazard due to its stability, it is also important to allocate uncertainties in the proper category." While it is true that the median curve is often preferred to the mean curve, a clear rationale for this practice or, more generally, a procedure for dealing with epistemic uncertainty in decision making is not presented in the SSHAC report. Finally, in Section 7.6 reference is made to the need for multiple hazard curves in the context of probabilistic risk assessment studies.

It is not the purpose of this discussion to analyze in detail each of the reasons for quantifying epistemic uncertainty. However, the panel observes that different uncertainty representations are appropriate for

different applications. To add focus to this discussion, we consider and contrast three main uses of quantified epistemic uncertainty in PSHA:

1. In the *elicitation and experts/model combination process*, quantitative estimates of epistemic uncertainty are used to characterize the credibility of alternative hypotheses and models, to assess the statistical variability of parameters, and to communicate this information among the experts and between the experts and the TFI.
2. In the course of a properly conducted analysis, the effect of epistemic uncertainty on the final hazard is used to *assess the relative importance of different models* (e.g., of the seismicity model versus the ground motion model) *and parameters* and to guide the analyst in seeking further information (data, expert opinion, etc.) to reduce uncertainty in the most cost-effective way.
3. A project's sponsor typically accounts for uncertainty in a hazard when *making decisions* (e.g., about the design of a new facility or the retrofitting of an existing one).

For ease of reference, we label these three phases of uncertainty consideration as the elicitation/combination phase, the PSHA planning phase, and the final utilization phase. Different needs for uncertainty representation characterize these phases.

In the *elicitation/combination phase*, experts need to be aware of all pertinent sources of uncertainty, including parameter and model uncertainties and their correlations, and the limitations and errors of the available data, so that they can make an informed assessment of the validity of alternative hypotheses, the accuracy of alternative models, and the value of data and can convey such uncertainties to the TI/TFI. **The panel finds the type of epistemic uncertainty analysis recommended by SSHAC to be most useful at this stage of a PSHA study.**

In the *PSHA planning phase* (which refers to resource allocation for the purpose of maximizing the reduction of uncertainty on the final hazard results), there is no need for a detailed analysis of uncertainty. In fact, such analysis is usually not available when the PSHA effort is structured. **For this purpose it may be sufficient to conduct limited sensitivity analyses, using bounding hypotheses, and to consider the level of effort that would be required to substantially reduce each component of uncertainty.**

The final *utilization phase* is critically important and arguably the one phase that should drive the level of uncertainty analysis and mode of uncertainty representation in a properly conducted PSHA. SSHAC's position is that the final results of a study should represent the epistemic uncertainty of the informed scientific community. This is roughly defined by SSHAC as the average of the uncertainties of the experts that make up the community (possibly weighted according to their degree of expertise, their outlier status, etc.).

A fundamental problem with this way of presenting the final results is that, as previously noted, the epistemic uncertainty in the hazard depends on which among many legitimate models one uses—for example, a deterministic or stochastic model of earthquake occurrence. What changes with the model is not the mean hazard but the amount of epistemic uncertainty and, therefore, all the fractile hazard curves—including the median. Therefore, any decision that is based on the fractile curves rather than the mean curve depends on the essentially arbitrary choice of how much epistemic uncertainty is included in the seismicity and ground motion models. This well-known fact has often been taken to mean that the only admissible decision rules are those based on the mean hazard and that other decision rules are wrong and should be excluded. In fact, this is not quite correct. As the study by Veneziano (1995) quoted in the SSHAC report shows:

1. If the mean hazard can be assumed to remain constant over the lifetime of the project (e.g., because only a small amount of relevant new information is expected to become available in the near future), decisions should be based exclusively on the present mean hazard.
2. On the other hand, if the mean hazard cannot be assumed to remain constant over the lifetime of the project, decisions should depend on possible future fluctuations of the mean hazard (Veneziano, 1995, p. 121).

These results show why the common practice of using mean probabilities is appropriate in certain cases but also explain why in other cases one should act conservatively. Notice that **the distinction does not depend on the total amount of current epistemic uncertainty but on the amount of total uncertainty that might be explained in the future and thus might cause the mean hazard to fluctuate.**

This is consistent with intuition. As a classic example of the irrelevance to decision making of the aleatory/epistemic classification, the betting attitude of a rational individual on the outcome of a coin flip should not change from before flipping, when all the uncertainty is aleatory, to after flipping (but before the outcome is revealed), when the same total amount of uncertainty is epistemic. On the other hand, the importance of temporal fluctuations of a mean hazard may be illustrated by considering the retrofitting problem, which occurs when, at some time after completion of a project, the estimated mean hazard changes and exceeds a regulatory limit. The reason why future volatility of the mean hazard should in this case affect present decisions is that the utility of each decision depends in an asymmetric way on future positive and negative changes in the mean hazard: large penalties are associated with retrofitting if the mean hazard increases, whereas only modest gains may result from future reductions in the mean hazard. The decision maker should consider the potential future volatility of the mean hazard and include it in his/her deliberations.

In the future, fundamental advances in PSHA may come from adopting this time-dependent view of earthquake safety decisions. However, explicit quantification of future volatility of a mean hazard would require a level of analysis even more sophisticated than that proposed by SSHAC, and the panel does not advocate such an extension at the present time, even for critical facilities.

Short of explicitly quantifying the future variability of the mean hazard, what could be done to provide the decision maker with a useful representation of epistemic uncertainty? One possibility, but certainly not the only one, is to calculate the mean hazard according to the uncertainty of each participating expert, when that expert acts as an evaluator (not integrator) of alternative models, data sets, etc. **To the degree that the beliefs held now by different members of the scientific community reflect possible future fluctuations in the overall community mean hazard, this should be useful input to the decision maker.** For example, this information would allow the decision maker to see how the decision he/she must make would vary if different experts in the informed scientific community had to make that same decision. Notice that the hazard curves derived from each expert do not suffer from the limitations of the fractile curves observed earlier; each of them is a mean hazard curve and therefore is insensitive to the choice of model type used by the expert.

Some observations should be made on presenting the final hazard results through the community mean hazard and the interexpert variability in the mean hazard, as just described.

1. One might argue that full epistemic uncertainty quantification is needed anyway, to calculate the mean hazard of the community and the mean hazard of the individual experts. However, this is true only in theory, as it is clear that different amounts of information are needed to estimate with confidence the mean value of a random variable, as opposed to its complete distribution. For example, the use of best estimates for recurrence and ground motion models often leads to hazard values that are close to the mean hazards obtained by considering a large number of alternative models. Moreover, there is no need when calculating the mean hazard to label accurately each component of uncertainty as epistemic or aleatory, provided that the total uncertainty is accounted for. Therefore, the elaborate machinery needed to carefully separate uncertainties of different types is no longer needed.
2. Much emphasis is given in the SSHAC report to intensive interaction among experts, discussion of alternative models, and exclusion or downweighting of outliers. These are all appropriate and remain valid under the format proposed here. **In essence, what changes is that the TFI quantifies not the total uncertainty of the scientific community, as done in the SSHAC approach, but the variability of the mean hazard according to the experts that make up that community.** In so doing, weights can be applied and outliers can be removed for the same reasons and in the same way as discussed by SSHAC.
3. The multiple interpretations, models, and model parameters at the basis of the elicitation process are not "lost." They remain part of the documentation of the PSHA study and should be made available to interested users. The panel anticipates that users will primarily be technical experts—for example, in the context of a regulatory review or an update of a PSHA study. However, that information should, for the most part, be irrelevant to the decision maker.

