



NUREG/CR-7214

Toward a More Risk-Informed and Performance-Based Framework for the Regulation of the Seismic Safety of Nuclear Power Plants

Office of Nuclear Regulatory Research

AVAILABILITY OF REFERENCE MATERIALS IN NRC PUBLICATIONS

NRC Reference Material

As of November 1999, you may electronically access NUREG-series publications and other NRC records at NRC's Library at www.nrc.gov/reading-rm.html. Publicly released records include, to name a few, NUREG-series publications; *Federal Register* notices; applicant, licensee, and vendor documents and correspondence; NRC correspondence and internal memoranda; bulletins and information notices; inspection and investigative reports; licensee event reports; and Commission papers and their attachments.

NRC publications in the NUREG series, NRC regulations, and Title 10, "Energy," in the *Code of Federal Regulations* may also be purchased from one of these two sources.

1. The Superintendent of Documents

U.S. Government Publishing Office
Mail Stop IDCC
Washington, DC 20402-0001
Internet: bookstore.gpo.gov
Telephone: (202) 512-1800
Fax: (202) 512-2104

2. The National Technical Information Service

5301 Shawnee Rd., Alexandria, VA 22312-0002
www.ntis.gov
1-800-553-6847 or, locally, (703) 605-6000

A single copy of each NRC draft report for comment is available free, to the extent of supply, upon written request as follows:

Address: **U.S. Nuclear Regulatory Commission**
Office of Administration
Publications Branch
Washington, DC 20555-0001
E-mail: distribution.resource@nrc.gov
Facsimile: (301) 415-2289

Some publications in the NUREG series that are posted at NRC's Web site address www.nrc.gov/reading-rm/doc-collections/nuregs are updated periodically and may differ from the last printed version. Although references to material found on a Web site bear the date the material was accessed, the material available on the date cited may subsequently be removed from the site.

Non-NRC Reference Material

Documents available from public and special technical libraries include all open literature items, such as books, journal articles, transactions, *Federal Register* notices, Federal and State legislation, and congressional reports. Such documents as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings may be purchased from their sponsoring organization.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at—

The NRC Technical Library

Two White Flint North
11545 Rockville Pike
Rockville, MD 20852-2738

These standards are available in the library for reference 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—

American National Standards Institute

11 West 42nd Street
New York, NY 10036-8002
www.ansi.org
(212) 642-4900

Legally binding regulatory requirements are stated only in laws; NRC regulations; licenses, including technical specifications; or orders, not in NUREG-series publications. The views expressed in contractor-prepared publications in this series are not necessarily those of the NRC.

The NUREG series comprises (1) technical and administrative reports and books prepared by the staff (NUREG-XXXX) or agency contractors (NUREG/CR-XXXX), (2) proceedings of conferences (NUREG/CP-XXXX), (3) reports resulting from international agreements (NUREG/IA-XXXX), (4) brochures (NUREG/BR-XXXX), and (5) compilations of legal decisions and orders of the Commission and Atomic and Safety Licensing Boards and of Directors' decisions under Section 2.206 of NRC's regulations (NUREG-0750).

DISCLAIMER: This report was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any employee, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for any third party's use, or the results of such use, of any information, apparatus, product, or process disclosed in this publication, or represents that its use by such third party would not infringe privately owned rights.



United States Nuclear Regulatory Commission

Protecting People and the Environment

NUREG/CR-7214

Toward a More Risk-Informed and Performance-Based Framework for the Regulation of the Seismic Safety of Nuclear Power Plants

Manuscript Completed: September 2013

Date Published: May 2016

Prepared by:

Robert J. Budnitz

Energy Geosciences Division

Lawrence Berkeley National Laboratory

University of California

Berkeley, CA 94720

Michael W. Mieler*

Department of Civil and Environmental Engineering

University of California

Berkeley, CA 94720

*Now at Department of Civil Engineering

Johns Hopkins University

Baltimore, MD 21218

NRC Technical Monitors

Scott P. Stovall

Richard F. Rivera-Lugo

Annie Kammerer

NRC Job Code V6159

Office of Nuclear Regulatory Research

ABSTRACT

Dozens of seismic probabilistic risk assessments (SPRAs) exist that have analyzed large nuclear power plants (NPPs). A major insight from these SPRAs is that, although generally the plants are adequately safe against earthquake threats, the way the industry and the United States Nuclear Regulatory Commission (NRC) currently design, build, operate, analyze, and regulate seismic NPP safety may not be optimal. This suboptimal situation means that for both operating NPPs and new NPP designs not yet built, they may fail to take advantage of possible additional safety insights and improvements, may be more difficult to analyze and to regulate than they need to be, and may cost more than they otherwise would. There is room for improvement in several areas. The major topics covered herein are the variations in residual seismic risk from plant to plant; the unbalanced risk profiles and incomplete defense in depth achieved at many plants; the impact of structures and components being designed individually for seismic performance rather than taking a systems approach; the variations in margin that exist among the industry consensus codes and standards for seismic design and analysis; and the observation that the design codes and the NRC regulations do not work together well. Each of these topics is analyzed in detail, and case-study examples based on the seismic risk profiles of two operating plants are used to illustrate the issues. Suggestions are presented that would improve the seismic framework in each of these areas.

FOREWORD

The staff believes that the current approach used by the U.S. nuclear industry to design, build, operate, and analyze large nuclear power plants to achieve seismic safety, and the approach used by the U.S. Nuclear Regulatory Commission to regulate seismic safety are adequate. Traditional engineering analyses and dozens of seismic probabilistic risk assessments that have analyzed the seismic safety of nuclear power plants both in the U.S. and abroad support this conclusion. Design approaches, analysis methods, regulations, and other regulatory positions that support the seismic safety of the NPP fleet were mostly developed many years ago, and in a number of ways they do not reflect modern engineering design philosophies that use more risk-informed thinking and more performance-based approaches. This has been recognized broadly in the community of seismic-safety experts and indeed some changes have already been implemented to modify both the industry's design and analysis approaches and the NRC's regulatory framework to rely more on performance-based and risk-informed ideas.

The NRC seismic research has been addressing improved estimates of seismic hazards as well as the capability of SSCs to perform their safety functions for seismic events greater than their design basis. Risk-informed and performance-based approaches developed under this program with concurrent industry and Department of Energy research resulted in methods that provide a thorough overall picture of a plant seismic capacity that can and are being used for seismic reevaluations. Those results also informed advances in the design and review of new plants.

To provide a technical analysis of the opportunities for further reliance on risk-informed and performance-based concepts in relation to seismic safety of new designs, the NRC's Office of Nuclear Regulatory Research initiated a project to review the existing regulatory framework related to seismic safety and to develop options for advances that, if further researched and implemented, could provide added efficiency and clarity. Besides possible new safety insights, this would enable, for example, focusing resources promptly where they matter the most for safety so that regulations could be met with overall fewer resources, and could lead to designs that are overall less costly to design, construct, operate and maintain. This report describes the results of that review and analysis. Because the number of technical areas involved is large and because they interact, no simple approach is likely to emerge. Rather, advances will likely need to account for a complex interaction among several different technical issues, which would need to be considered together.

To obtain the necessary insights on possible options and related advances, the study covered two major topics. The first topic deals with the variations in residual seismic risk from plant to plant, risk balance profiles, and related defense in depth aspects achieved at various plants. The second topic addresses the impact of designing structures and components individually for seismic performance rather than taking a systems approach as well as consideration of the variations in margin that exist among the industry consensus codes and standards for seismic design and analysis. The study analyzed each of these topics and used case-study examples based on the seismic risk profiles of two plants to illustrate the issues.

The report provides the analysis and the case study examples as well as options for advances that could provide added clarity and efficiency to the regulatory framework in relation to each of these two major topics. It is anticipated that these analysis and options for advances will inform, for example, research and other staff activities as they relate to risk-informed regulatory improvements within the current regulatory framework, as well as staff interactions with standards development organizations pursuing the development of risk-informed and performance-based seismic design and analysis standards for nuclear installations.

TABLE OF CONTENTS

ABSTRACT	iii
FOREWORD	v
LIST OF FIGURES	ix
LIST OF TABLES	xi
EXECUTIVE SUMMARY	xiii
ACKNOWLEDGMENTS	xvii
ABBREVIATIONS	xix
1 INTRODUCTION	1
2 PROBLEM STATEMENT	3
2.1 Variations in Residual Seismic Risk	3
2.2 Unbalanced Seismic Risk Profiles	3
2.3 SSCs Are Designed Individually for Seismic Performance	3
2.4 Variations in Margin Among Industry Design Codes	4
2.5 Coordination Between Design Codes and NRC Regulations	4
3 ATTRIBUTES OF AN IDEAL FRAMEWORK FOR ASSURING SEISMIC SAFETY	7
3.1 High Level Regulatory Philosophy and Approach for Seismic Safety	8
3.2 Seismic Input	9
3.3 Design Guidance	10
3.4 Analysis Guidance	11
4 COMPARISONS OF VARIOUS ATTRIBUTES TO THE IDEAL	13
4.1 Summary of Evaluation of Framework Attributes	13
4.2 In-depth Evaluation of Framework Attributes	19
4.2.1 High Level Regulatory Philosophy and Approach for Seismic Safety	19
4.2.2 Seismic Input	22
4.2.3 Design Guidance	23
4.2.4 Analysis Guidance	25
5 A PROPOSAL FOR INCREMENTAL PROGRESS: DEFENSE IN DEPTH	27
5.1 Step A: Complete a New NPP Design	28
5.2 Step B: Perform a PRA that Includes a Full SPRA	28
5.3 Step C: Identify Leading Accident Sequences and SSCs	28
5.4 Step D: Compare the Risk Profile to a Safety Target	29
5.5 Step E: Evaluation for an Imbalance in Defense in Depth	30
5.6 Implementation	32
5.6.1 Upgrading Options	33
5.6.2 Using Tradeoff Analysis Tools	33
5.6.3 An Iterative Process	33
5.6.4 Issues with a Firm Decision Criterion	33
6 OPPORTUNITIES FOR FURTHER ADVANCES – PATHS FORWARD OVER THE LONGER TERM	35
6.1 Core Damage vs. Large Early Release	35
6.2 Confidence Level	35
6.3 Providing an Option to Use a Performance Target	36
6.4 Systems-analysis Issues	36
6.5 SSCs Are Designed Individually for Seismic Performance	37
6.6 Variations in Margin Among Industry Design Codes	38
6.7 Broad Institutional Barriers	38

	6.7.1	NRC Regulations.....	38
	6.7.2	Code Committees.....	39
	6.7.3	NRC and Industry Working Together.....	39
	6.8	Relative Roles of Risk from Earthquakes vs. Other Initiators.....	40
7		CASE STUDIES	41
	7.1	Introduction.....	41
	7.2	Background for the Case Studies.....	43
		7.2.1 Plant A.....	44
		7.2.2 Plant B.....	48
		7.2.3 Summary.....	54
	7.3	Analysis.....	55
		7.3.1 ASCE 43-05 Configuration.....	55
		7.3.2 Enhanced Balance Configuration.....	62
	7.4	Implications.....	68
8		SUGGESTIONS FOR A PATH FORWARD	73
	8.1	Summary of Issues and Potential Paths Forward.....	73
		8.1.1 A Balanced Risk Profile and Defense in Depth.....	73
		8.1.2 Core Damage vs. Large Early Release.....	74
		8.1.3 Confidence Level.....	74
		8.1.4 Providing an Option to Use a Performance Target.....	74
		8.1.5 Systems-analysis Issues.....	74
		8.1.6 SSCs Are Designed Individually for Seismic Performance.....	75
		8.1.7 Variations in Margin Among Industry Design Codes.....	75
		8.1.8 Broad Institutional Barriers.....	75
	8.2	Conclusion.....	76
9		REFERENCES	77

LIST OF FIGURES

Figure 5-1	Seismic fragility curves and HCLPF point	30
Figure 7-1	Mean seismic hazard curves for both case study sites.....	43
Figure 7-2	Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the baseline configuration of Plant A. The ordinate is a frequency in units of 10^{-5} per year.	46
Figure 7-3	Effect of strengthening ELEC1 on the seismic risk profile of baseline configuration of Plant A.....	47
Figure 7-4	Effect of strengthening STRUC2 on the seismic risk profile of baseline configuration of Plant A.....	48
Figure 7-5	Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the baseline configuration of Plant B. The ordinate is a frequency in units of 10^{-5} per year.	52
Figure 7-6	Effect of strengthening ELEC14 on the seismic risk profile of baseline configuration of Plant B.....	53
Figure 7-7	Effect of strengthening STRUC4 on the seismic risk profile of baseline configuration of Plant B.....	54
Figure 7-8	Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the “ASCE 43-05” configuration of Plant A. The ordinate is a frequency in units of 10^{-5} per year.....	58
Figure 7-9	Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the “ASCE 43-05” configuration of Plant B. The ordinate is a frequency in units of 10^{-5} per year.....	61
Figure 7-10	Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the “enhanced balance” configuration of Plant A. The ordinate is a frequency in units of 10^{-5} per year.....	64
Figure 7-11	Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the “enhanced balance” configuration of Plant B. The ordinate is a frequency in units of 10^{-5} per year	67
Figure 7-12	Seismic risk profile for each of the three configurations of Plant A	69
Figure 7-13	Seismic risk profile for each of the three configurations of Plant B	70

LIST OF TABLES

Table 4-1	Status and evaluation of the various attributes of today's Framework for seismic design, analysis, and regulation of SSCs in NPPs.....	14
Table 7-1	Properties of SSCs in the baseline configuration of Plant A	44
Table 7-2	Properties of dominant seismic accident sequences in the baseline configuration of Plant A.....	45
Table 7-3	Properties of SSCs in the baseline configuration of Plant B	49
Table 7-4	Properties of dominant seismic accident sequences in the baseline configuration of Plant B.....	51
Table 7-5	Changes in seismic capacities of SSCs in the "ASCE 43-05" configuration of Plant A (see Table 7-1 for SSCs not listed).....	56
Table 7-6	Changes in properties of dominant seismic accident sequences in the "ASCE 43-05" configuration of Plant A.....	57
Table 7-7	Changes in seismic capacities of SSCs in the "ASCE 43-05" configuration of Plant B (see Table 7-3 for SSCs not listed).....	59
Table 7-8	Changes in properties of dominant seismic accident sequences in the "ASCE 43-05" configuration of Plant B.....	59
Table 7-9	Changes in seismic capacities of SSCs in the "enhanced balance" configuration of Plant A (see Table 7-1 for SSCs not listed).....	62
Table 7-10	Changes in properties of dominant accident sequences in the "enhanced balance" configuration of Plant A	63
Table 7-11	Changes in seismic capacities of SSCs in the "enhanced balance" configuration of Plant B (see Table 7-3 for SSCs not listed).....	65
Table 7-12	Changes in properties of dominant seismic accident sequences in the "enhanced balance" configuration of Plant B.....	66

EXECUTIVE SUMMARY

In the literature, there are dozens of seismic probabilistic risk assessments (SPRAs) that have analyzed nuclear power plants (NPPs), including not only plants in the United States but also NPPs abroad with designs similar to those of U.S. NPPs. An SPRA is a probabilistic analysis of the seismic safety of an NPP whose bottom-line result is the overall risk profile of the plant. Here the “risk profile” means (i) the overall annual frequency of earthquake-induced core-damage accidents and earthquake-induced large-early-release accidents, (ii) the identification of the most important accident sequences that could produce such damage and the relative contributions of each, and (iii) the identification of the most important structures, systems, and components (SSCs) whose failures contribute to those sequences. An SPRA is, by design, intended to be as realistic an assessment as feasible, not an assessment that embeds unnecessary conservatism or other distortions. It is also intended to analyze and understand the overall uncertainties associated with each of the three major results that comprise the risk profile.