As observed previously, the correct way to represent epistemic uncertainty for decision making would be through the uncertain fluctuations of the mean hazard in future assessments. The expert-to-expert variability of the mean hazard at the time of the analysis is only a surrogate for this variability and is not entirely satisfactory because using

it this way implies that, during the time interval of interest, new evidence and knowledge may end up "proving right" one member of the present group of experts. While this may not be a valid assumption, documentation of the expert-to-expert variability in the mean hazard may be preferable to the full display of epistemic uncertainty proposed by SSHAC.

4
Seismic Source Characterization

Chapter 4 of the Senior Seismic Hazard Analysis Committee's (SSHAC) report, entitled "Methodology for Characterizing Seismic Sources," describes the key elements of a seismic source characterization (SSC): the seismic source requirements for a probabilistic seismic hazard analysis (PSHA), the uncertainties in seismic source characterization, and guidance on expert elicitation for seismic source description. The chapter presents a good description of the state of practice for SSC in a PSHA, as shaped chiefly by guidance on methodology from the seismic hazard programs of the Lawrence Livermore National Laboratory and the Electric Power Research Institute (EPRI), as well as from other PSHA exercises modeled on those programs, for many other critical facilities. In the panel's judgment, **practitioners of PSHA should be aware of and free to use other valid approaches to SSC.**

SCIENTIFIC VALIDITY AND CLARITY OF PRESENTATION

A primary concern of the panel is the overall scientific validity of the procedures recommended by SSHAC. The basic methodology for SSC described in the SSHAC report has been validated by extensive peer review of prior projects in which such a methodology was used. The SSHAC report correctly states that a seismic source is a construct developed for seismic hazard analysis as a means of approximating the locations of earthquake occurrences. Insofar as SSC involves a simplified representation of real-world complexity, the validity of the simplifications is always an issue. Such validity is generally tested as part of sensitivity analyses, which are an essential part of a PSHA, as correctly advocated in SSHAC's report. With regard to modeling real-world complexity, the

classification of seismic source types (Section 4.2) is nonunique, and the categories described in the report are admitted to be arbitrary. Nevertheless, they provide a useful framework for discussion and guidance on methodology.

The practitioner experienced in PSHA will have no trouble understanding SSHAC's Chapter 4. However, the nonpractioner scientist may be confused by the subtleties between differing concepts of a "seismic source" presented in chapters 4 and 5. Chapter 4 describes a seismic source as a geologic structure or as a domain within which the spatial and temporal occurrences of earthquakes are approximately uniformly distributed. Chapter 5, on ground motion, describes seismic source basically as a dynamic excitation in the earth that causes ground motion at the surface.

Readers of the SSHAC report should be aware that two different terms, upper-bound and maximum magnitude, and two symbols, m_u and M_{max}, are used Section 2.1 and in Chapter 4 to denote the largest-magnitude earthquake that a particular seismic source is capable of producing. This magnitude is the upper bound of the frequency of occurrence magnitude curve used in the analysis. A value for this parameter must be specified in order to carry out the integration over all relevant magnitudes when calculating seismic hazard. The problems encountered and conventional procedures used in the selection of M_{max} (m_u) and the specification of the substantial epistemic uncertainty often associated with it are discussed in Sections 4.2.2 and 4.3.2 of the SSHAC report.

If one accepts the basic formalism of uncertainty analysis presented in Section 2.2 of the SSHAC report, the approaches for characterizing uncertainties in SSC (Section 4.3) will seem logically consistent and well established in practice. Similarly, the guidance described in Section 4.4 for the expert elicitation process follows one's acceptance of the decision science methodology laid out in Chapter 3.

A notable gap in Chapter 4 of the SSHAC report is the absence of discussion on and guidance for earthquake catalogs. In Section 4.4 the technical facilitator/integrator (TFI) or the technical integrator (TI) is given responsibility for providing a comprehensive and uniform data base to the experts for use in the PSHA. The only guidance given, under the subheading "Area Sources" in Section 4.2.3, is the recommendation that "seismicity catalogs should be reviewed for uniformity in designation of magnitudes and for completeness as a function of magnitude, location,

and time. The association of older historical events with particular seismic sources should be assessed bearing in mind the location uncertainties."

Earthquake catalogs can play a major, even dominating, role in determining the outcome of a PSHA, particularly in the central and eastern United States, where information on active faults and other geologic structures is generally lacking. There are many problems hidden in earthquake catalogs that need be sought out and identified. There may be improper or mistaken entries, particularly for historic earthquakes. In many cases, locations and sizes were assigned to historic earthquakes based on inadequate or incomplete information. Unfortunately, modern earthquake catalogs often do not indicate which events have been critically reexamined and which have been carried forward without question from original catalog compilations.

Uniformity of the data with time is also variable even in times of instrumental monitoring. Changes in network configurations and sensitivity and changes in the procedures for computing event magnitudes reported in earthquake catalogs (often not documented in an easily available form) should be sought out and carefully considered in a PSHA. Tests are available for identifying time-varying systematic shifts in reported magnitudes. Declustering or decomposing earthquake catalogs into main and secondary events (foreshocks, aftershocks, swarm events) is a nontrivial procedure that also requires careful attention.

Recognizing that earthquake recurrence relationships based on seismicity depend critically on factors such as those described above, EPRI undertook major efforts to address these and other earthquake data base issues, which are still of great importance in PSHA—both in principle and in continuing practice. **Those who utilize the SSHAC procedures should be aware of these requirements for preparation of their earthquake catalog for PSHA. To the panel's knowledge, a comprehensive study of the effects of systematic changes in earthquake catalogs on the results of a PSHA has not been done.**

Most of Chapter 4 of the SSHAC report is well organized and well written, and the presentation should be easy for general readers to follow. The text refers to Appendixes H and I, each of which provides some ancillary pertinent material. Appendix H describes the results of a workshop on expert elicitation of seismic source (zone) information, while Appendix I describes effects of a nonuniform spatial distribution of seismicity in a seismic source (zone). Both of these appendixes are informative.

The table in Section 4.2.1 is important for guidance, but it is confusing. The lines beginning with "Faults" and "No faults" should be understood to be "if" statements, recognizing "fault" to mean a "Type 1 seismic source" (i.e., "If no Type 1 fault source within 50 km of a site, then . . .").

Because the SSHAC report is intended for general PSHA guidance, the following question arises: Is the EQPARAM code (which is introduced as an important element of the methodology in Section 4.3.5) readily available or is it proprietary to EPRI? If the latter, it should have been described as such. This question illustrates the concerns of the panel about software availability expressed in the previous discussion of documentation.

CONTRIBUTIONS TO THE DEVELOPMENT OF PSHA

Because SSC is such a major component of a PSHA, the comprehensive methodology for expert elicitation presented in Section 4.4 of the SSHAC report is an important contribution. On first reading, the material in Chapter 4 may appear to be just a restatement of Chapter 3. However, SSHAC is correct in noting in Section 4.4 that the elicitation procedures and methods for SSC differ from those for ground motion characterization. Further, "lessons learned" from past SSC exercises are incorporated into major PSHA projects (Appendix H).

Another important contribution of Chapter 4 and its accompanying appendixes is the practical guidance provided for carrying out sensitivity analyses to determine "what drives the seismic hazard" and "what contributes significantly to uncertainties in hazard." Basic discussion relevant to SSC is presented in Section 4.3.6, but important details are given in Appendix G and Section 7.8.

A third major contribution of Chapter 4 is the exposition in Section 4.3.5 (bolstered by Appendix I) of the effects of spatial variations in seismicity within a seismic source vis-à-vis the assumption of homogeneous seismicity. The analysis techniques date from the EPRI program (EPRI, 1989, as cited in the SSHAC report), but the detailed discussion and examples presented there forcefully demonstrate how the usual assumption of homogeneous seismicity for seismic sources can, under certain predictable cases, significantly affect both the mean seismic hazard and its statistical uncertainty.

THE OUTLOOK FOR EVOLUTION OF SSC

While affirming the scientific validity and practical effectiveness of the SSC methodology set forth in the SSHAC report, the panel recognizes that the scientific community will naturally strain against the confines of SSHAC's prescriptions for SSC. The panel applauds SSHAC's perspective that "[its] formulation should not be viewed as an attempt to 'standardize' PSHA in the sense of freezing the science and technology that underlies a competent PSHA, thereby stifling innovation" (Section 1.2 of the SSHAC report). A few brief examples suffice to illustrate current trends in the scientific community that may influence the evolution of SSC. Diverse trends lead to advocacy for both greater simplification and greater complexity.

Frankel (1995) proposes a method for PSHA that uses spatially smoothed representations of historic seismicity instead of seismic source zones to directly calculate probabilistic seismic hazard. Insofar as he demonstrates the capability to produce values of mean seismic hazard similar to those from the more complicated EPRI methodology, his simple methodology offers understandable attraction. The applicability obviously pertains to cases where seismicity "drives the hazard"—either for specific regions or for definable exposure periods.

In terms of modeling earthquake occurrence with greater complexity, one example is the multidisciplinary approach (e.g., Ward, 1994), in which data from space geodesy and synthetic seismicity are added to the traditional information from geology, paleoseismology, and observational seismology. Main (1995) examines the implications if earthquake populations are really an example of a self-organized critical phenomenon. If this is correct, the *a priori* assumption of the Gutenberg-Richter frequency-magnitude distribution is no longer valid in some cases, and Main provides evidence for questioning the use of only the Poisson distribution in seismic hazard analyses, based on the accumulating evidence of local or long-range interactions of earthquakes. It should be pointed out that PSHA is not limited to the use of the Gutenberg-Richter relationship. Alternate estimates of the frequency-magnitude distribution are, and have been, used in probabilistic analyses.

Main (1995) also discusses an independent approach to the vexing problem of estimating the maximum-magnitude earthquake that is "credible" for a seismic source zone, based on his suggested distribution of moment release and the long-term slip rate on the causative fault

system. Geophysicists are becoming increasingly aware of the nonstationarity of earthquake occurrence, particularly in light of observations of fault interactions leading to "triggered" or "encouraged" earthquakes. As earth scientists improve their ability to assess time-varying earthquake potential on active faults, SSC will evolve correspondingly. Indeed, "time-variable seismic hazard" is already a topic of special sessions at geophysical society meetings.

5
The Estimation of Earthquake-Generated Ground Motion

Chapter 5 of the Senior Seismic Hazard Analysis Committee's (SSHAC) report, entitled "Methodology for Estimating Ground Motions on Rock," addresses the basic building block of a well-executed probabilistic seismic hazard analysis (PSHA) that has the surest observational and theoretical foundation. The past two decades have brought significant theoretical advances in ground motion models, as well as significant new data sets with which to test the new models. Fundamental to the stability of state-of-the-art high-frequency (f = 1 Hz) ground motion estimates is the essential constancy of earthquake stress drops. This allows the substantial experience developed from California and elsewhere to be transferred to the eastern United States (EUS) with little modification.

There are, to be sure, real variations in earthquake stress drops, and recent data for the EUS point to some anomalous magnitude-dependent high-frequency excitation (Atkinson, 1993). The EUS data set on the excitation and propagation of earthquake ground motion for the purposes of PSHA is still very sparse. Model predictions of EUS earthquake ground motion, whether empirical or theoretical, can vary significantly across the magnitude, distance, and frequency range of interest.

SCIENTIFIC VALIDITY AND CLARITY OF PRESENTATION

SSHAC's Chapter 5, together with the supporting Appendixes A and B (Ground Motion Workshops I and II), is an impressive synthesis of current knowledge about estimating high-frequency ground motions and their uncertainties in the EUS. The reader experienced in SHA will note

that site-response issues, including nonlinear effects, are not addressed, on the grounds that they can only be incorporated on a site-specific basis.

Chapter 5 is itself a well-written primer on the essentials of ground motion estimation, valid for any region in which earthquakes occur. It begins with basic ground motion measures; provides the fundamentals of magnitude, distance, and site response; and describes the essentials of empirical and theoretical predictions of earthquake ground motion. It explicitly warns against the use of fixed spectral shapes anchored by peak ground acceleration (PGA) alone, and then progresses to a discussion of uncertainty in ground motion predictions. A fourfold decomposition of uncertainty for the Hanks and McGuire (1981) point-source, stochastic model, the simplest physical model used in these predictive exercises, is demonstrated in this discussion. Readers should study this decomposition carefully (Table 5-1, Section 5.5.1). It is difficult, and, if this example is not well understood, similar attempts at uncertainty decomposition for more sophisticated and parametrically complicated models will be frustrating.

Section 5.7, "Specific Expert-Elicitation Guidance for Obtaining Ground Motion Values," is based on the results of Workshops I and II, reported in detail in Appendixes A and B. Figure 5-5, reproduced as Figure 2.1 in this report, is intended to guide readers through the process. Regrettably, it is not well keyed to the description in the text.

CONTRIBUTIONS TO THE DEVELOPMENT OF PSHA: SUMMARY OF THE GROUND MOTION WORKSHOP RESULTS

The comprehensive treatment of ground motion estimation in Appendixes A and B is an important contribution to the SSHAC effort. Workshop I provided for the presentation of four basic ground motion estimation models: (1) intensity-based models presented by M. D. Trifunac, (2) empirical models presented by K. W. Campbell, (3) stochastic or random-vibration models presented by G. M. Atkinson, and (4) the empirical source-function method presented by C. Saikia. These proponents of the models were asked to evaluate the models in the company of 10 additional experts, the "invited participants" listed in Table A-1 of the SSHAC report. The principal result of Workshop I was rejection of intensity-based models for estimating ground motion in the EUS (SSHAC Table A-2). Additional information was collected on the

applicability or validity of all models as a function of frequency, magnitude, and distance (SSHAC Tables A-3 and A-4). These polls of the assembled experts also show a distinct preference for the stochastic models.

Workshop II proceeded to actual ground motion numbers and their uncertainties on the basis of the "selected models" resulting from Workshop I. The threefold elicitation exercise that constituted Workshop II, described below, provided for pre-, co-, and postworkshop estimates. Prior to the workshop, the four proponents were asked to provide estimates of peak acceleration and spectral accelerations based on the ground motion models they actually use, along with the corresponding estimates of epistemic and aleatory uncertainties. The distances, frequencies, and magnitudes for which estimates were requested are listed in an unnumbered table in "Instructions for Proponents," Appendix B. In keeping with the Workshop I preference for stochastic models, two of the four Workshop II proponents supported stochastic models (Atkinson and Silva), although there are significant differences between their models.