The general finding about seismic safety from the body of SPRAs in the literature, and also from traditional engineering analyses, is that U.S. NPPs (as well as similar plants abroad) are generally adequately safe; the U.S. plants generally meet all applicable NRC seismic regulations with adequate margin, including meeting the NRC’s safety goals with adequate margin. (This broad statement should not be taken, however, as applying to any individual NPP, whose seismic safety may or may not comport with the overall performance of the larger fleet.) Overall, accidents initiated by earthquakes are often important contributors to the total residual risk, and the SPRAs provide important information as to which sequences contribute most to risk and why. SPRA results indicate that the dominant contributors to the seismic risk can vary substantially from one plant to the next; and that a significant fraction of the important accident sequences involve non-seismic failures and/or human errors in addition to failures caused by the earthquake. That is, these sequences would not lead to an accident in the absence of the non-seismic failures and/or human errors. The SPRA results also demonstrate that only a small fraction of the seismic sequences leading to a core-damage accident are associated with a large early release of radioactivity, which is similar to the finding for sequences that start with other categories of initiating events.

Another insight from these SPRAs, supplemented by other knowledge gained over the past decade or so, is that the way the industry and the NRC currently design, build, operate, analyze, and regulate seismic NPP safety is not optimal. In part, this is because the design approaches, analysis methods, regulations, and other regulatory positions that support the seismic safety of the NPP fleet were mostly developed many years ago, and in a number of ways they do not reflect modern engineering philosophies that use more risk-informed thinking and more performance-based approaches. This has been recognized broadly in the community of seismic-safety experts for many years. Indeed, some changes have already been implemented to modify both the industry’s design and analysis approaches and the NRC’s regulatory framework to rely more on risk-informed and performance-based ideas.

Despite these changes, there is still lots of room for improvement in several areas, which, if taken advantage of, could make operating NPPs safer, more efficient and economic, and better understood. Increasing understanding means that increased safety can be objectively demonstrated. For new NPPs not yet built, this suboptimal situation means that although the new designs are generally adequately safe against earthquakes with a high degree of margin,

the designs may fail to take advantage of possible additional safety insights and improvements, are more difficult to analyze and to regulate than they need to be, and may cost more than they otherwise would.

In summary, although the current approach has produced a fleet of operating NPPs and also a group of new designs that are generally demonstrably safe against earthquakes with adequate margin, the approach has in many areas failed to take advantage of several modern approaches to designing, analyzing, and regulating NPP seismic safety that could make the NPPs demonstrably safer, as well as more efficient.

To provide insight into possible advances by evaluating both their potential and their limitations, Lawrence Berkeley National Laboratory (LBNL), supported by the NRC's Office of Nuclear Regulatory Research, has reviewed the existing framework and has developed options for advances that, if implemented, could improve that framework. This report describes the results of that review. One of this report's major objectives is to provide detailed analyses of several technical and policy issues as input to future deliberations by not only the NRC (Commissioners and staff) but also the regulated industry and the public.

Some of the policy initiatives suggested in this report overlap with or coincide with other industry and NRC initiatives that are currently in the works or under active debate. These initiatives are moving seismic design and assessment of NPPs toward a more performance-based approach that accounts for risk information wherever practical. What is apparent from a review of the current state of practice and development activities as discussed in this report is that in the technical arena concerned with the safety of large NPPs against earthquakes, the situation is already more advanced than in many other safety areas, and is therefore closer to being "ripe" for even further advances. These advances could pave the way to the broader advances in other technical areas that are now under discussion.

The body of this report begins with an exploration of several technical issues and current policy positions that could benefit from modern thinking. It then describes the attributes of an ideal framework for seismic NPP safety. This is followed by a comparison of today's situation with that ideal.

This is followed in turn by a discussion of a proposed path forward that in the relatively short term (a few years) can provide some useful advances and benefits. A longer-term path forward is then described that could lead to even greater benefits in design, analyzability, and regulation and that could positively affect both safety and cost. Next, the report contains case-study examples from two currently operating NPPs that demonstrate many of the insights laid out in the earlier parts of this report. Finally, the report presents key findings and suggestions for the path forward.

As noted, there is room for improvement in several areas. The major topics covered herein are the variations in residual seismic risk from plant to plant; the unbalanced risk profiles and incomplete defense in depth achieved at many plants; the impact of structures and components being designed individually for seismic performance rather than taking a systems approach; the variations in margin that exist among the industry consensus codes and standards for seismic design and analysis; and the observation that the design and analysis codes and the NRC regulations do not work as well together as they could. Each of these topics is analyzed in detail, and case-study examples based on the seismic risk profiles of two operating plants are used to illustrate the issues. Suggestions are presented that would improve the seismic framework in each of these areas.

A major conclusion of the review is that because the number of technical areas involved is large and because they interact, no simple approach is likely to emerge. Rather, advances will likely need to account for a complex interaction among several different technical issues, which would need to be considered together.

ACKNOWLEDGMENTS

The authors wish to acknowledge the support of Annie Kammerer, Richard F. Rivera-Lugo, and Scott P. Stovall of the United States (US) Nuclear Regulatory Commission's (NRC) Office of Nuclear Regulatory Research, who served as the technical contacts on the NRC staff, and Jose Pires of the staff, who provided very valuable input and assistance during the final stages of preparing this report.

This manuscript has been written by authors at the Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy (DOE). The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes. This project was supported by the U.S. Nuclear Regulatory Commission under a Federal Interagency Agreement with DOE. While this document is believed to contain correct information, neither the U.S. Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California. Lawrence Berkeley National Laboratory is an equal opportunity employer.

ABBREVIATIONS

ACI	American Concrete Institute
ANS	American Nuclear Society
ASCE	American Society of Civil Engineers
ASME	American Society of Mechanical Engineers
CDF	core-damage frequency
CFR	US Code of Federal Regulations
DBE	design-basis earthquake
DOE	US Department of Energy
HCLPF	high-confidence-of-a-low-probability-of-failure
IEEE	Institute of Electrical and Electronic Engineers
ISRS	in-structure response spectrum
LBNL	Lawrence Berkeley National Laboratory
LERF	large-early-release frequency
LWR	light water reactor
N/A	not applicable
NPP	nuclear power plant
NRC	US Nuclear Regulatory Commission
PRA	probabilistic risk assessment
PSHA	probabilistic seismic hazard analysis
R&D	research and development
SMA	seismic margin assessment
SPRA	seismic probabilistic risk assessment
SSC	structure, system, or component
SSE	safe shutdown earthquake
SSI	soil-structure interaction

1 INTRODUCTION

In the literature, there are dozens of seismic probabilistic risk assessments (SPRAs) that have analyzed nuclear power plants (NPPs), including not only plants in the United States (US) but also NPPs abroad with designs similar to those of US NPPs. An SPRA is a probabilistic analysis of the seismic safety of an NPP whose bottom-line result is the overall risk profile of the plant. Here the “risk profile” means (i) the overall annual frequency of earthquake-induced core-damage accidents and earthquake-induced large-early-release accidents, (ii) the identification of the most important accident sequences that could produce such damage and the relative contributions of each, and (iii) the identification of the most important structures, systems, and components (SSCs) whose failures contribute to those sequences. An SPRA is, by design, intended to be as realistic an assessment as feasible, not an assessment that embeds unnecessary conservatism or other distortions. It is also intended to analyze and understand the overall uncertainties associated with each of the three major results that comprise the risk profile.

The general finding about seismic safety from the body of SPRAs in the literature is that US NPPs, as well as similar plants abroad, are generally adequately safe; the US plants generally meet all applicable NRC seismic regulations with adequate margin, including meeting the NRC’s safety goals (NRC 1986) with adequate margin¹. Overall, accidents initiated by earthquakes are often important contributors to total risk, and the SPRAs provide important information as to which sequences contribute most to risk and why. SPRA results indicate that the dominant contributors to the seismic risk can vary substantially from one plant to the next; and that a significant fraction of the important accident sequences involve non-seismic failures and/or human errors in addition to failures caused by the earthquake. That is, these sequences would not lead to an accident in the absence of the non-seismic failures and/or human errors (NRC 2001). The SPRA results also demonstrate that only a small fraction of the seismic sequences leading to a core-damage accident are associated with a large early release of radioactivity, which is similar to the finding for sequences that start with other categories of initiating events (NRC 1990).

Another insight from these SPRAs, supplemented by other knowledge gained over the past decade or so, is that the way the industry and the NRC currently design, build, operate, analyze, and regulate seismic NPP safety is not optimal. This suboptimal situation means that there is room for improvement in several areas, which, if taken advantage of, could make operating NPPs safer, more efficient and economic, and better understood. Increasing understanding means that increased safety can be objectively demonstrated. For new NPPs not yet built, this suboptimal situation means that although the new designs are generally adequately safe against earthquakes with a high degree of margin, the designs may fail to take advantage of possible additional safety insights and improvements, are more difficult to analyze and to regulate than they need to be, and may cost more than they otherwise would.

In summary, although the current approach has produced a fleet of operating NPPs and also a group of new designs that are demonstrably safe against earthquakes with adequate margin,

¹ This broad statement should not be taken, however, as applying to any individual NPP, whose seismic safety may or may not comport with the overall performance of the larger fleet.

the approach has generally failed to take advantage of several modern approaches to designing, analyzing, and regulating NPP seismic safety that could make the NPPs demonstrably safer, as well as more efficient.

Some of the policy initiatives suggested in this report overlap with or coincide with other industry and NRC initiatives that are currently in the works or under active debate. These initiatives are moving seismic design and assessment of NPPs toward a more performance-based approach that accounts for risk information wherever practical. What is apparent from a review of the current state of practice and development activities as discussed in this report is that in the technical arena concerned with the safety of large NPPs against earthquakes, the “Framework” is already more advanced than in many other safety areas, and is therefore closer to being “ripe” for even further advances. These advances could pave the way to the broader advances in other technical areas that are now under discussion.

In this report, which will concentrate on the seismic safety of modern NPPs, Section 2 explores several technical issues and current policy positions that could benefit from modern thinking. Section 3 then discusses the attributes of an ideal Framework for seismic NPP safety. This is followed by a comparison in Section 4 of today’s situation with that ideal. Section 5 discusses a proposed path forward that in the relatively short term (a few years) can provide some useful advances and benefits. A longer-term path forward that could lead to even greater benefits in design, analyzability, and regulation and that could positively affect both safety and cost is discussed in Section 6. Section 7 contains case-study examples from two currently operating NPPs that demonstrate many of the insights laid out in the earlier parts of this report. Finally, Section 8 presents key findings and suggestions for the path forward.

2 PROBLEM STATEMENT

This Section explores several technical issues and current policy positions that will be addressed in subsequent Sections of this report.

2.1 Variations in Residual Seismic Risk

The risk profiles of the many nuclear power plants (NPPs) analyzed with seismic probabilistic risk assessments (SPRAs) show a lot of variation from one plant to the next in terms of the overall seismic-induced core-damage frequency (seismic CDF). They also reveal a wide variation in the overall seismic margin of the plants above their individual seismic design bases (NRC 2001). “Seismic margin” is defined herein as the magnitude of ground motion the plant can withstand, beyond the ground motion represented by the design basis, before it starts to get into trouble vis-à-vis accident sequences leading to core damage².

2.2 Unbalanced Seismic Risk Profiles

The risk profiles from the SPRAs also show significant variation in how the various NPPs respond to earthquake ground motion, particularly as it relates to the structures, systems, or components (SSCs) that contribute the most to calculated seismic CDF. For many of the NPPs, the seismic risk profile is unbalanced in the sense that the seismic failure of a single SSC, or of a very small number of SSCs, contributes disproportionately and dominates the risk profile. See Section 4.2.1(F) and Section 5 for further discussion. Section 7 provides some actual case-study examples that illustrate this. Such an unbalanced risk profile means that the plant does not possess as much defense in depth as is typically desired. That is, a plant with strong defense-in-depth attributes would not possess a risk profile dominated by a single failure or a very small number of failures. Rather, an objective of the defense-in-depth philosophy is an NPP design in which the risk is more balanced among a variety of different contributors to overall risk.

2.3 SSCs Are Designed Individually for Seismic Performance

Using today’s approach, each SSC within an NPP is designed individually to achieve an adequate seismic capacity and performance. Both the design codes and the NRC’s regulations (NRC 1977; NRC 1996) take this approach for every SSC that is determined to be “safety related” and thus to need a specific design for earthquake loads³. Insufficient account is taken of the role of a given safety-related SSC in contributing to the overall seismic safety of the plant as a system. The systems view of the plant does not generally play a strong enough role in how

² The “seismic margin” is often defined in terms of the High Confidence of Low Probability of Failure (HCLPF) capacity, which is defined as the ground motion level at which there is about 95 percent confidence that the likelihood of overall failure is about 5 percent (NRC 1985). The HCLPF capacity turns out to be about equal to the ground motion level at which there is about a 1 percent overall likelihood of failure, accounting for all of the uncertainties in the analysis of the likelihood of failure. The HCLPF capacity concept can be used for an individual SSC, for an individual accident sequence, or for the seismic capacity of the plant-as-a-whole (NRC 2012c).

³ The term “safety-related” has a specific meaning in NRC regulations. In this report, the term is used more loosely to designate those SSCs that play a role in important accident sequences and hence merit special attention by the designer, in contrast to the large number of other SSCs in any NPP that do not play such a role. Throughout this report, when the design and analysis of SSCs is discussed, the subset of SSCs that are safety-related SSCs is intended.

individual SSCs are now designed and regulated against earthquakes. This is discussed further in Section 4.2.4, especially in (Q), and in Section 6.5.

2.4 Variations in Margin Among Industry Design Codes

The design of any individual SSC for service in a nuclear power plant is generally governed by some industry code or standard that has usually been endorsed by the NRC for its particular application. These consensus codes are developed and maintained by a number of code committees organized under different standards-development organizations (e.g., American Nuclear Society, American Concrete Institute, American Society of Civil Engineers, Institute of Electrical and Electronic Engineers, and others). These codes, which generally rely on an externally-specified design-basis ground motion as the starting point for the design, use a variety of different approaches to address the question of how much seismic margin above the design basis should be embedded in an SSC designed using the code. That there are differences in embedded margin for SSCs designed using various codes is not surprising, given that the code committees have generally worked independently and that the consensus codes represent different philosophies of design that exist in the different fields of engineering. For example, it would only be through serendipity that the design of electrical components against earthquake ground motions and the design of concrete shear walls against earthquake ground motions would have taken similar approaches to embedded margin, given how different the design problems and design solutions are, and how differently the code committees in those areas went about developing the requirements. Additionally, not all objectives or target margin levels are explicitly stated. This is discussed further in Section 4.2.3(K) and Section 6.6.

This is not a criticism of the work of the various code committees. It is merely an observation that there are differences in approaches and hence in the outcomes. Absent intervention by an outside entity (for example, a regulatory agency like the NRC) that could impose some consistency of approach, consistency of outcomes in terms of embedded seismic margin would not be expected. As a result of the history of code development to date, that consistency has not existed and current does not exist.

2.5 Coordination Between Design Codes and NRC Regulations

NRC's overall regulatory scheme for seismic safety involves much more than criteria for design: it also provides criteria for construction, maintenance, operations, analyses (both analyses for regulatory compliance and realistic analyses to understand performance), and inspections. In an ideal NPP regulatory framework, the design of individual safety-related SSCs against earthquake loads would be consistent with and support the NRC's safety regulatory scheme to achieve an overall NPP design with more than adequate safety margin. Judging the adequacy of that margin would specifically fall within the purview of the NRC, and ideally the NRC would specify the desired margin and the desired confidence level, rather than leaving the choice of target margin to each of the various code committees.

Specifically, in an ideal Framework, an individual safety-related SSC designed according to a given code (say, a tank designed to an ASME code or an electrical component designed to an IEEE code) would use the Design Basis Earthquake (DBE) ground motion defined by the NRC as input. Each code embeds a certain amount of seismic margin in the design, for reasons noted in Section 2.4 and elsewhere in this report, so that only a ground motion somewhat larger than the DBE ground motion can cause even the beginning signs of failure of the SSC. Only a ground motion much larger than the design ground motion would have a high probability of causing the SSC to fail to perform its required safety function. Certain construction,

maintenance, inspection, and analysis requirements would also be seamlessly integrated and specified to assure that safety objectives are met. The result would be an SSC whose seismic capacity would meet (or exceed) an NRC-established target capacity needed for overall NPP safety. In this approach, an industry consensus design code and the NRC guidance and regulations would work together to achieve a prescribed safety target set by the NRC.

This approach is not in place today, as discussed further in Section 6.7. This is because historically the NRC has neither had nor used explicit quantitative safety targets, nor was an attempt made by the NRC historically to bring about any measure of uniformity in the amount of margin between the seismic design basis and the seismic capacity achieved from one industry code to the next. The NRC instead focused on assuring that margins were acceptably large or larger. Also, this in no way implies that the engineering community (including the NRC) was somehow historically derelict in their duty – far from it. Seismic engineering was a rapidly maturing field and practitioners were acutely conscious of the limitations that existed at any given time, including the variability of outcomes that would result in the code development process. When most of the operating US NPPs were designed, the philosophical underpinnings of the more integrated approach discussed above did not exist, nor did the analysis tools exist that could be used routinely to ascertain whether the desired outcome (in terms of seismic capacity, described probabilistically by a fragility curve) would be achieved.

However, today there has been an evolution on both fronts: (i) First, in the philosophy of NRC regulation, which is continually becoming more risk-informed and performance-based, and which has a goal to integrate design, construction, maintenance, and inspection activities more fully; and (ii) second, in the ability of the code committees to embed an approach that could better integrate into an ideal regulatory framework with the confidence that both analysis of the outcome of the design and a comparison against a target are feasible.

3 ATTRIBUTES OF AN IDEAL FRAMEWORK FOR ASSURING SEISMIC SAFETY

The title of this section uses the phrase “ideal Framework” although it is recognized that no Framework can fully reach that goal. However, trying to delineate what such an ideal Framework might look like can help to focus attention on where and why advances are beneficial and feasible.

In an ideal Framework, achieving overall acceptable seismic safety in a large NPP requires successfully implementing three steps:

(1) A Design-Basis Earthquake (DBE) ground motion, characterized according to NRC regulations and guidance and including a frequency spectrum and other attributes, needs to be selected by the regulators and used by the plant’s designers. This DBE ground motion must be site-specific, as required by the US Code of Federal Regulations and described in NRC guidance (NRC 1997; NRC 2007). A method for how the ground motion is used in the design also needs to be specified. This method must account for the fact that incoming seismic energy produces different seismic excitations in different plant locations (both across the plant and within structures). The DBE ground motion needs to be large enough (and hence its annual frequency of exceedance needs to be small enough) that when used in the design process the outcome will meet the regulators’ criteria for *adequate seismic safety*.

(2) Design codes need to exist, each of which uses the DBE ground motion as the basis for the design of individual SSCs, along with procedures for their use. These codes and procedures will necessarily differ from one design problem to another. For example, there need to be different design codes for concrete shear walls than for electrical components. When used in the design process, each code should produce designs for individual SSCs that meet the regulators’ criteria for assuring adequate seismic capacity. These design codes also need to be written in a way so that they can be used successfully by large numbers of designers and equipment manufacturers working in the industry. A design code that can be implemented successfully only by a small cadre of cutting-edge experts is not an adequate code and does not meet the needs of the NRC or industry. Also, any design code for use with NPPs needs to be coordinated with the NRC’s current approach in 10 CFR 50.69 (NRC, 2004) for risk-informed categorization and treatment of SSCs.

(3) An overall NPP systems design needs to be executed in a manner so that the individual SSCs, as they come together to comprise the overall plant design, produce an NPP having a *transparent and demonstrably adequate overall seismic capacity*, or an *acceptable overall seismic risk profile*.

How a nuclear regulatory authority judges the overall adequacy of an NPP’s seismic design can take one of three broad approaches.

- In the first approach, the regulator can determine overall adequacy by assuring that each SSC has adequate seismic capacity, and then the regulator can judge overall adequacy by asserting that a plant with adequate individual SSCs will be adequately safe overall.
- In the second approach, the regulator can determine overall adequacy by assuring that each SSC has adequate seismic capacity, and then the regulator can judge overall adequacy by assessing the adequacy of the overall system using some additional analysis judged against some additional criterion.

- In the third approach, the regulator can determine overall adequacy by assessing the adequacy of the overall system using an overall assessment method (judged against a specified criterion) but without necessarily passing judgment on the adequacy of each individual SSC's seismic capacity.

This Section and the next will explain why the second of these three approaches is to be preferred.

An ideal Framework for designing, analyzing, deploying, and regulating for seismic safety of an NPP would have the following 19 attributes and be supported by the following information and analyses. The 19 attributes fall into 4 broad categories: those involving high-level regulatory philosophy, those involving seismic input, those involving design guidance, and those involving analysis.

3.1 High Level Regulatory Philosophy and Approach for Seismic Safety

(A) Safety objective or performance target

A *safety objective* or *performance target* (or more than one)⁴ would exist against which the plant's overall seismic safety is judged, and it would be used in the regulatory decision process. Decisions by the designer at the design stage and by the NRC at the final regulatory-approval stage would consider these safety objectives or performance targets as a key criterion.

(B) Roles of seismic vs. other initiators

The safety objective or performance target would account for the relative roles of earthquake-initiated accidents and of accidents initiated by other causes in judging the adequacy of the seismic part of the overall design.

(C) Core damage vs. large early release

The safety objective or performance target would account for the relative importance of designing against potential seismic-initiated core-damage accidents and designing against seismic-initiated accidents leading to a large early radioactive release.

⁴ There is a difference between a safety objective and a performance target. As used here, a *safety objective* is some public-policy objective for overall plant safety, such as found in the NRC's "safety goals" (NRC, 1986) or subsidiary objectives, specifically in terms of the likelihood of offsite public-health consequences. A *performance target*, as used here assures that the annual frequency of some undesired consequence is less than a certain target. For example, for the plant as a whole, the seismic performance target might be assuring that the core-damage frequency from earthquakes is less than, say, 1×10^{-5} /year, with the assurance based on an analysis using seismic-PRA methods. For an individual SSC, it might take the form of assuring that its overall annual probability of failing seismically so that it could not perform its safety function is less than 1×10^{-4} , arrived at by integrating the seismic hazard at the input to the SSC with the seismic fragility function for the individual SSC item. Another form that a seismic performance target might take is achieving a specified high level of confidence that the SSC will perform its safety function if a specified large earthquake were to occur. It is likely that, as a matter of policy, the "performance target(s)" would be derived from and would be subsidiary to the "safety objective(s)."

(D) Analysis methods

Demonstrating that the safety objective(s) or performance target(s) are met would need to be feasible using established and readily accessible analysis methods.

(E) Confidence level

The confidence level required for meeting the seismic-safety regulations would be explicitly stated, used as an element of the decision criteria, and accounted for in the compliance analysis. That is, regulatory decisions and the criteria governing them would explicitly require a specified quantitative confidence level as an acceptance criterion. In the analysis of the confidence level, account would be taken of both epistemic uncertainty and aleatory variability in understanding the phenomena and performance being analyzed.

(F) Defense in depth

The defense-in-depth principle would be incorporated by requiring a “balanced” seismic risk profile in which no one aspect or single failure governs the profile⁵. A criterion for deciding whether the seismic design possesses adequate defense in depth would be established and used. That criterion would differentiate between moving toward a more “balanced” risk profile by decreasing the overall absolute risk and moving toward it by changing the relative contributions of different “risk contributors.”

3.2 Seismic Input

(G) Use of the design-basis earthquake concept

The design guidance would use a specified design-basis earthquake (DBE) ground motion (or a set of them) that is site-specific and is formulated in a way that typical designer engineers⁶ can use (see Attributes (H) and (I) below).

(H) Probabilistic seismic hazard analysis

A site-specific probabilistic seismic hazard analysis (PSHA) focused on assessing the best estimate and uncertainty in the ground motion would be required and would be used as the basis for choosing the design-basis earthquake (DBE) ground motion at a given site. The method to be used in selecting the DBE ground motion(s) based on a PSHA would be established by regulation, by regulatory guidance, or by a consensus code/standard.

⁵ See Section 7 for a discussion of the opposite of this, which is illustrated by some actual case-study examples.

⁶ The term “typical design engineer” as used herein is intended to describe a skilled design engineer typically working in a large design organization. The engineering community demands that such design engineers have certain credentials and that their skills include the ability to understand and use the consensus design codes and standards applicable to the task at hand. However, the typical design engineer is not expected to go beyond the design requirements (or at least not much beyond) in the day-to-day work contemplated herein.

(I) *Design-basis earthquake ground motion selection*

The DBE ground motion would be selected based on the PSHA in (H) above (using the method in (H) and with an appropriately chosen target probability of exceedance) so that the overall safety objective(s) or performance target(s) (see Attributes (A), (B), and (C) above) are achieved when the DBE ground motion is used in conjunction with the various coordinated design codes and standards, and accounting for other elements of the overall system (the idea of consistency between the design ground motion and various codes and guidance is explained in Section 2.5). This would require that the method for DBE ground motion selection (see Attribute (H) above) and the design codes and standards (see Attribute (K) below) be developed so that they are consistent. Again, achieving this coordination is not straightforward in either a technical or a policy sense.

(J) *Realistic analysis of SSI and ISRS*

The analysis of the seismic motion (excitation) at various locations at a given NPP facility would be probabilistic, and based on a realistic analysis of how the earthquake energy enters the site and then the structure. This would include realistic soil-structure-interaction (SSI) analysis and a realistic structural analysis, leading to realistic probabilistic in-structure response spectra (ISRS) at all relevant locations in the facility.