In advance of Workshop II these ground motion estimates were sent to three additional experts. These experts were asked to provide their own estimates of ground motion and uncertainties for the same distances, frequencies, and magnitudes, on the basis of what the proponents had provided, as well as any other information they considered relevant. Significantly, the four proponents were also asked to perform as experts; as such, their ground motion estimates were generally not the same as those they provided as proponents. These pre-Workshop II ground motion estimates and uncertainties are labeled as Expert 1 results, examples of which are shown in SSHAC Figure B-3, reproduced here as Figure 5.1a.

The second stage of the elicitation process occurred at the workshop, attended by all proponents and experts, the integration team, and several observers (SSHAC Table B-1). The principle of "active listening" was put to work, the idea being that all proponents and experts were to understand what every other proponent and expert was doing, whether or not he/she agreed with it. The panel concludes that this worked very well, revealing significantly different interpretations of key terms and procedures. It is noteworthy that Workshop II deliberations also revealed considerable misunderstandings about the differences between epistemic and aleatory uncertainties.

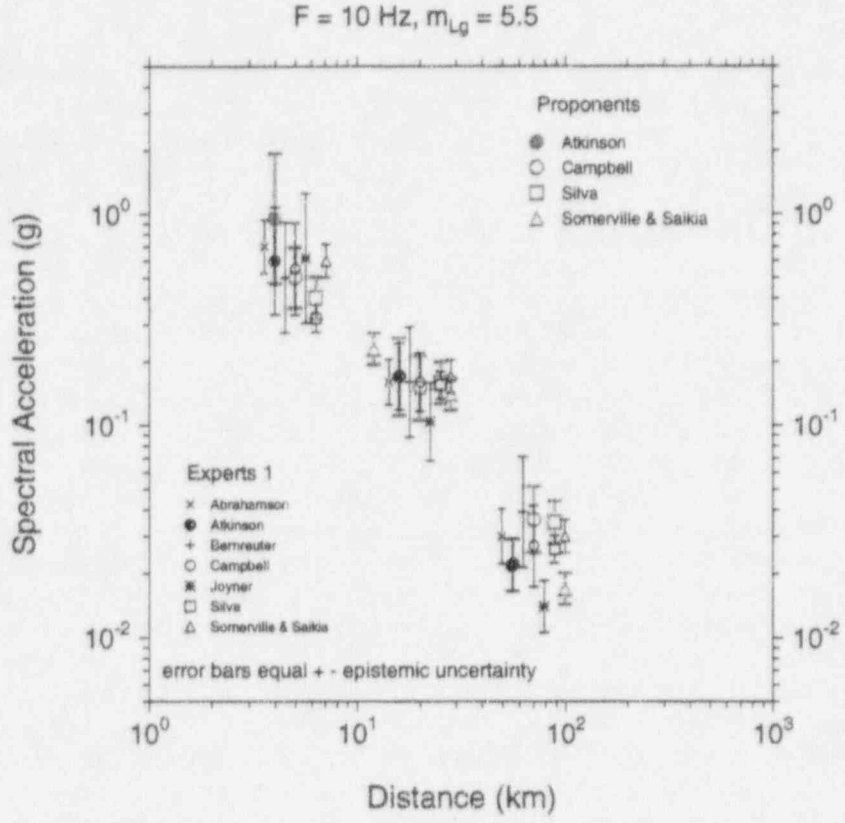

FIGURE 5.1a Comparison of proponents' estimates (gray) to Experts 1 estimates (black) of 10-Hz spectral acceleration for $m_{bLg} = 5.5$. The error bars represent $\pm \sigma_{epistemic}$ range.

Next, experts (at this stage all proponents were now experts) were asked to reconsider overnight their estimates of ground motion and uncertainties. This led to the Expert 2 results, which are compared to the Expert 1 results. An example (SSHAC Figure B-7) is reproduced here as Figure 5.1b. The differences are modest to zero at $f = 10$ Hz and somewhat greater at $f = 1$ Hz.

Two activities followed the workshop. First, all experts were invited to change their estimates one more time. Only a few did, and no one offered significant changes. An example of the integrated Expert 3 (postworkshop) results is shown here in Figure 5.1c (SSHAC Figure B-21). The second postworkshop activity was the manipulation of the Expert 3 results by the Integration Team. The results of the seven experts were weighted equally (SSHAC Table B-8, shown here as Table 5.1), and the results of the four proponents were weighted unequally (SSHAC Table B-9). The former are the preferred results, but the differences in median values and epistemic and aleatory uncertainties are slight.

IMPLICATIONS FOR FUTURE GROUND MOTION ESTIMATION

The many successes and few limitations of the Workshop II elicitation/integration process are summarized in Section B.5, "Concluding Observations and Discussion," of the SSHAC report. The panel is impressed with the success of this process in two principal ways, one of which SSHAC recognized and the other it did not.

SSHAC recognized explicitly that "the Proponents and Experts exhibited a striking amount of agreement. . . ." **Once freed from the thicket of unintentional disagreements, mutual misunderstandings, and individual egos, the group of specialists who participated found that what it knows about ground motion estimation is impressively consistent. The panel doubts that this degree of consistency and agreement could have been achieved without this highly interactive elicitation/integration process.**

There may be some who will believe that this agreement is illusory, that in some unspecified way it was cajoled or coerced. The panel finds no evidence of this. Doubters should note the workshop finding that "the estimated values of aleatory uncertainty for 10 Hz

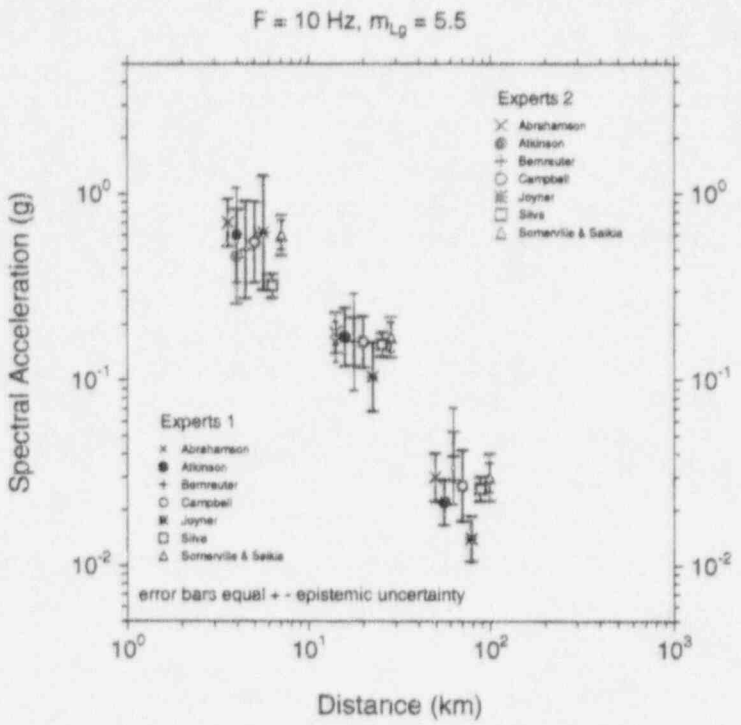

FIGURE 5.1b Comparison of Experts 2 results (gray) to Experts 1 results (black) for 10-Hz spectral acceleration at $m_{bLg} = 5.5$ as a function of distance.

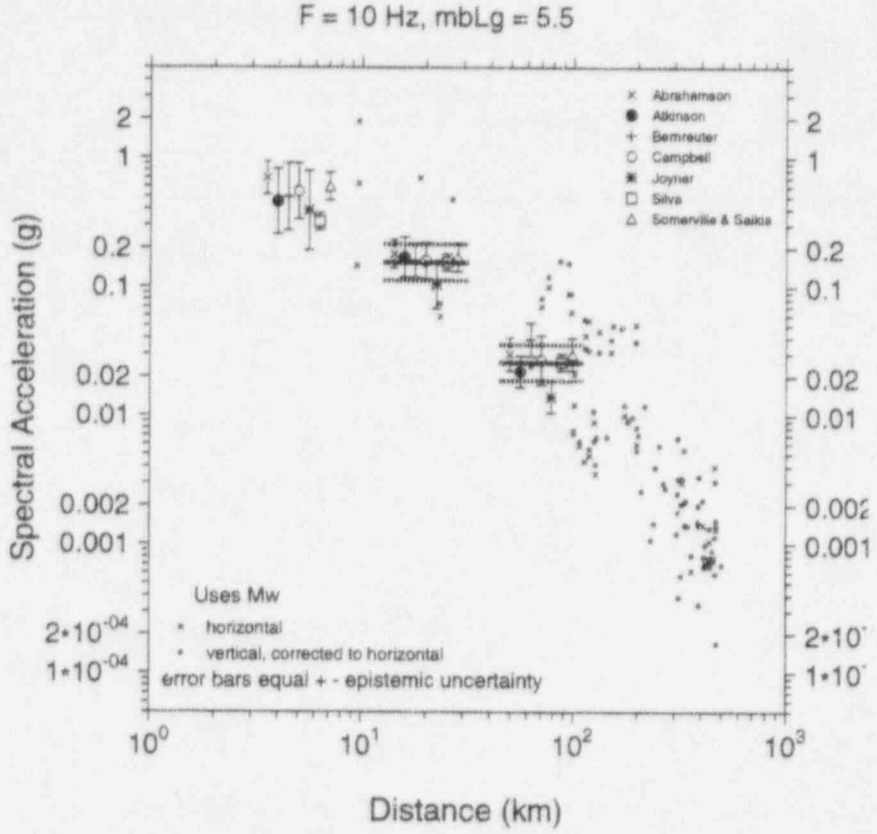

FIGURE 5.1c Experts 3 results, together with mean values and variances obtained from equally weighting the Experts 3 results for 10-Hz spectral acceleration at $m_{bLg} = 5.5$ as a function of distance. Small circles and crosses represent instrumental data.

TABLE 5.1 Results of Integrating Experts' Estimates with Equal Weights (Table B-8, Appendix B, SSHAC Report)

f (Hz)	m_{bLg}	R (km)	Median Amplitude (g)	Epistemic Std. Dev.	Aleatory Std. Dev.
1	5.5	20	1.09E-02	0.48	0.80
	5.5	70	2.27E-03	0.46	0.80
	5.5	200	9.36E-04	0.37	0.80
	7.0	20	1.67E-01	0.66	0.78
	7.0	70	4.50E-02	0.71	0.78
	7.0	200	1.82E-02	0.73	0.79
2.5	5.5	20	4.17E-02	0.34	0.77
	7.0	20	3.67E-01	0.53	0.73
10	5.5	20	1.55E-01	0.32	0.73
	5.5	70	2.58E-02	0.32	0.75
	7.0	20	8.45E-01	0.52	0.70
	7.0	70	1.88E-01	0.53	0.72
25	5.5	20	2.13E-01	0.34	0.73
	7.0	20	1.07E+00	0.51	0.70
PGA	5.5	70	1.28E-02	0.41	0.75
	7.0	70	9.36E-02	0.51	0.70

and PGA are, however, significantly higher than [the] values obtained using western North America strong-motion data, especially for large magnitudes."

SSHAC did not comment on the extent to which the workshop ground motion estimates and uncertainties can actually be used in future PSHA studies, at any level. The panel recognizes that there is a certain incompleteness about Table 5.1. Considerable interpolation and some extrapolation of the results in that table will be required to cover the many distances, frequencies, and magnitudes that must be considered in even the lowest-level PSHA. Unfortunately, the elicited results for R = 5 km, where R is the distance between the seismic source and the affected area, are not presented by SSHAC, presumably because of problems with the interpretations of "closest distance."

Even if the SSHAC ground motion results are not suitable for further use in their present form, the panel wonders how many times this information will be reelicited in the future. **The panel believes that community consensus on PSHA-type ground motion issues, at any level of PSHA, may well be close at hand, at least within the limits of the ground motion models and data sets available in 1994.** The broad agreement resulting from the two SSHAC ground motion workshops led to this opinion of the panel. With further consideration of some additional distances, frequencies, and magnitudes, together with appropriate interpolation schemes, ground motion matters of concern to PSHA could well be resolved at least for the next few years.

6
Summary and Conclusions

GENERAL APPRAISAL OF THE SSHAC REPORT

The Senior Seismic Hazard Analysis Committee's (SSHAC) report offers substantial contributions to the foundations and practice of probabilistic seismic hazard analysis. But the primary focus of the report is not on how to create an assessment from the inputs; only in Chapter 2, in an introductory fashion in Chapter 6, and in Appendix J is a methodology for calculating the hazard estimates and their uncertainties addressed.

Instead, the central theme of SSHAC is guidance on the process of eliciting and aggregating expert opinion on seismic sources, seismicity within these sources, and ground motion attenuation, as well as the associated uncertainties and final estimates of the hazard. SSHAC focused on this theme based on its conclusion that the reason for some serious discrepancies in the results of prior studies is differences in ways in which these inputs were derived, even though the work was done by competent specialists working from the same or similar data bases. **In the panel's view, SSHAC's most important message is that the quality of a PSHA using multiple experts can be enhanced by careful and wise choice of experts and skillful facilitation of expert discussion and interaction through workshops and other meetings.**

The panel believes it very important to emphasize what the SSHAC report is and what it is not. The report presents a procedure for using experts in seismic hazard evaluation and for determining the uncertainties at key stages of the hazard analysis process. Its primary domain of application is to nuclear and other critical facilities. According to SSHAC, if a project sponsor and the analysts choose to do a probabilistic hazard analysis, its procedures will yield stable results. The SSHAC report is not a defense of the probabilistic approach to hazard assessment. In particular, SSHAC explicitly excludes any discussion of the nonprobabilistic methods of seismic hazard assessment. The panel

accepts this decision of SSHAC on the grounds that an evaluation of the relative effectiveness of the two approaches, or their relationship, was not in the committee's charge. The full-blown version of the SSHAC procedure, utilizing the technical facilitator/integrator (TFI) technique where needed, is costly and will almost certainly be used only for major critical facilities. The SSHAC report offers useful guidelines as to the level of effort required for various kinds of problems and for various levels of information already available to analysts. In the view of the panel, simpler methods of probabilistic hazard analysis are appropriate for application to noncritical facilities.