3.3 Design Guidance

(K) *Design guidance must be specific*

Design is performed for one SSC at a time, and is usually done by a working design engineer. Therefore, the ideal Framework would need to include design guidance (based on regulations, codes, and design standards) that is specific enough that it can be used by the working designer to design a specific safety-related SSC item without the need for constant and routine reference to the high-level safety philosophy and the higher-level regulations. This means that there would need to be a “translation” from the high-level (plant-as-a-whole) safety objective(s) or performance target(s) (in (A), (B), and (C) above) to specific design guidance at the SSC level. This translation is not straightforward in either a technical or a policy sense. Other attributes of the Framework touch on this issue (see Attribute (I) above).

In an ideal Framework, the design guidance would be formalized in a set of consensus design codes and standards. This set would need to provide the typical design engineer with specific guidance at the SSC level that has been developed and demonstrated to meet the high-level (plant-as-a-whole) safety objective(s) or performance target(s) (in (A), (B), and (C) above).

As noted above at the beginning of Section 3, any design code for use with NPPs needs to be coordinated with the NRC’s current approach in 10 CFR 50.69 (NRC, 2004) for risk-informed categorization and treatment of SSCs.

(L) *Design guidance and performance targets*

The design guidance that is intended for routine use in designing a specific safety-related SSC item would not require the designer to demonstrate by analysis that the SSC has met a specific “performance target,” but a specific performance target would underlie the guidance, such that if a designer follows the design guidance there would be high confidence that the performance target has been met for that safety-related SSC.

(M) Design of an SSC against a performance target (optional)

As an alternative to the design approach in (L), the designer would be offered the option of demonstrating directly that the underlying performance target has been met for a given SSC, based on specified acceptance criteria tied to specified methods of demonstration (e.g., by analysis, test, field experience, etc.).

(N) Design codes based on realistic analyses and data

The consensus design codes for all of the various categories of safety-related SSCs would be based on realistic analyses and/or adequate tests or experience data. They would insert conservatisms only with an explicitly stated rationale for each, such as accounting for material variability or for variations in outcomes when different engineers use ostensibly similar design practices. (Of course, in practice the extent to which a “realistic” analysis can be performed and relied on is often limited by available knowledge, including not only phenomenological uncertainties but limitations in data and modeling. Here the attributes of an “ideal” system are being described.)

3.4 Analysis Guidance

(O) Realistic response analysis and fragility analysis

The regulations would require the applicant/licensee to perform a realistic analysis of the response behavior of each individual safety-related SSC in postulated earthquake scenarios, coupled with a realistic probabilistic capacity/fragility analysis. Conservative or bounding analysis, rather than realistic analysis, would be permitted if the outcome of such analysis is demonstrably highly conservative, and if a realistic analysis is not needed for other purposes.

(P) Seismic probabilistic risk assessment (SPRA)

The regulations would require the applicant/licensee to perform a realistic analysis of the role that each safety-related SSC plays as part of the overall system that achieves adequate seismic safety for the facility. This means that performing a SPRA is a necessity. The SPRA would be used to guide technical decisions on the topics covered in (Q), (R), and (S) below. (The same comment applies here that was made above under (N), namely that in practice the extent to which a “realistic” analysis can be performed and relied on is often limited by available knowledge, including not only phenomenological uncertainties but limitations in data and modeling. Here the attributes of an “ideal” system are being described.)

Using an iterative approach to design: At this point it should be noted that the Framework, even the ideal Framework, is intended to function in practice through an iterative approach: After the design of individual SSCs is completed (see Attributes (K), (L), and (M) above) based on the inputs (see Attributes (G), (H), (I), and (J) above), a seismic PRA (see (P) above) would reveal whether any of the higher-level criteria (Attributes (A) through (F) above) is not met. If so, iteration of the design is required. The defense-in-depth criterion, Attribute (F), is particularly susceptible to the need for such iteration at a relatively late stage. Also, the graded approach (see Attribute (Q) below) may lead to iteration of the design.

(Q) *Design tailored to overall safety role of each SSC*

The regulations would require that the seismic design of each safety-related SSC would account for (that is, would be tailored to) that SSC's role in overall plant safety, as determined by the seismic PRA. This implies a graded approach to seismic design for those SSCs, because some have a greater impact than others on the risk profile of the plant. This also implies that an iteration of the original design (or more than one) may be needed, based on the results of the PRA analysis of the original design.

(R) *Accounting for non-seismic failures and human errors*

The regulations would require that account be taken of the role of non-seismic failures and human errors in the accident sequences considered in the seismic safety analysis, and would require the design to account for this role.

(S) *Accounting for correlations*

The regulations would require that account be taken of the correlations among seismic failures and the interactions among SSCs, as well as of the role of potential cascading and consequential failures of multiple SSCs, and would require the design to account for these factors. This is discussed in more detail in Section 4.2.4 (S).

4 COMPARISONS OF VARIOUS ATTRIBUTES TO THE IDEAL

Historically, the approach to the challenge of achieving an adequate overall NPP seismic design has not possessed every attribute of the ideal Framework presented above. Some of the attributes of the current Framework fall short compared to the ideal approach. This section discusses and evaluates the current Framework against the 19 attributes of the ideal Framework previously discussed (see Attributes (A) through (S) in Section 3). The section begins with a summarized evaluation in tabular form (Table 4-1), followed by a more in-depth evaluation of the several Framework Attributes (Section 4.2).

Before presenting this evaluation, it should be noted that there is a danger that the message of the following discussion might be misunderstood. Specifically, the reader may question the basis for the broad confidence today that US operating nuclear power plants have adequate safety against large earthquakes, if the shortfalls from an ideal Framework are as described below.

The basis has two pillars. First, the current regulatory Framework for seismic safety has large conservatisms in many areas. Every analysis performed to support the regulatory decisions on plant safety (in all areas, which include seismic hazard; analysis of seismic motion making its way from subsurface into the buildings and then up into the structures; analysis of the seismic response and capacity of structures and of equipment; systems analysis) has demonstrably large conservatisms.

Second, compared to when the current operating reactors were designed, it is feasible now to perform much more realistic analyses supported by data and experience, both engineering analyses in the several areas noted just above and probabilistic risk assessments (PRAs) of overall plant seismic safety and of its major constituent “parts.” These SPRAs, for which there now exists an industry consensus methodology standard (ASME-ANS 2013), indicate to the community of experts in this field that the fleet of operating plants is generally adequately safe, both when judged against the specific details in NRC’s seismic regulations and when compared to the NRC’s broad safety goals.

This general statement does not necessarily apply to any individual plant. Each NPP is different and must be judged separately because an individual plant may “fall short” in one or more of these areas. However, taken as a whole, and with some caution against overconfidence, there is a strong basis for concluding that the fleet of plants generally exhibits broadly acceptable seismic safety, and in any event “falling short” in one area does not necessarily mean that an individual plant’s seismic safety is inadequate overall. Each case is different.

4.1 Summary of Evaluation of Framework Attributes

Table 4-1 summarizes the evaluation of the current regulatory Framework with respect to the 19 attributes of an ideal Framework defined in Section 3.

Table 4-1 Status and evaluation of the various attributes of today’s Framework for seismic design, analysis, and regulation of SSCs in NPPs

	ATTRIBUTE	CURRENT STATUS	EVALUATION: comparison against the ideal	SUGGESTIONS for further development
HIGH-LEVEL REGULATORY PHILOSOPHY AND APPROACH FOR SEISMIC SAFETY				
A.	Safety objective or performance target	Such a target exists for the seismic design of new plants: to meet NRC criteria -- the plant-level HCLPF capacity must exceed 1.67 times the SSE.	Acceptable but lacking enough scope, in that the NRC criteria do not differentiate between CDF vs. LERF. Consideration of the relative roles of seismic vs. other initiators is also not part of current NRC policy.	New regulatory policy development would be needed to address LERF and to account for seismic vs. other initiators.
B.	Roles of seismic vs. other initiators	Analysis of this aspect is feasible.	See note in A above.	New regulatory policy development would be needed.
C.	Core damage vs. large early release	Analysis of this aspect is feasible.	See note in A above.	New regulatory policy development would be needed.
D.	Analysis methods	Analysis methods exist and are readily accessible	Adequately close to the ideal	N/A
E.	Confidence level	Analysis of confidence level is feasible.	The NRC plant-level HCLPF criterion ($HCLPF \geq 1.67$ times the SSE ground motion) implicitly embeds a confidence level in the HCLPF concept, but the scope does not cover B and C above.	New regulatory policy development would be needed.

Table 4-1 Status and evaluation of the various attributes of today’s Framework for seismic design, analysis, and regulation of SSCs in NPPs (Cont’d)

	ATTRIBUTE	CURRENT STATUS	EVALUATION: comparison against the ideal	SUGGESTIONS for further development
F.	Defense in depth	Analysis methods exist that can identify possible defense-in-depth vulnerabilities.	No defense-in-depth criterion exists.	New regulatory policy development would be needed.
SEISMIC INPUT				
G.	Use of the design basis earthquake ground motion concept	Methodology for this exists.	Adequately close to the ideal.	N/A
H.	Probabilistic seismic hazard analysis as the basis for selecting the DBE ground motion	The methodology for realistic site-specific PSHA exists, including extensive guidance, and is in common use. Using PSHA to select the DBE ground motion is now part of NRC guidance.	Adequately close to the ideal.	N/A
I.	Design basis earthquake ground motion selection: working together with design codes.	The methodology exists.	Far from the ideal: Some codes work together with the DBE ground motion well (for example, ASCE 43-05) but others do not.	Code committee work and new regulatory policy development would be needed.

Table 4-1 Status and evaluation of the various attributes of today’s Framework for seismic design, analysis, and regulation of SSCs in NPPs (Cont’d)

	ATTRIBUTE	CURRENT STATUS	EVALUATION: comparison against the ideal	SUGGESTIONS for further development
J.	Realistic analysis of SSI (soil structure interaction) and ISRS (in-structure-response spectrum)	<p>Long-standing SSI methods exist that are functional and useful but have shortcomings. R&D is under way to develop advanced SSI methods.</p> <p>Probabilistic ISRS methods exist but are not in widespread use.</p>	<p>Adequate for use today.</p> <p>More advanced SSI methods would make feasible more realistic probabilistic SSI analysis.</p> <p>More advanced probabilistic ISRS methods should soon come into more common use.</p>	N/A
DESIGN GUIDANCE				
K.	Design guidance must be specific	<p>Design codes and standards exist, but do not generally “work together” with other guidance and regulations.</p> <p>ASCE 43-05 has a model for how to develop specific performance-based design guidance for the practicing design engineer that meets higher-level regulations and requirements.</p>	<p>Far from the ideal: Some codes work together well with the DBE ground motion (for example, ASCE 43-05) but others do not.</p> <p>Adequate for interim use in an iterative design process, but significant progress would be needed in policy development.</p> <p>A design target has not been adopted as an NRC numerical regulatory criterion.</p>	<p>To bring about the development of specific performance-based design guidance for the practicing design engineer that meets higher-level regulations and requirements, code-committee work and new regulatory policy development would both be needed.</p>
L.	Design guidance and performance targets	This situation is current practice.	Adequate	N/A

Table 4-1 Status and evaluation of the various attributes of today’s Framework for seismic design, analysis, and regulation of SSCs in NPPs (Cont’d)

	ATTRIBUTE	CURRENT STATUS	EVALUATION: comparison against the ideal	SUGGESTIONS for further development
M.	Design of an SSC against a performance target (optional)	Design methods and analysis to demonstrate meeting a performance target exist but are rarely exercised, and would be difficult to use by the typical practicing design engineer.	Case-by-case application of this option means that no technical development would be required.	New regulatory policy development would be needed.
N.	Design codes based on realistic analyses and data	Current practice is largely acceptable.	Advances are needed to assure that all codes and standards are based on realistic input and realistic analysis, and that conservatism are only inserted with an explicitly stated rationale.	Code committee work and new regulatory policy development would both be needed.
ANALYSIS GUIDANCE				
O.	Realistic response analysis and fragility analysis	<p>The methodology for probabilistic response analysis exists, including extensive guidance, and is in common use.</p> <p>The methodology for probabilistic seismic capacity (fragility) analysis (the “fragilities” aspect of seismic PRA) exists, including extensive guidance, and is in common use.</p>	Adequately close to the ideal	N/A

Table 4-1 Status and evaluation of the various attributes of today’s Framework for seismic design, analysis, and regulation of SSCs in NPPs (Cont’d)

	ATTRIBUTE	CURRENT STATUS	EVALUATION: comparison against the ideal	SUGGESTIONS for further development
P.	Seismic probabilistic risk assessment	This is the systems aspect of SPRA. This methodology exists, including extensive guidance, and is in common use.	Adequately close to the ideal.	N/A
Q.	Design tailored to overall safety role of each safety-related SSC	No such requirement exists: each safety-related SSC is currently designed individually.	Far from ideal: insufficient account is taken of the role of an individual safety-related SSC.	New regulatory policy development would be needed.
R.	Accounting for non-seismic failures and human errors	The methodology exists for accounting for these issues.	Far from ideal	New regulatory policy development would be needed.
S.	Accounting for correlations	The methodology exists for accounting for correlation. However, major uncertainties exist in the analysis because not enough is known about what governs the extent of correlation.	Far from ideal	New technical research would be needed, followed by new regulatory policy development.

4.2 In-depth Evaluation of Framework Attributes

This Section provides an evaluation in greater depth of the various Framework attributes described in Section 3 and then discussed in brief in Table 4-1.