GENERAL SHORTCOMINGS AND LIMITATIONS OF THE SSHAC REPORT

The SSHAC report, with its appendixes, is a lengthy and complex document that requires careful reading. Many important ideas, including clarification of the limitations of the SSHAC procedures, are distributed throughout the text. A casual scanning of the document may leave readers with incorrect impressions as to what SSHAC has recommended, especially with regard to nonnuclear facilities. Most importantly, the report appears to have been written for those already quite familiar with PSHA methods, offering guidance on a preferred way to get stable results from a PSHA.

SSHAC's Executive Summary will be useful to administrators and project sponsors who are not specialists in hazard analysis methodology, but it includes nothing about the excellent earth science materials that are in the report and its appendixes.

SSHAC provides an up-to-date procedure for obtaining stable results from the application of PSHA principles that have been established in past practice. It does provide a consistent and systematic approach to elicitation and aggregation of diverse expert opinion and the uncertainties that arise therefrom, but this is not the same as the calculation of seismic hazard from the information elicited.

The SSHAC report does not make reference to nuclear reactors or other nuclear facilities, thereby lending an air of generality to its final report and the applicability of its recommended procedures. The panel believes, nevertheless, that the flavor of the report is strongly influenced by concern for applications to nuclear facilities and this generality is more

apparent than real. In response to recommendations in the panel's March 1995 letter report (Appendix B) to the U.S. Nuclear Regulatory Commission, SSHAC did attempt to narrow the scope of the applications for which its recommended procedure is intended. Disclaimers are included in several places that are technically adequate to protect a practitioner who chooses not to use the SSHAC prescription against the need to defend that decision in a regulatory situation. Nevertheless, it seems clear that the report was written to support the highest, most sophisticated level of PSHA practice. Because the concept of the TFI is held by SSHAC to be one of its most important contributions to PSHA practice, a great deal of space is devoted to this topic, even though there are repeated comments that it is not needed for many of the issues that arise. The impression is given that this highest level of operation is really the key to success in general.

The panel concludes that the SSHAC contention—namely, that all PSHA projects should share the same basic principles and goals—should be taken as an overarching postulate for project design. But this contention should not be taken as implying or imposing the full elaborate and demanding methodology for application to every PSHA study. That alternate simpler methods may well be adequate for noncritical facilities is acknowledged by SSHAC, but they are not discussed nor is guidance offered as to where readers can learn about them.

In meetings and in its letter report of March 1995 (Appendix B), the panel urged SSHAC to document in adequate detail the manner in which lessons leading to the recommended SSHAC procedures were learned from the study of prior PSHA studies. Although the SSHAC report states that its conclusions are based on a thorough review of a number of such studies, the requested details are not offered and no previous PSHA analyses other than the Lawrence Livermore National Laboratory and Electric Power Research Institute studies are referenced.

The panel's evaluation of SSHAC's treatment of uncertainty is presented in detail in Chapter 3 of this report. The panel acknowledges that recognition of the two kinds of uncertainty is useful in elicitating expert opinion and in making decisions about where additional data gathering and research are likely to lead to reduced uncertainty about hazard estimates. However, as discussed in Chapter 3, the panel has reservations about how this distinction is ultimately helpful to final users, especially because the distinction between uncertainty types is sometimes

ambiguous and the amount of epistemic uncertainty regarding a hazard depends on the type of models used in the analysis.

Moreover, it is the impression of the panel that the statistical analysis and uncertainty separation procedures recommended in the SSHAC report are, at times, more sophisticated than is warranted by the data on which such analysis is based or the purposes for which the results are used.

The problem of integrating the opinions of a group of experts is difficult. It is treated in greatest detail in Appendix J of the SSHAC report. The panel found that this treatment is not easy to follow and that specific aggregation models described are not exhaustive. **Therefore, the panel recommends that the quantitative methods of Appendix J be used as examples and not be regarded as prescriptive procedures.** Given the current state of the art in formal expert aggregation and the difficulties specific to the earthquake hazard problem, the panel suggests that judgmental combination rules may be at least as valid as quantitative procedures.

SOME CONTRIBUTIONS OF SSHAC TO HAZARD ASSESSMENT

The contributions that the SSHAC report makes to the hazard assessment process are discussed in detail in the preceding sections of this report. A few key items are highlighted here.

The TFI Methodology

SSHAC considers the TFI methodology to be the centerpiece of its work and developed it from lessons it learned from prior hazard analysis studies and from workshops conducted as part of its study. The panel is favorably impressed with the concept and its implementation in the two ground motion workshops (SSHAC's Appendixes A and B). Readers of the SSHAC report should keep in mind that use of a TFI is not recommended or needed for all hazard assessments and should not even be viewed as a rigid prescription for a high-level PSHA. The TFI elicitation procedure is not synonymous with PSHA methodology.

Clear Definition of Experts' Distinct Roles as Proponents, Evaluators, and Integrators

It is important that experts be educated to the significance of their distinct role as proponents of a particular position or as evaluators. The panel is not sure that experts can truly assess the view of the whole informed community on the entire range of relevant issues.

Results of SSHAC-Sponsored Workshops

SSHAC held workshops on seismic source characterization, ground motion estimation, and earthquake magnitudes. The outputs of these workshops (Appendixes A, B, C, H), especially those on ground motion, are a valuable contribution of the SSHAC effort and led to the formulation of many of the recommended procedures in the committee's report.

Considering the broad consensus on ground motion modeling that was reached at the end of Workshop II, the panel believes that a real opportunity exists now to formulate, with further work to fill in necessary details, a ground motion model that can be used as a standard in the eastern United States for PSHA until new data or future theoretical developments warrant a reevaluation. The results of this effort would eliminate the need to elicit again ground motion input for each hazard analysis and could be used as a baseline for more detailed studies as needed for specific problems.

References

Atkinson, G.M. 1993. Earthquake Source Spectra in Eastern North America. *Bull. Seismol. Soc. Am.* 83:1778-1798.

Bernreuter, D.L., J.B. Savy, and R.W. Mensing. 1987. *Seismic Hazard Characterization of the Eastern United States: Comparative Evaluation of the LLNL and EPRI Studies.* NUREG/CR-4885, Lawrence Livermore National Laboratory, Livermore, CA.

Coppersmith, K.J., R.C. Perman, R.R. Youngs, T.A. Crampton, L.A. DiSilvestro, and J.C. Nesbit. 1995. Probabilistic Volcanic Hazard Analysis at Yucca Mountain, Nevada, Using Expert Judgment. In *Proceedings of the U.S. DOE Natural Phenomena Hazards Mitigation Symposium,* Denver, Colo. Nov. 13-14.

Frankel, A. 1995. Mapping Seismic Hazard in the Central and Eastern United States. *Seismol. Res. Lett.* 66:8-21.

Hanks, T.C., and C.A. Cornell. 1994. Probabilistic Seismic Hazard Analysis: A Beginner's Guide. In *Proceedings of the. Fifth Symposium on Current Issues Related to Nuclear Power Plant Structures, Equipment and Piping,* pp. I/1-1 to I/1-17, North Carolina State University, Raleigh.

Hanks, T.C., and R.K. McGuire. 1981. The Character of High-Frequency Strong Ground Motion. *Bull. Seismol. Soc. Am.* 71:2971-2995.

Keeney, R.L., and D. von Winterfeldt. 1991. Eliciting Probabilities from Experts in Complex Technical Problems. *IEEE Trans Eng. Manage.* 38:191-201.

Main, I.G. 1995. Earthquakes as Critical Phenomena: Implications for Probabilistic Seismic Hazard Analysis. *Bull. Seismol. Soc. Am.* 85:1299-1308.

McGuire, R.K. 1995. Probabilistic Seismic Hazard Analysis and Design Earthquakes: Closing the Loop. *Bull. Seismol. Soc. Am.* 85:1275-1284.

Senior Seismic Hazard Analysis Committee (SSHAC). 1997. *Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts.* U.S. Nuclear Regulatory Commission, NUREG/CR-6372, Washington, DC.

Veneziano, D. 1995. Uncertainty and Expert Opinion in Geologic Hazards. *Proceedings. of the Earth, Engineers, and Education: A Symposium in Honor of Robert V. Whitman,* October 7-8, 1994, Massachusetts Institute of Technology, Cambridge.

Ward, S. N. 1994. A Multidisciplinary Approach to Seismic Hazard in Southern California. *Bull. Seismol. Soc. Am.* 84:1293-1309.

Appendix A

ACRONYMS

DOE	Department of Energy
DSHA	deterministic seismic hazard analysis
EPRI	Electric Power Research Institute
EUS	eastern United States
LLNL	Lawrence Livermore National Laboratory
NRC	National Research Council
PGA	peak ground acceleration
PSHA	probabilistic seismic hazard analysis
SHA	seismic hazard analysis
SSHAC	senior seismic hazard analysis committee
SSC	seismic source characterization
TI	technical integrator
TFI	technical facilitator/integrator
USNRC	United States Nuclear Regulatory Commission

Appendix B

LETTER REPORT OF THE PANEL ON SEISMIC HAZARD EVALUATION, MARCH 1995

Committee on Seismology, National Research Council
Comments on SSHAC Draft Report of 11 November 1994
Based on the Panel Meeting of December 9-10, 1994

The Panel on Seismic Hazard Evaluation of the Committee on Seismology, National Research Council (NRC), is charged with reviewing the report to be produced by the Senior Seismic Hazard Analysis Committee (SSHAC) under the sponsorship of the U.S. Nuclear Regulatory Commission (USNRC), the Department of Energy (DoE), and the Electric Power Research Institute (EPRI). The USNRC prescribed that the Panel provide feedback to SSHAC as they prepare their report, but in such a way as not to compromise the objectivity of the Panel in providing its review of the final product. SSHAC submitted for review a draft of their report in mid-November, 1994, and the Panel met, with all SSHAC members present, on December 9, 1994, for discussion of the draft.

Unfortunately the draft was not complete, missing some key appendices, some sections of text, and an executive summary. It should be understood that the Panel may have comments with regard to the missing material when it is available for the final review. The discussions of December 9 were carried out in the presence of representatives of the sponsoring organizations. The Panel met in executive session on December 10 to continue its review. The resulting comments and recommendations are submitted to the USNRC.

The suggestions made are offered as guidance to SSHAC on the issues at this stage of their work, in accord with the request of the USNRC. They should not be interpreted as a substitute for the final report to be developed by the Panel.

GENERAL COMMENTS

The Panel believes that the draft report is a basis for a useful final product that has the potential to advance the process of Probabilistic Seismic Hazard Analysis (PSHA). However, the Panel feels that the introduction to the report must be expanded to make clear the purpose and scope of the report, and specifically to state what the report is not. As it stands, the report implies that the methodology is applicable to a broader range of facilities than can be justified. The full range of alternative approaches is not discussed, let alone taken into account.

From the discussions, it appears that there may be a conflict between the expressed needs of the USNRC for a single unified, fully prescribed regulatory method of seismic hazard analysis (SHA) and the attempt by SSHAC to produce a general consensus methodology. The USNRC wants a prescribed procedure that is based on what has been learned from past PSHA experiences. The USNRC recognizes that the way in which input from experts was obtained is a main reason for the discrepencies between the analyses made by Lawrence Livermore and EPRI.

The Panel recognizes the strengths of the report and the significant contributions it offers to PSHA. As applied to nuclear regulations the SSHAC report breaks new ground in its discussion of the Technical Integrator (TI)/Technical Facilitator Integrator (TFI) approaches. However, as discussed in detail below, the presentation of these ideas needs to be made more clear to eliminate some apparent contradictions and advise the users of the report when the full TFI treatment is called for. The TI/TFI approach has the potential to overcome some aspects of past PSHA applications that have led to objections by critics of the whole process.

Because the focus of the report is on process for PSHA, rather than on the underlying earth science, the detailed attention to the treatment of uncertainty is appropriate. However, as discussed below, the motivation for this careful treatment of uncertainty and the way in which the results will be applied are not made clear to the potential user.

Again without yet having the benefit of full discussion of the subject, the Panel feels that the recommendation that behavioral aggregation of expert input be employed is sound, because mechanical aggregation algorithms, if used as "black boxes," may lead to poor results.

CONCERNS AND PROBLEMS

Recommendations for Improvement

Some suggestions for revision and restructuring of the report were given orally to SSHAC during the Panel meeting. The most essential of these, which the Panel feels cannot be neglected during revision of the report, are repeated here for completeness of the record.

The word "Consensus" should be removed from the title, perhaps replaced by a more appropriate adjective.

An excellent executive summary is essential for the success of this report. The report is lengthy and detailed. The key findings and recommendations of SSHAC must be assembled in concise, easily understandable form if they are to be accessible to others than the experienced practioneer of PSHA.

The draft as submitted is overly repetitious. Unnecessary redundancy should be eliminated, to reduce the length substantially without loss of content.

The specific criticisms to follow all can be categorized as due to one or more of the following: inadequate *focus* of the report, absence of the *history* of evolution of the key concepts and recommendations, or lack of a presentation of the *context* within which the report was developed and is to be understood and applied.

Motivation. The reader should be offered better motivation for adopting the procedures required or recommended in the report. In addition, the context for the procedures should be framed in such a way that the PSHA analyst who follows other procedures for any of a number of valid reasons is not put in a position of having to defend in a regulatory situation the failure to carry out the SSHAC prescription in every detail.

PSHA methodologists often have sound reasons for introducing new concepts and approaches, but have not always included in their reports the background reasoning that has led to these innovations. Where it exists in this report, this shortcoming must be overcome if the final SSHAC product is to be widely accepted and applied. In particular, the report should say how the results are to be used as motivation for the great emphasis on the distinction between aleatory and epistemic uncertainty and the need to separate the two in SHA.

Space and emphasis devoted to the TFI approach. Scattered through the text, and asserted by SSHAC members at the December 9 meeting, is the key idea that the full TFI approach is required only for some complex issues for which a review of the published literature cannot produce satisfactory input to the PSHA process. However, the great detail in which the recommended TFI approach is depicted tends to obscure this principle. The reader is left with the impression that the use of the TFI is dominant in a properly executed PSHA.

- SSHAC must carefully set out the criteria for deciding if an issue requires a TFI. What are the operational criteria for deciding if an issue is of type A, B, C?
- SSHAC must state its perception of the qualifications required of the TFI. The recommendation for use of a *strong* TFI for prescribed issues, without clearly expressed qualifications, contradicts one of the stated criteria for success: that the recommended methodology, when applied independently by different groups, should always yield comparable results.
- The Panel is concerned that the TFI is empowered to act as a "super expert," able to overrule the diverse views of the experts from whom input is elicited. It is not prudent to generate an apparent consensus unless consensus among the experts is really achieved. It is not necessary that the TFI agree with the outcome of the process; the TFI can stand behind that outcome as the result of thorough interaction among experts.