4.2.1 High Level Regulatory Philosophy and Approach for Seismic Safety

- (A) Safety objective or performance target**
- (B) Roles of seismic vs. other initiators**
- (C) Core damage vs. large early release**

Attributes (A), (B), and (C) will be discussed together here because they are closely related. As noted in Table 4-1, for a new NPP design the NRC has promulgated an acceptance criterion for the overall adequacy of the seismic part of the design (NRC 1993). Specifically, as discussed starting three paragraphs below, the plant-level HCLPF capacity must exceed 1.67 times the safe-shutdown earthquake (SSE) ground motion. This criterion satisfies the policy imperative in Attribute (A).

However, as useful as it is, this criterion does not go far enough. Specifically, it does not include criteria that would judge the adequacy of a design against both a core-damage endpoint and a large-early-release endpoint (Attribute (C)). Also, the NRC has no policy that would indicate whether the relative roles of earthquake-initiated accidents and of accidents initiated by other causes should be considered in a judgment of overall adequacy (Attribute (B)). For Attributes (B) and (C), addressing these issues would require new policy development at the NRC. See Section 6.1.

Concerning the safety objective or performance target (Attribute (A)): For new NPP designs, the NRC has developed and currently applies a target that must be met to support an NRC finding that the plant has acceptable seismic safety. The target is expressed in terms of a HCLPF seismic capacity for the plant-as-a-whole (NRC 1993). This is a major philosophical advance.

The analysis needed to produce risk values that can be compared with the NRC target value proceeds as follows. When the seismic design of the NPP is complete, the analyst performs an SPRA or a Seismic Margin Assessment (SMA) using PRA systems-analysis methods (ASME-ANS 2013 and NRC 2010). The outcome of this analysis is a seismic “fragility curve” for the plant-as-a-whole. This plant-level fragility curve (Kennedy and Ravindra 1984; Reed and Kennedy 1994; EPRI 1991; EPRI 2002) is developed as follows. First, the individual seismic failures come together into seismic-initiated accident sequences, each of which has an accident-sequence-level fragility curve. Based on this, an associated HCLPF seismic capacity is developed using the usual computational rules for HCLPF capacity analysis (EPRI 1991). If there is one dominant accident sequence, the HCLPF capacity of the plant-as-a-whole is numerically equal to the HCLPF capacity of that sequence. If several sequences are “close together,” then the usual computational rules are used to determine the fragility curve (and hence the HCLPF capacity) of the plant-as-a-whole. The entire approach can account rigorously for any non-seismic failures or human errors in any of the accident sequences.

The HCLPF capacity of the plant-as-a-whole is then compared numerically to the NRC target, which is equal to 1.67 times the SSE ground motion (NRC 1993). The plant is judged adequately designed if the plant-level HCLPF capacity exceeds the target. It is important to

notice that because this approach relies on a SPRA or an SMA using PRA systems-analysis methods, it intrinsically embeds probabilistic logic and relies on probabilistic analysis, including accounting explicitly for both epistemic and aleatory uncertainty.

Although the NRC has a risk target for the plant-as-a-whole for new plant designs, no comparable risk target for defining acceptable seismic safety has been established by the NRC for use in the design of individual SSCs, nor is there a target for any individual accident sequence(s). Such an SSC (or accident-sequence) target might be framed either in probabilistic terms (as a target for an acceptable frequency per year of overall failure) or in capacity margin terms (in terms of how much seismic margin above the design-basis ground motion an individual SSC must have to be acceptably strong, and with what confidence level). A decision criterion framed in probabilistic terms could be phrased something like, “An SSC is acceptably designed if its annual mean frequency of seismic-induced failure, determined by a mathematical convolution of its seismic capacity with the seismic hazard it experiences at its location in the plant, is less than $A \times 10^{-N}$ per year.”⁷ A decision criterion framed in capacity-margin terms could be established at the HCLPF capacity, for example something like, “An SSC is acceptably designed when its HCLPF capacity (or its median capacity) exceeds the DBE ground motion by a factor of F.”⁸ Alternatively, similar decision criteria could be framed at the accident-sequence level.

(D) Analysis methods

As noted in Table 4-1, analysis methods now exist for demonstrating that a safety objective or performance target is met. These methods are seismic PRA or the closely related seismic margin assessment (SMA) methodology using PRA-based systems-analysis methods. These methodologies are both accessible and widely used.

Unfortunately, a shortcoming today is that SPRA and SMA methods are not yet in widespread use, so the insights from them are not widely considered in decision-making.

As a general matter, both plant managers and regulators often make safety decisions that rely on a judgment as to the relative importance to overall plant risk of the potential failures of the various SSCs and systems within their purview. All US nuclear plants currently have a modern internal-events PRA, but many plants do not have a modern SPRA or an SMA using PRA-based systems-analysis methods. Many plants don’t use the insights from an earlier one if they have one, because the earlier one has typically not been kept up-to-date. Thus in the seismic arena the NPP’s safety decision-makers (or its regulators) are often forced to work with incomplete or missing information on the subject of relative importance when it comes to seismic issues.⁹

This means that on seismic-safety issues their thinking is often informed mainly by the logic of the era in which the plant was designed and its operating procedures established – an era of compartmentalization rather than systems thinking, and an era in which adequacy was defined in terms of “adequate conservatism” (however judged), rather than in terms of understanding

⁷ Here A and N are numerical values.

⁸ Here F is a numerical factor greater than one, such as 1.3 or 1.8.

⁹ In the U.S., this situation will change for the better in the next few years. As a follow-on to the lessons-learned after the 2011 Fukushima accident in Japan, perhaps half of the U.S. NPPs plan to perform SPRAs soon. Most of these will be performed in response to the NRC post-Fukushima generic letter (NRC, 2012b) while others are being undertaken on an individual plant’s own initiative.

realistic performance and accounting realistically for the uncertainties in knowledge, and then judging the SSC's adequacy (or that of the plant as a whole) against performance targets.

(E) Confidence level

As noted in Table 4-1, analysis methods exist now for understanding the confidence level of an analysis. This is achieved through the methodology of SPRA with full uncertainty analysis. The methodology is both accessible and widely used. Nevertheless, there is no explicit statement in regulation or regulatory policy as to what confidence level is sought or required in order for a proposed design to meet the seismic regulations.

However, the NRC has made important progress along the lines of this idea. Specifically, the criterion discussed above under Attribute (A), that the plant-level HCLPF seismic capacity of a new NPP must exceed 1.67 times the SSE ground motion (NRC 1993), contains an implicit confidence level. This is because the HCLPF capacity is the ground motion at which the analyst has about 99 percent confidence that the item (in this case, the plant-as-a-whole) will have a seismic capacity above the HCLPF level. This translates roughly into 99 percent confidence that failure will not occur for a ground motion level equal to 1.67 times the SSE.

Nevertheless, this confidence-level argument has not been translated directly into an NRC statement of confidence. Rather, the confidence level is implied or implicit. To remedy this and bring Attribute (E) closer to the ideal Framework discussed herein would require new policy development at the NRC. See Section 6.2 for further discussion.

(F) Defense in Depth

As noted in Table 4-1, analysis methods exist that can identify possible defense-in-depth vulnerabilities. These methods are SPRA methods.

The issue or problem can be described as follows: based on a large number of SPRAs of existing plants (NRC 2001), the current design approach has produced operating NPPs whose seismic *risk profile* is at many plants characterized by a single *leading* or *dominant* SSC whose seismic failure is the most important in contributing to overall seismic risk. This is in addition to the seismic vulnerability of offsite power, which is assumed in the analysis to fail in almost any ground motion large enough to cause other failures. This dominance of a single SSC, or sometimes a very small number of them, is inconsistent with the overall NRC philosophy of defense in depth, which seeks a more balanced risk profile. See Section 5 for further discussion of this, and Section 7 for some actual examples.

It is useful to note here the NRC's formal definition of defense in depth in the *Glossary* of the agency's Strategic Plan (NRC 2012):

Defense in depth: an element of the NRC's Safety Philosophy that employs successive compensatory measures to prevent accidents or lessen the effects of damage if a malfunction or accident occurs at a nuclear facility. The NRC's safety philosophy ensures that the public is adequately protected and that emergency plans surrounding a nuclear facility are well conceived and will work. Moreover, the philosophy ensures that safety will not be wholly dependent on any single element of the design, construction, maintenance, or operation of a nuclear facility.

The objective of this philosophy has not been met for many of the operating NPPs to the extent that the SPRAs reveal that the seismic part of the risk profile is dominated by a single SSC, or

perhaps a very small number of them. This situation will be called an “unbalanced seismic risk profile,” and it is present at many of the operating NPPs.

The reason why an unbalanced risk profile might be a problem has to do with confidence (or the related ideas of overconfidence or hubris), and/or with the fact that uncertainties not only exist but may not be well characterized. If one admits to the possibility that an error in either design or analysis has occurred, or that even if the design or analysis is of high quality that the inherent uncertainties may not be well characterized, then defense in depth is a way to provide an additional layer of assurance that there is not a “cliff” where serious safety problems may occur with a small additional load (e.g., a slightly higher ground motion). However, if a single SSC dominates seismic risk, (even if the SSC fully meets all regulations), then if one supposes that the SSC is actually even weaker because of a design, analysis, or maintenance error, then the plant may be far more vulnerable than if the design were more balanced in terms of contributions to the risk profile from various SSCs.

This is perhaps where one of the greatest opportunities exists for a near-term advance, as discussed in Section 5.

However, evolutions of policy on this issue must account for, or differentiate between, moving toward a more “balanced” risk profile by decreasing the overall absolute risk and moving toward it by changing the relative contributions of different “risk contributors.”

It is recognized that the “defense-in-depth” philosophy encompasses considerably more than the single attribute of a “balanced risk profile.” However, the several other complex issues involved with meeting the broad goals of “defense-in-depth” are not addressed herein, involving as they do issues well beyond the scope of this project.

4.2.2 Seismic Input

(G) Use of the design-basis earthquake ground motion concept

This is current practice, and the methodology exists.

(H) Probabilistic seismic hazard analysis (PSHA) as the basis for selecting the DBE ground motion

PSHA as a technical discipline is well established, and using PSHA to select the DBE ground motion is now part of NRC guidance in Regulatory Guide 1.208 (NRC 2007). This is therefore current practice.

(I) Design-basis earthquake ground motion selection

As noted earlier, satisfying this Attribute would assure that the selection of the DBE ground motion is done in a way that is consistent with the rest of the Framework, and crucially with the design codes, so that the safety objective(s) or performance target(s) are met.

Historically, the DBE ground motions chosen by the NRC for use at the existing individual NPP sites were not chosen with a consistent numerical risk target in mind. When these DBE ground motions were selected decades ago, this way of thinking had not yet even been described, nor, of course, accepted. The basis for their selection, as embedded in 10 CFR Part 100 Appendix A (NRC 1977) and other regulatory positions, did not incorporate a risk criterion. In fact, the

DBE ground motions for the various US operating plants have very different mean annual frequencies of occurrence, ranging from about 2×10^{-6} per year to about 1×10^{-4} per year (NRC 1997), although this was not understood at the time of their selection.

(J) *Realistic analysis of SSI and of ISRS*

As noted earlier, long-standing soil structure interaction (SSI) methods and in-structure response spectra (ISRS) methods exist and have been extensively used and validated. They are adequate today. More advanced probabilistic SSI methods are under development now, and more advanced ISRS methods should soon come into more common use.

4.2.3 Design Guidance

(K) *Design guidance must be specific*

As noted earlier, design is performed for one SSC at a time, and is typically done by typical practicing design engineers. Therefore, the ideal Framework would need to include design guidance (based on regulations, codes, and design standards) that is specific enough that it can be used by the typical practicing design engineer to design a specific safety-related SSC item without the need for constant and routine reference to the high-level safety philosophy and the even higher-level regulations.

Historically, the design codes for various categories of SSCs, developed by different consensus code committees (under the American Society of Mechanical Engineers, American Nuclear Society, American Concrete Institute, American Society of Civil Engineers, Institute of Electrical and Electronic Engineers, etc.), have produced different outcomes in terms of the seismic margin of the SSCs from category to category. There are wide differences in some cases. This arises mostly because the different code committees historically made different judgments about the design “targets” in terms of their likelihood of failure above the design basis, and these targets were sometimes not explicitly expressed, at least not in a numerical sense. Also, the code committees chose to incorporate different amounts of conservatism, to account for different types of issues (some of which were known issues and some of which were uncertainties) and/or different treatments of them.

A hypothetical but realistic example may be helpful. A code committee typically instructs the designer to use certain specified design rules along with a specified site-specific DBE ground motion. These design rules are selected to produce an acceptable outcome every time while accounting, for example, for uncertainty and variability in material properties, variability from designer to designer in applying the methods, and uncertainties in the design outcome for otherwise similar items depending on their size, within a specified size range. The designer need not concern himself or herself with these issues; he or she need only follow the code’s prescriptions. Assuring an acceptable outcome for every item designed according to the code, and accounting for the several issues mentioned, is the committee’s responsibility. Crucially, how the various code committees have gone about judging the adequacy of the various margins used in the analyses supporting the code has varied from one design code to the next.

Fortunately, one major advance has occurred in recent years in the important industry consensus code ASCE 43-05 (ASCE 2005). ASCE 43-05 provides specific design guidance for the practicing engineer at the SSC level that can be demonstrated to meet the high-level (plant-as-a-whole) safety objective(s) and performance target(s) (in Attributes (A), (B), and (C) above).

Specifically, ASCE 43-05 identifies and uses performance targets for various categories of SSCs so that, if the code is applied properly, the SSC designed according to the code achieves one or another of these performance targets. There are several performance targets which are “graded,” so that a safety-related SSC to be used in an NPP requires more seismic capacity than a similar SSC intended for, say, service in a facility with less safety importance or less potential for large seismic-induced consequences. For NPPs, the performance target in ASCE 43-05 is an annual frequency of earthquake-induced failure of about 1×10^{-5} per year. This is achieved by selecting a DBE ground motion at a higher annual frequency (for example, a site-specific DBE ground motion frequency of 1×10^{-4} per year) and then following the code’s design rules, which embed certain conservatisms to produce an SSC that meets or exceeds the performance target.