The issue of breadth. The statement on breadth of application on page 1-7 of the draft report and other statements related to the intended breadth of application of the recommended methodology are the cause of much uneasiness among the Panel. *A clear statement of the purpose and scope of the report should be included early in the introduction.*

- It should be made clear that the recommended methodology is based on a study of the experiences with LLNL and EPRI procedures. This should be brought out in the history-context material called for above. In the appropriate places, specific references to the lessons learned by examination of previous PSHA projects should be cited. The studies from which the recommended methodology was derived should be clearly described, even though the intent of the report is not to address the reconciliation of the LLNL/EPRI studies. The reader should be made

aware of the lessons learned from the evaluation of those (and other?) studies that have gone into the formulation of this report. The reader should be told explicitly that alternate PSHA approaches were not assimilated and that this report is not based on a consensus of a broad sample of practioneers.

- Some statement of costs would be in order. What a hazard evaluation can deliver is often a matter of how many dollars are available. Cost estimates may be beyond SSHAC's scope, but even this could be mentioned.
- The Panel anticipates that the full procedure recommended in this report will not be applied to the seismic regulation of all critical facilities. It is not a general methodology that will be applied step-by-step in all situations. Therefore, criteria or guidelines are needed in the report, to aid the project sponsor and the PSHA analysts in deciding when the full procedure is justified. A statement is needed about what can be delivered with different levels of PSHA, so the buyer can make an informed decision as to what will and will not be produced. As stated above, the analyst who chooses for sufficient reasons to use other procedures should not be put by this report in a position of having to defend that decision in a regulatory setting. He or she, of course, must be prepared to defend the procedures that were adopted.

SCIENTIFIC AND TECHNICAL CONCERNS

The Panel questions whether the links between SSHAC's recommended methodology and its applications are spelled out in sufficient clarity. Although SSHAC is not charged with specifying the use of hazard numbers in engineering design, a brief treatment is needed pointing to how the results can be used, and, in particular, what the knowledge of highly refined uncertainty estimates contributes to applications. A clear and unequivocal definition of aleatory and epistemic uncertainty is needed, as well as a clear and readily applied prescription

for separating the two. This is needed because of the emphasis on this subject in the report.

Although not as yet the subject of full panel evaluation, the following example illustrates the need for SSHAC to be very clear on the value and the method of application of their categorization of uncertainty. "What should count for decision is not the aleatory/epistemic distinction, but the temporal variation in the total uncertainty (in the total or predictive distribution of A_T, maximum peak ground acceleration and spectral values at the site in the next T years) during the lifetime of the project." According to this viewpoint:

- There is no need to label uncertainty as epistemic or aleatory.
- If one sees total uncertainty as being contributed by different sources (e.g., by uncertainty on model type or on various parameters), then it is reasonable to expect that the uncertainty associated with each source will evolve in its own way in time. Making a binary distinction between epistemic and aleatory uncertainty corresponds to assuming that each source will be either explained totally (epistemic components) or will remain constant over the lifetime of the system (aleatory components.)
- One can formulate rational ways to make decisions accounting for the possible temporal evolution of uncertainty. The Panel member responsible for these comments is not, on the other hand, aware of any convincing method to make decisions based on the aleatory/epistemic decomposition. The amount of conservatism displayed by decisions under time-varying uncertainty depends on the nature of the problem (essentially on the degrees of asymmetry in the rewards and penalties associated, respectively, with future possible decreases and increases in the calculated risks).

The SSHAC report will be strengthened by addressing these concerns in a straightforward way.

Intensity data from historic strong earthquakes in the central and eastern United States is not incorporated in the ground motion models. The relation between m_{bLg} and int[illegible]he eastern United States, first established by Nuttli, should not b[illegible]

"Seismic source zones", a key concept in the prescribed source characterization procedure, should be explicitly recognized as an artificial construct introduced to make hazard calculations tractable. They are not real physical entities.

NRC FORM 335 (2-89) NRCM 1102, 3201, 3202

U.S. NUCLEAR REGULATORY COMMISSION

BIBLIOGRAPHIC DATA SHEET

(See instructions on the reverse)

1. REPORT NUMBER (Assigned by NRC, Add Vol., Supp., Rev., and Addendum Numbers, If any.)
NUREG/CR-6372
UCRL-ID-122160
Vol. 1

2. TITLE AND SUBTITLE
Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts

Main Report

3. DATE REPORT PUBLISHED
MONTH: April
YEAR: 1997

4. FIN OR GRANT NUMBER
L2503

5. AUTHOR(S)
Senior Seismic Hazards Analysis Committee (SSHAC)
R.J. Budnitz(Chairman, G. Apostolakis, D.M. Boore, L.S. Cluff, K.J. Coppersmith, C.A. Cornell, P.A. Morris

6. TYPE OF REPORT

7. PERIOD COVERED *(Inclusive Dates)*

8. PERFORMING ORGANIZATION - NAME AND ADDRESS *(If NRC, provide Division, Office or Region, U.S. Nuclear Regulatory Commission, and mailing address; if contractor, provide name and mailing address.)*
SSHAC under contract to:
Lawrence Livermore National Laboratory
Livermore, CA 94550

9. SPONSORING ORGANIZATION - NAME AND ADDRESS *(If NRC, type "Same as above"; if contractor, provide NRC Division, Office or Region, U.S. Nuclear Regulatory Commission, and mailing address.)*
Division of Engineering Technology
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001

10. SUPPLEMENTARY NOTES

11. ABSTRACT *(200 words or less)*

Probabilistic Seismic Hazard Analysis (PSHA) is a methodology that estimates the likelihood that various levels of earthquake-caused ground motion will be exceeded at a given location in a given future time period. Due to large uncertainties in all the geosciences data and in their modeling, multiple model interpretations are often possible. This leads to disagreement among experts, which in the past has led to disagreement on the selection of ground motion for design at a given site. The Senior Seismic Hazards Analysis Committee (SSHAC) reviewed past studies, including the Lawrence Livermore National Laboratory and the EPRI landmark PSHA studies of the 1980's and examined ways to improve on the present state-of-the-art. The Committee's most important conclusion is that differences in PSHA results are due to procedural rather than technical differences. Thus, in addition to providing a detailed documentation on state-of-the-art elements of a PSHA, this report provides a series of procedural recommendations. The role of experts is analyzed in detail. Two entities are formally defined – the Technical Integrator (TI) and the Technical Facilitator Integrator (FI) – to account for the various levels of complexity in the technical issues and different levels of efforts needed in a given study.

12. KEY WORDS/DESCRIPTORS *(List words or phrases that will assist researchers in locating the report.)*
Probabilistic Seismic Hazard Analysis (PSHA), earthquakes, geosciences data, ground motion,

13. AVAILABILITY STATEMENT
unlimited

14. SECURITY CLASSIFICATION
(This Page)
unclassified
(This Report)
unclassified

15. NUMBER OF PAGES

16. PRICE

Federal Recycling Program

UNITED STATES
NUCLEAR REGULATORY COMMISSION
WASHINGTON, DC 20555-0001

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE, $300

SPECIAL STANDARD MAIL
POSTAGE AND FEES PAID
USNRC
PERMIT NO. G-67

120555139531 1 1AN1RA
US NRC-OIRM
PUBLICATIONS BRANCH
TPS-PDR-NUREG
2WFN-6E7
WASHINGTON DC 20555