The ASCE 43-05 approach uses probabilistic thinking in several of its steps, starting with a probabilistically chosen DBE ground motion (from a PSHA) and ending with a design that is compared to a probabilistic performance target. The design rules themselves, however, employ a set of engineering approaches that follow the philosophy of the more traditional design codes. For example, ASCE 43-05 has rules for margins, rules for analysis, rules for selecting the cases to be analyzed, and specified acceptance criteria. The rationale for this (which makes good sense in the world of real engineers designing real SSCs) is that the typical practicing design engineer needs specific deterministic design rules and specific engineering acceptance criteria.

It is important to note that no other nuclear design codes approach even this part of the problem in this way. Crucially there is no overall coordination across the many different design codes. To accomplish this broad coordination would almost certainly require a high-level coordination function to be undertaken by an industry-wide group or by an agency like the NRC. The NRC would need to adopt as a policy objective some sort of performance target that each of the several code committees could then use as its own benchmark target(s) in modifying that code to work together with the design rules to achieve the desired objective. This implies not only that the NRC would adopt a risk target but also that it would adopt a confidence level that it would ask each design code committee to assure would be met by a practicing designer who uses that code.

As noted earlier, this translation is not straightforward in a technical sense, and the coordination is also not straightforward in a policy sense. Also, despite the obvious philosophical advantages of the approach in ASCE 43-05, the NRC has not adopted this ASCE 43-05 approach as an element of its own regulatory guidance.

Also, the ASCE 43-05 approach is for the design of an individual SSC, to be done as a separate design problem without reference to the role of that SSC in the overall plant. While that in and of itself is a shortcoming (see Attribute (Q)), ASCE 43-05 goes a long way toward the broad objective described here.

As noted above at the beginning of Section 3 and in Section 3.3(K), any design code for use with NPPs needs to be coordinated with the NRC’s current approach in 10 CFR 50.69 (NRC, 2004) for risk-informed categorization and treatment of SSCs.

(L) *Design guidance and performance targets*

This is current practice and is adequate.

(M) Design of an SSC against a performance target (optional)

This Attribute is moot at the current time because in any event there is currently no performance target. In an ideal Framework, this Attribute would allow a designer (as an option) to demonstrate by analysis and/or test that the performance target has been met. This would be on a case-by-case basis. Because this would require major new policy development, and in any event the analysis would be tailored to the specific case, the technical feasibility of this Attribute is not an issue *per se*. See Section 6.3 for further discussion.

(N) Design codes based on realistic analyses and data

Current practice is largely acceptable. Further work is needed to assure that each consensus code or standard is based on realistic input and realistic analysis, and that conservatisms are only inserted with an explicitly stated rationale. (The same comment on realism made above applies here: That is, in practice the extent to which a “realistic” analysis can be performed and relied on is often limited by available knowledge, including not only phenomenological uncertainties but limitations in data and modeling.)

4.2.4 Analysis Guidance

(O) Realistic response analysis and fragility analysis

The methodology for probabilistic response analysis exists, including extensive guidance, and is in common use. Also, the methodology for probabilistic seismic capacity analysis (the fragilities aspect of SPRA) exists, including extensive guidance, and is in common use. This area is therefore adequately close to what an ideal Framework would require.

(P) Seismic probabilistic risk assessment (SPRA)

This topic refers to the systems aspect of SPRA. This methodology exists, including extensive guidance, and is in common use. This area is therefore adequately close to what an ideal Framework would require.

(Q) Design tailored to overall safety role of each SSC

Today, although the long list of SSCs in an NPP is divided into those that are safety-related and those that are not, there is no requirement or guidance that would require that the seismic design of each safety-related SSC accounts for (that is, is tailored to) that SSC’s role in overall plant safety, as determined by the SPRA. Specifically, although the role of individual safety-related SSCs in overall plant seismic safety differs considerably even for identical components situated differently in the plant design, these differences are seldom accounted for in either design or regulation.

This represents a major opportunity for an advance. Clearly, some SSCs have more impact than others on the risk profile of the plant. Those that stand out as more important should be designed and regulated so that more safety margin is embedded in their design than would be the requirement for those of lesser importance. Notice that this is not only an opportunity for a relaxation for those SSCs of lesser importance – it is also a vehicle for tightening for those of greater importance. Cutting both ways, the graded approach implied by this Attribute of an ideal

Framework would, if implemented, go a long way to bringing about a much more nearly performance-based and risk-informed plant safety design. See Section 6.5 for further discussion.

This also implies that an iteration of the original design may be needed, based on the results of the PRA analysis of the original design.

(R) *Accounting for non-seismic failures and human errors*

The methodology for accounting for these types of failures and errors exists as a part of SPRA, and has been implemented widely. However, the role that these failures and errors play in individual seismic-initiated accident sequences and hence in overall seismic safety is not accounted for in the design of individual SSCs. This is far from ideal, because insights from this type of analysis can help understand which SSCs require more stringent seismic design, and which may need less.

New regulatory policy development, supported by technical analysis, would be needed to bring about a more risk-informed approach in this area. See Section 6.4 for further discussion.

(S) *Accounting for correlations*

The methodology for accounting for correlations among seismic-induced failures, especially but not limited to identical co-located SSCs, exists as a part of SPRA, and has been implemented widely. However, there are major uncertainties in the analyses, because not enough is known about how correlated failures occur (or not), even for seemingly identical co-located SSCs, and less is known for other types of correlations among failures. New technical research is needed to take advantage of safety insights that might be available on this topic. Later, new regulatory policy development, supported by technical analysis, would be needed to bring about a more risk-informed approach in this area. See Section 6.4 for further discussion.

5 A PROPOSAL FOR INCREMENTAL PROGRESS: DEFENSE IN DEPTH

This section will lay out a *proposal for incremental progress*, based on the evaluation in Section 4, which has examined the current status of the various attributes of the current Framework for *designing, analyzing, deploying, and regulating for seismic safety* of an NPP, and has compared them with an “ideal Framework” as described in Section 3. In particular, a major target for interim progress is Attribute (F), concerning an initiative to improve defense in depth.

While recognizing that significant policy development and some technical work will be needed to make important strides toward the “ideal Framework” (see Section 4), all elements of the proposed *interim scheme* are believed to be fully feasible technically, and should also not arouse intense regulatory controversy. Crucially, the scheme described here is fully in line with the regulatory philosophy that is proposed in the NRC’s recent 2012 publication NUREG-2150, “*A Proposed Risk Management Regulatory Framework*” (NRC 2012a). This philosophy proposes, among other things, to use PRA-type and other risk information much more extensively to support NRC regulatory positions and regulatory decisions.

The discussion below will assume that the problem at hand is *to carry out the seismic aspects of the design of a new NPP, optimizing the plant’s safety and cost while meeting all regulatory requirements with margin*. The several Steps that need to be completed will be briefly outlined, in sequence. Where the proposed scheme differs from the current approach and why it is an improvement will be pointed out.

However, to dispel any suspense, the essence of the interim proposal will be described briefly here. Specifically, the proposal is that, once a tentative design has been accomplished, an SPRA would be performed, and the risk insights from the PRA would be used to support certain iterations to improve the design. The improvements would have as their aim three separate objectives: (i) to enhance the design’s “balance” vis-à-vis defense in depth; (ii) to enhance its safety without adding important costs (and in fact perhaps with less cost); and (iii) to provide a more transparent safety case; or some combination of these.

A cautionary note is in order here. Specifically, any analysis of a plant in the design phase has certain intrinsic limitations that can reduce the value of the insights gained. This is true not only because the plant has not been built and operated yet so plant-specific data do not exist, but also because certain decisions on layouts and equipment have not yet been made, nor can such a plant be studied using walkdown methods to understand those issues that only arise after construction and operation.

Also, as noted earlier, caution is called for: Any evolution in policy must differentiate between moving toward a more “balanced” risk profile by decreasing the overall absolute risk and moving toward it by changing the relative contributions of different “risk contributors.”

It is recognized that the “defense-in-depth” philosophy encompasses considerably more than the single attribute of a “balanced risk profile.” However, the several other complex issues involved with meeting the broad goals of “defense-in-depth” are not addressed herein, involving as they do issues well beyond the scope of this project.

The “interim proposal” encompasses 5 “Steps,” as described in Sections 5.1 through 5.5.

5.1 Step A: Complete a New NPP Design

Step A involves completing a new NPP design using existing codes and standards, existing engineering practice, and existing regulations. This Step follows existing practice. It could be described in some detail, but that is not necessary here. Several complete new NPPs have been designed both in the US and abroad, have passed the regulatory review processes in their country, and are under construction now. Therefore, there is a basis for confidence that the methods to implement this Step exist and are widely used.

5.2 Step B: Perform a PRA that Includes a Full SPRA

Step B involves performing a probabilistic risk assessment (PRA) that includes a full SPRA, or that includes a Seismic Margin Assessment (SMA) using PRA-type systems-analysis methods. This Step also follows existing practice. There are a few SPRAs and SMAs for newly designed plants, and more are being done each year. There is, of course, a PRA Methodology Standard developed under ANS and ASME (ASME-ANS 2013) that includes SPRA and SMA, although that Standard is for the analysis of operating LWRs. Furthermore, for design-certification applications to the NRC for new plant designs, the NRC's Interim Staff Guidance ISG-20 requires a PRA-based SMA analysis (NRC 2010).

There is also a new PRA methodology standard under development for analyzing advanced LWR plants in the design stage (ASME-ANS 2015). The same consensus committee that maintains the existing PRA standard for operating plants is writing this new standard. It is expected to be adopted and available in final form by late 2015, and to be published initially for "trial use."

When the term "PRA" has been used here, the intent is typically to describe a Level 1 PRA (whose end point is CDF) plus enough of a Level 2 PRA to identify the important large-early-release accident sequences. A full Level 2 or Level 3 PRA is not envisioned, although of course if one is available then additional insights can be derived. This same comment applies to an SMA that is developed to support the same type of safety decisions.

The work to implement Step B has been accomplished numerous times using SPRA methods and a few times using SMA methods, and therefore there is confidence that the methods to implement this Step exist, are becoming more widely used, and hence are not a problem.

5.3 Step C: Identify Leading Accident Sequences and SSCs

Step C involves using the SPRA or SMA to identify the "leading" sequences and "leading" SSCs arising due to potential large earthquakes. A major objective of a design-stage SPRA or SMA is to develop an accurate understanding of the seismic "risk profile" of the new plant. As noted in an earlier section above, in practice this comes down mainly to identifying those accident sequences that are the most important contributors to "seismic risk," often defined in terms of an annual core-damage frequency (CDF) and an annual large-early-release frequency (LERF). There are other figures of merit that may be used as well. One way or the other, the analyst, perhaps guided by the regulatory authority, will identify the leading accident sequences. Within each seismic sequence, the PRA will identify the several SSCs that participate in that accident sequence. Again, in some of these sequences a human error will contribute or the sequence will include a so-called non-seismic failure, namely a failure that occurs randomly (meaning not associated with the earthquake).

The objective of Step C is the identification of these leading sequences and SSCs. For each, the seismic capacity, as measured by the probabilistic fragility curve, will be the way that the seismic failure is characterized, along with the specific failure mode, any common-cause failures that link this seismic failure in a correlated way with other failures, and any dependent or cascading failure issues, such as SSC #1's seismic failure causing the failure of SSC #2.

A further and important aspect of the PRA is that it can identify those few core-damage sequences that, because of various failures, lead to a so-called "large early release" rather than to a core-damage accident that does not entail such a release. Avoiding such large early releases is a major objective of the design of any NPP, of course.

The above three steps provide the analyst with the "seismic risk profile" of the plant design, which is the end-point of Step C here.

5.4 Step D: Compare the Risk Profile to a Safety Target

Step D involves comparing the risk profile and its seismic portion to an overall "safety target." The underlying reason for performing a design-phase PRA, including an SPRA or an SMA, is to provide the basis for a judgment as to whether the new design has an adequate measure of "safety." This involves more than one figure of merit to enable an evaluation of "how much" has been achieved in the way of "overall" safety. Nowadays in the US, this evaluation begins with a comparison to the NRC's "safety goal" targets in terms of impacts on public health (NRC 1986). However, these safety-goal targets are almost always met with lots of margin by new-plant designs. (However, the limitations of any design-phase PRA noted above in the introductory paragraphs of this Chapter 5 must be borne in mind.)

More important in practice, in terms of the adequacy of the seismic part of the design, is meeting the NRC target found in the Commission's policy (NRC 1993). Specifically, the designer of a new plant seeking design certification must demonstrate by analysis that the "plant-level HCLPF capacity" of the newly designed plant exceeds 1.67 times the Safe Shutdown Earthquake (SSE) ground motion, which is the design-basis earthquake. This plant-level HCLPF capacity is determined by working out the "plant-level fragility curve" and from that curve identifying the "HCLPF capacity point" that is roughly the 1 percent point on the mean fragility curve. Figure 5-1 shows this point on a typical seismic fragility curve. This analysis is done using standard analytical techniques (Kennedy and Ravindra 1984; Reed and Kennedy 1994; EPRI 1991; EPRI 2002), and the SPRA or SMA must meet the ASME-ANS PRA standard (ASME-ANS 2013) or follow ISG-20 (NRC 2010). If any shortfalls are found compared to this requirement, a design change is required.

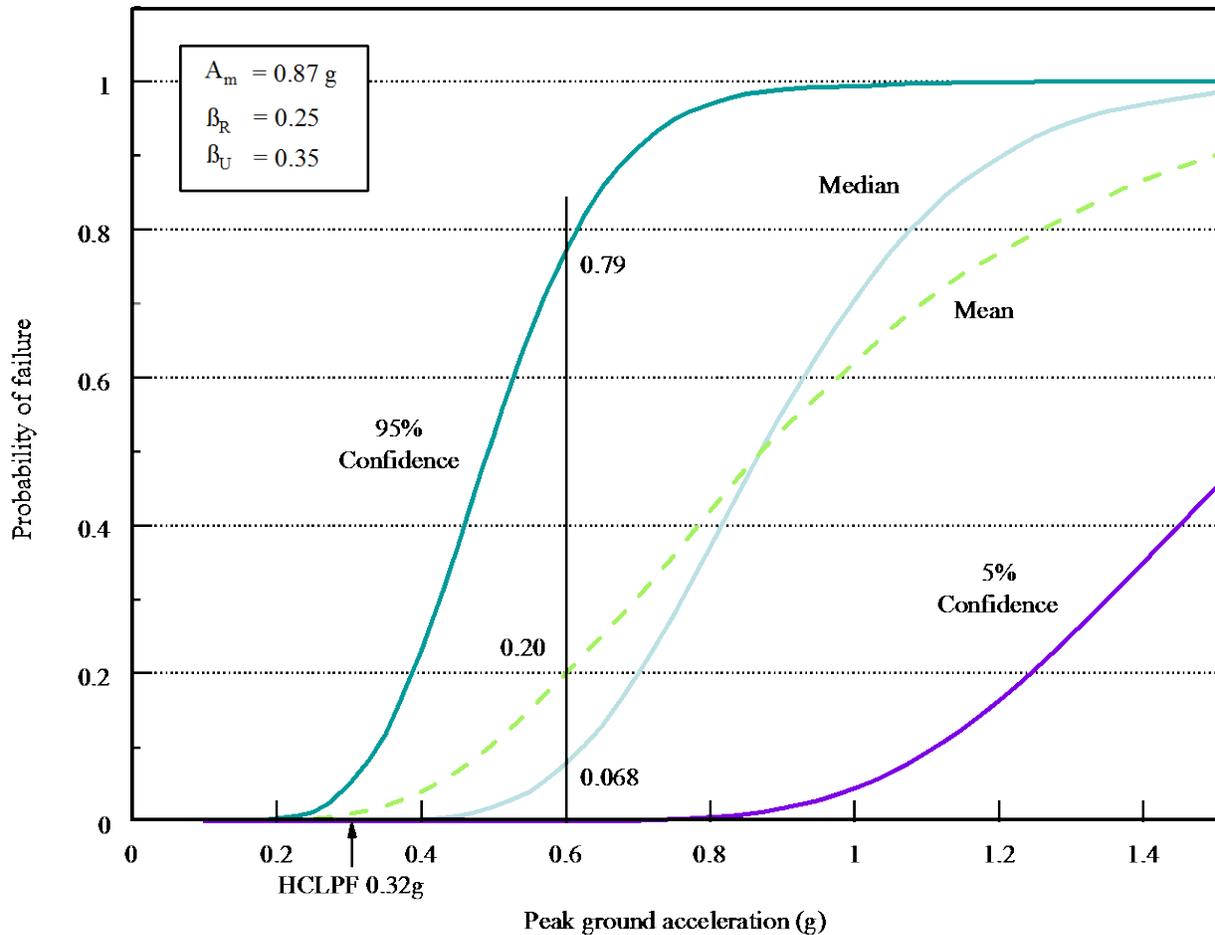


Figure 5-1 Seismic fragility curves and HCLPF point

5.5 Step E: Evaluation for an Imbalance in Defense in Depth

Step E involves evaluating the adequacy of the design vis-à-vis an imbalance in defense in depth. This Step cannot now be accomplished using existing regulatory policies, because no such policies exist.

There are two central problems or issues in relation to defense in depth. The first is that there are a number of different definitions of defense in depth. The second is that no criterion exists that can distinguish between a design for which defense in depth is adequate and one for which it is inadequate. Although these problems could lead to an impasse, the following discussion explores the issue and provides suggestions for a coherent approach in the context of the seismic design.

What is the issue or problem of defense-in-depth “balance” in the seismic safety area under discussion here? The issue is easy to explain, as follows. There are many SPRAs for NPPs in the literature. They contain a wide variety of conclusions and findings about the “seismic risk profile” of the specific plant being studied. As mentioned earlier, this “profile” consists of the set of the most important seismic-initiated accident sequences, for each of which the PRA has identified the (usually small number of) seismic-induced failures, human errors, and non-seismic failures that participate in that sequence. The PRA also provides a numerical core-damage

frequency (CDF) and/or a large-early-release frequency (LERF) for each such sequence. The summing of these represents the plant-wide CDF and LERF frequencies due to earthquakes.

The PRA will have provided an analysis or description of the analyst's estimate of the uncertainty in the numerical values of CDF and/or LERF. The PRA will also have given insights into which of the various failures and errors are most important, and can provide comparisons between the seismic-initiated sequences and similar sequences initiated by other causes such as internal plant faults, other external hazards, or human errors.

Taken as a whole, the SPRA can also help a decision-maker address whether the plant is "safe enough," whatever that might mean to the individual decision-maker. Here it could mean "safe enough against earthquakes" or "safe enough" in a broader sense.

Back to defense in depth: One major insight from the numerous seismic PRAs in the literature for NPPs is that sometimes – in fact, often – the seismic risk profile is dominated by a single seismic failure in addition to the loss of offsite power. This failure can be, for example, the earthquake-induced damage to a major building (the auxiliary building or the turbine building), or the loss of all service water. In the jargon of PRA, this is a "singleton." See Section 7 for a discussion of this, which is illustrated by some actual examples.

Suppose that the "leading" CDF seismic sequence is such a "singleton" sequence. Also suppose that, even so, the overall CDF (and LERF) is acceptably small, in a range that is within the "comfort zone" of safety decision-makers, whether in a regulatory agency or within the power plant's management. This fact, if true, would normally mean that the plant is "acceptably safe against earthquakes," whatever those words mean. However – and this is the point of this discussion – *the whole rationale for a defense-in-depth approach to plant safety is that the overall acceptability of the plant's safety should not rest on the strength of a singleton, nor rely on a single line of defense, nor be dependent on the correctness of the analysis of a single item, nor be overly optimistic about the understanding of uncertainties in the analysis.*

To put a fine point on it, assume now that due to an error of some kind, the singleton is actually not as strong against earthquakes as the PRA analysis says it is. There might have been an analysis error, the item as found in the field might be different from the item as analyzed from the drawings. Anchorage might be degraded, or fabrication errors might have occurred, or maintenance might have left the item in a degraded state, or the item might be vulnerable to failure due to the unrecognized potential for seismic damage caused by failure of another nearby structure or piece of equipment. Whatever the reason, the singleton is somehow not as strong as the PRA says it is.

This error, if true, would mean that *the bottom-line findings and insights of the entire seismic PRA are incorrect*. However, neither the analyst nor the decision-maker knows it. *This type of situation is exactly why the overall philosophy of defense in depth has been developed and deployed – to avoid the situation where a single error related to a single item (or to a single event, for example a multiple failure due to a common cause) means that the NPP is more vulnerable than analysts and decision-makers think it is.*

Such a plant is out of balance vis-à-vis defense in depth. To restore an appropriate level of defense in depth, the plant design or layout or operations must be modified.

The policy question remains: What criterion should be used to determine "how much imbalance" is an undesirable amount of imbalance? Might it be some sort of percentage figure of merit, like "defense-in-depth imbalance exists if a singleton sequence comprises more than 50 percent of

the seismic risk profile” or “more than 20 percent?” Might it be some sort of absolute CDF figure of merit like “defense-in-depth imbalance exists if a singleton accident sequence has a mean sequence-specific CDF greater than 1×10^{-5} /year or 5×10^{-6} /year”? A different policy matter is whether the CDF and LERF endpoints should be treated separately.

In today’s NRC regulatory scheme, there are no criteria beyond a general qualitative exhortation to seek out and avoid imbalances when they are identified.

The solution, as a practical matter, is for the NRC to initiate a technical debate on this subject, ultimately leading to some sort of policy guidance – perhaps a Commission policy position, perhaps a rulemaking, perhaps more general or less binding guidance. Among possible policy options might be to incorporate appropriate levels of conservatism where needed to account for the vulnerability, or to try to actually quantify defense-in-depth partially if not fully, at least as part of the uncertainty/sensitivity analysis. That something like this is needed is manifest if one accepts that many of the SPRAs in the literature contain singleton-type accident sequences as the leading (or nearly leading) contributors to the seismic risk profile. It goes without saying that this is also likely to be true of the many operating LWRs that have not yet been studied using SPRA methods.

As noted above, caution is necessary: Any evolution of policy in this area must differentiate between moving toward a more “balanced” risk profile by decreasing the overall absolute risk and moving toward it by changing the relative contributions of different “risk contributors.”

In summary, a new policy proposal is advanced herein, that would bring to bear a more fully developed defense-in-depth philosophy to the design and analysis of the seismic performance of a new NPP. It would be aimed at reducing or eliminating a situation in which a new NPP’s seismic risk profile is out-of-balance because the profile is unduly dominated by a seismic singleton – a single seismic-induced failure which makes an inordinately large contribution to the overall seismic risk arising from the plant. To achieve the goal outlined here would require new NRC policy development, in the form perhaps of a policy statement, a regulatory guide, or an actual regulation. How the form of this policy development would evolve is beyond the scope of the proposal described here.

5.6 Implementation

If the defense-in-depth criterion in Step E is adopted by the NRC as a safety requirement, or even as guidance, the way in which it is formulated is important. One possibility is that the NRC will judge this requirement as *necessary to provide adequate protection*. In that case, the plant design must comply. Another possibility is that the defense-in-depth criterion is promulgated as a requirement but on a *safety-enhancement basis*, meaning that the plant’s safety level is judged to meet the adequate-protection criteria, and therefore a safety improvement would only be required if it could meet the NRC’s *cost-benefit criteria*, as promulgated by the NRC (NRC 2004a). Either way, the regulatory foundation for this is well established, and therefore this Step can be performed using existing methods. A third possibility is that the defense-in-depth issue discussed above in Step E is to be dealt with by following NRC *guidance*, which would be non-binding. This approach would leave implementation to the operating power plant’s management, but accompanied by an exhortation or by specific guidelines. How those would be formulated is all for the future, of course.

The following subsections discuss several important implementation issues.

5.6.1 Upgrading Options

In dealing with accident sequences initiated by large earthquakes, it is important to understand that a safety improvement need not involve upgrading of the seismic adequacy of a structure or component. From the SPRAs performed for our existing NPPs, a major insight is that a significant fraction of the important seismic accident sequences involve one or more seismic-induced failures in combination with one or more human errors or non-seismic (random) equipment failures. Hence, a safety improvement to decrease either the frequency or the consequences of one of these seismic-initiated accident sequences might involve improving the (non-seismic) reliability of equipment, or decreasing the likelihood or the consequences of one or more human errors – errors in the control room, errors elsewhere in the plant, etc. The designer can use the PRA to support such a decision on a case-by-case basis. This emphasizes how important it is for the analyst to take into account the *entire plant risk profile* when considering which are the most efficacious options for upgrading the seismic risk profile. Not considering the entire plant-as-a-whole runs the risk that an upgrade of the seismic safety will be sub-optimal for the plant-as-a-whole.

5.6.2 Using Tradeoff Analysis Tools

Various cost-benefit analysis tools exist, so the techniques required for this Step are well exercised and straightforward. The most obvious tradeoffs are of two kinds: (a) tradeoffs between two or more different ways to accomplish a change in the plant's configuration (or procedures) to enhance safety, or (b) tradeoffs involving costs (usually financial costs) vs. benefits (usually safety benefits). This latter tradeoff issue is governed by established NRC guidance (NRC 2004a), whereas the former is accomplished using often-used and well-established engineering methods. Therefore, implementing this should not be a problem.

5.6.3 An Iterative Process

The process above, if implemented properly, will often end up in an iterative cycle. Specifically, Steps A, B, and C produce a design and an analysis of it. Steps D and E are evaluations of the design leading to proposals for possible design changes. After this process has run its course, the analysts need to go back to Steps B and C to do a new analysis, leading to a new evaluation.

This iterative process is a natural one in all design work, of course. Therefore, in a way, nothing is new here. However, what is novel is the use of risk information (but not exclusively!) and the need to balance that information against other non-risk considerations. As is now widely recognized, this is exactly what the adjective “risk-informed” is supposed to describe.

5.6.4 Issues with a Firm Decision Criterion

A problem, albeit not a new one, arises in iterative design if there is no firm bright-line decision criterion differentiating when a design change is needed from when it is not, based on insights from the iteration. Such a criterion would presumably be a threshold imposed by (or suggested by) the regulatory agency, perhaps along the lines discussed in Section 5.5 (for example, a limit on CDF, or on a fraction of CDF, or on an absolute contribution from a given accident sequence). There are benefits and liabilities to such a bright-line criterion. The benefit of such a criterion, specified in advance, is that decision-making would be more uniform and less susceptible to interpretation. The downside could be that the criterion might not anticipate all of

the issues, some of which could be complicated, if the technical situation is multi-faceted. In that case, such a criterion could be more constraining than liberating and less a help than a hindrance.

On the other hand, the absence of a specific criterion could be a recipe not only for variations in outcomes based on judgments (which will vary from one decision-maker to another), but also for controversy about those judgments. An intermediate position might be that the decision criterion is promulgated as a suggestion or as guidance but not as a regulatory constraint. But supposing that such a decision criterion is not ultimately adopted in regulation even as only a suggestion or as guidance, there would be benefit merely from having done the analysis to call attention to this issue. Specifically, there is a potential benefit to plant safety even if a seismic risk-profile imbalance has only been noticed but not remedied, because plant safety decisions later on can be made in cognizance of this imbalance.

6 OPPORTUNITIES FOR FURTHER ADVANCES – PATHS FORWARD OVER THE LONGER TERM

Besides the opportunity for an advance represented by the “defense-in-depth imbalance” issue discussed above in Section 5, there are a few other shortcomings in the current Framework where an opportunity exists for an advance. Each of these was discussed briefly in the earlier part of this report. They will be discussed briefly again next, in the context of possible paths forward.

For each of these issues, the path forward is likely to be more difficult, and therefore likely to take longer, than for the few issues discussed above in Section 5. For some of these issues, the technical background provides a firm basis for moving forward, although important policy issues need resolution. For others of these, there are still technical unknowns or uncertainties that require research to develop the needed new technical knowledge.

However, each of these issues needs to be resolved, sooner or later, in order to bring the Framework for seismic design and safety regulation closer to the “ideal” Framework described in Section 3 above.

6.1 Core Damage vs. Large Early Release

The NRC “target” for the seismic adequacy of new plants (plant-level HCLPF > 1.67 times SSE) applies to core damage but not to large early release. The fact that the NRC has promulgated a requirement for the seismic adequacy of new NPPs is an admirable step forward. It was adopted many years ago (NRC 1993) and, in fact, broke new ground in that no comparable requirement seems to exist in any of the several other areas of safety. Nevertheless, it seeks a plant-level HCLPF seismic capacity against the core-damage endpoint while ignoring the large-early-release endpoint.

This is obviously a gap, and filling that gap would be an advance. If promulgated in more-or-less the same way, in terms of a plant-level HCLPF seismic capacity, it would provide a second plant-level design target, to be used in parallel with the existing one. An advantage is that the ability to analyze such a HCLPF capacity is well within the capabilities of today’s SPRA analysis experts.

This is an area ripe for debate followed by the development of an NRC regulatory position.

6.2 Confidence Level

There is no explicit declaration of what confidence level an analysis must meet vis-à-vis the regulatory requirement(s) in order for the design being analyzed to be in regulatory compliance. This statement is not quite correct – the requirement for the plant level HCLPF capacity to exceed 1.67 times the SSE ground motion (NRC 1993) has an embedded confidence level because the HCLPF capacity is the point (in terms of earthquake “size”) on the fragility curve at which the analyst has about 99 percent confidence that success will occur. However, for none of the other NRC acceptance criteria for seismic safety is a confidence level provided and required. The typical staff regulatory judgment is based on having something that they loosely call “high confidence,” and the typical analysis requires the reporting of mean values, which tend in this area to be at or above one standard deviation above the median, meaning 85 percent.

But the mean value can sometimes be near the 90 percent confidence level or even higher. However, there seems not to be an explicit confidence level used.

An ideal Framework, as noted in Section 3, would provide an explicit confidence level as an integral part of the regulatory acceptance criterion.

This is an area ripe for debate followed by the development of an NRC regulatory position.

6.3 Providing an Option to Use a Performance Target

There is no provision today allowing the designer the option of designing explicitly against a performance target instead of using the design codes. A provision along these lines pre-supposes the existence of a regulatory “performance target.” That is perhaps the more difficult issue, but again seems amenable to both debate (extended if necessary) and then promulgation of a regulatory position. The sticking point is probably attaining confidence about the analysis to be done to demonstrate compliance with the target.

Nevertheless, the benefits of such a formal regulatory position are manifest. Even if the burden of proof is high indeed for the applicant, meaning that invoking this option would be rarely done, it seems advantageous to allow such an option. An example will perhaps support this line of argument: The containment structure for a new reactor might cost as much as a few hundred million dollars. Today’s designer-applicant must “meet the code” in all respects. Suppose that a better design solution – both safer and substantially less expensive – has been identified by the designer, but it doesn’t “meet the code” for one or another reason. In an ideal Framework, the designer would have the option of demonstrating that the design is safer overall, with an extremely high hurdle in terms of the burden-of-proof. Nevertheless, if it cost a few million dollars of testing and a comparable amount in analysis but could save much more than that in construction cost, and if the design is manifestly safer too, this would be a “win-win” for everyone, yet today’s Framework essentially precludes it.

There are, of course, good reasons to be somewhat skeptical of how far one can “push” a performance-target approach. It certainly should not be the sole or predominant basis for a decision on the adequacy of safety achieved, if only because that requires more hubris than is merited by today’s methodologies for analyzing performance. Nevertheless, moving in this direction seems to be well worth careful consideration on a case-by-case basis.

For what seem to be good reasons, including history, today’s Framework unduly stifles innovation in the name of consistency and predictability. There is really no reason in principle not to expand the options available under the regulations in the way discussed above.

Again, this is an area ripe for debate followed by the development of an NRC regulatory position.

6.4 Systems-analysis Issues

In trying to optimize the design, a designer has no explicit way in the current regulatory Framework to take into account the fact that non-seismic failures and human errors contribute to many important seismic accident sequences. Also, the fact that seismic failures of equipment (especially but not limited to identical or similar co-located SSCs) may fail in a correlated way is not accounted for either in the design guidance or in any regulatory positions. The SPRA methodology accounts explicitly for the fact that in many important seismic-initiated accident

sequences, one or more seismic failures or human errors contribute to the sequence. The SPRA methodology also accounts for the extent to which correlations in the response or fragility of co-located identical SSCs matter after a large earthquake. The fact that there is important uncertainty in this analysis shouldn't obscure the fact that the analysis provides useful insights.

Given the above, and given a design, an SPRA can reveal whether any of these contributions is important to the accident sequence, and if so how. This information, if available to the designer, would allow an optimal solution to be identified in terms of how best to reduce either the likelihood or the consequences of the given accident sequence. However, the current regulatory Framework neither asks for such an analysis explicitly, nor has any context for how to deal with the information if it were made available. In sum, neither the applicant/licensee nor the regulatory reviewer has any guidance on this set of subjects. Yet these issues can substantially affect the seismic risk profile.

An advanced Framework would not only recognize the existence of these issues as contributing to the risk profile but would provide guidance to the designer about how to use the information to improve seismic safety, or at least to demonstrate the achieved level of safety with greater confidence. The counterpart regulatory reviewer, in turn, has no guidance about how to use the information if it were presented. The guidance to both parties might take the form of citing acceptable analysis methods (presumably citing the ASME-ANS PRA standard) (ASME-ANS 2013), supplemented by provisions for using the information to support a decision to permit a design configuration that does not otherwise meet the regulations, if it can be shown with confidence that the proposed modification surely enhances safety.

Again, this is an area ripe for debate followed by the development of an NRC regulatory position.

6.5 SSCs Are Designed Individually for Seismic Performance

As discussed earlier (see Section 2.3), each SSC within an NPP currently is designed individually to achieve an adequate seismic capacity. To advance beyond this to account for the design and the safety of the plant-as-a-whole will likely require a consensus code committee to address head-on how to use the systems insights from SPRA to modify the design guidance to account for the broader issues. The general process to be used here is similar to the multi-step process outlined in Section 5, which covers how to go about addressing an unbalanced seismic risk profile. Specifically, one would complete the design using normal processes, then perform a full PRA that includes a seismic PRA, and then use insights from the risk profile of the plant-as-a-whole to account, where appropriate, for safety issues that the SPRA indicates are worthwhile considering. This might lead, for example, to the desire to strengthen the seismic design of one or more SSCs or alternatively to a possible lesser emphasis on certain other SSCs.

An approach like this could be part of the basis used by a consensus committee, charged with a seismic design code, to modify its approach to account for these types of plant-as-a-whole safety insights. To form the basis for the work of such a consensus committee, perhaps an individual researcher could take on the issue by writing a policy options paper that would become a focal point for discussion across the affected industry and the major professionals involved. The researcher could be a member of the NRC staff, or from an industry, academic, or consultant background. Perhaps an affected party could convene a workshop in which interested individuals can explore the issues thoroughly. Perhaps the NRC or EPRI could support a research project along these lines. It seems unlikely that progress can occur absent

one or another of these precursor activities. Ultimately, if any modification to the existing “framework” were to be proposed, it would be necessary to assure consistency with the NRC’s current approach in 10 CFR 50.69 (NRC, 2004) to risk-informed categorization and treatment of SSCs.

6.6 Variations in Margin Among Industry Design Codes

As discussed earlier in Section 2.4, “[t]he design of any individual SSC for service in a nuclear power plant is generally governed by some industry code or standard that has usually been endorsed by the NRC for its particular application. These consensus codes are developed and maintained by a number of code committees organized under different standards-development organizations (e.g., American Nuclear Society, American Concrete Institute, American Society of Civil Engineers, Institute of Electrical and Electronic Engineers, and others). These codes, which generally rely on an externally-specified design-basis ground motion as the starting point for the design, use a variety of different approaches to address the question of how much seismic margin above the design basis should be embedded in an SSC designed using the code. That there are differences in embedded margin for SSCs designed using various codes is not surprising, given that the code committees have generally worked independently and that the consensus codes represent different philosophies of design that exist in the different fields of engineering. This is not a criticism of the work of the various code committees. It is merely an observation that there are differences in approaches and hence in the outcomes.

To advance beyond this, so as to achieve greater harmony across the various code committees leading to greater uniformity in the margins achieved, will likely require one of two catalysts: either an overarching body like the NRC could enforce some consistency of approach, or a consortium of the major code committees could somehow get together to bring this about. While the former would be “cleaner” and perhaps “easier” administratively, the latter might produce a result that ultimately has broader overall stature and staying power. To provide the intellectual framework, perhaps an individual researcher could take on the issue by writing a policy options paper, or perhaps a workshop could be assembled involving all of the major technical stakeholders, to try to work out an agreed approach. Absent something like one of the above, a consistency of approach is unlikely to develop anytime soon.

6.7 Broad Institutional Barriers

There are two broad institutional barriers to making sweeping progress on the agenda put forth in this report. One involves the NRC and the other involves the code committees.

6.7.1 NRC Regulations

How ripe is the NRC for a reassessment and revision to embed more risk-informed and performance-based requirements into the seismic regulatory Framework, along the lines of this report? There has been major recent progress in the direction of overall new policy development at the NRC in this area. The recent publication in April 2012 of NUREG-2150, “*A Proposed Risk Management Regulatory Framework*,” (NRC 2012a) not only lays out a roadmap for several significant new NRC policies, but it explains how they would fit together into a coherent agency-wide framework. This is, of course, only the beginning of a long process involving significant interactions with the regulated industry and the public. However, there is no doubt that over the long term this is the direction that much new regulatory policy development will be following.

One issue of importance is whether the NRC's possible future approach will seek to differentiate between seismic design and analysis aimed at preventing core-damage accidents vs. aimed at preventing large early radioactive releases. This issue has not yet been joined through any extensive debate.

Recently, the NRC has sought to embed more risk-informed and performance-based approaches into their regulations and other regulatory positions. This bodes well for the future, albeit with some cautions. Clearly, a good deal of debate is required, and appropriate, before a consensus can emerge on which policy advance(s) might be best.

6.7.2 Code Committees

How ready are code committees for revisions along the lines of this report to their deterministic approaches to seismic design and analysis? Throughout the community of experts, the philosophy has been changing, slowly but inexorably, for years. Now this issue is ripe for even more movement toward performance-based design and analysis, if not risk-informed too, which is more difficult. It will take some leadership to follow up on the ideas embedded in NRC's NUREG-2150 report (NRC 2012a). One major piece of progress occurred when ASCE 43-05 paved the way. A new version of ASCE 4 (that will supersede ASCE 4-98) is coming out soon and perhaps it will take the next steps (ASCE 2013; Budnitz 2013). However, what is really needed is a common approach across the several different standards development organizations that are active in this arena, including the American Nuclear Society, American Society of Mechanical Engineers, American Concrete Institute, American Society of Civil Engineers, Institute of Electrical and Electronic Engineers, and others. This will take industry-wide leadership that probably can only come from the NRC, although it is possible that it might emerge from a consensus of industry experts instead, or in addition. No forum now exists, however, that could sponsor or encourage the discussions that might lead to such a consensus of leading experts.

The goal ought to be not only that each major industry consensus design and analysis code embed the advanced philosophy discussed here, but also that there be coordination, so that the different technical areas work together to achieve comparable outcomes in terms of design and analysis of different categories of SSCs.

6.7.3 NRC and Industry Working Together

How ready are these two major stakeholders to work together? As discussed earlier (see Section 2.5, Section 3, and Section 4), the industry consensus committees and the NRC do not yet "work together" to achieve a prescribed safety target, in major part because the NRC's current safety targets (see the discussion above in Section 5.4) are not written in a form amenable for direct use by the code committees, and in part because working together has not historically been the pattern – the various code committees themselves have seldom seen either the need or the motivation to put in place uniformity in the safety achieved. And old habits die hard.

To advance beyond this will require an explicit decision, probably by the NRC staff with policy input from management (including the Commissioners), that achieving the desired uniformity is important enough to give priority to the effort required.

Fortunately, as noted above in Section 2.5, "... today there has been an evolution on both fronts: (i) First, in the philosophy of NRC regulation, which is continually becoming more risk-informed and performance-based, and which has a goal to integrate design, construction, maintenance, and inspection activities more fully; and (ii) second, in the ability of the code committees to embed an approach that could better integrate into an ideal regulatory framework with the confidence that both analysis of the outcome of the design and a comparison against a target are feasible."

Crucially, one of this report's major objectives has been to provide detailed analyses of several technical and policy issues as input to future deliberations by not only the NRC (Commissioners and staff) but also the regulated industry and the public.

6.8 Relative Roles of Risk from Earthquakes vs. Other Initiators

This issue has been a source of contention for a very long time, at least since the time when the advent of PRA methods allowed analysts to determine the relative risks arising from different types of accident classes. This is known as the "risk allocation" problem.

Those in favor of some sort of allocation typically argue that it would lead to a more balanced design; that it would provide the specialists in each of the safety areas (for example, the fire-safety area vs. the seismic-safety area) with their own targets to work toward; and that it would make understanding the total risk profile easier for the public. Each of these seems persuasive on the surface. However, the counter-arguments are equally persuasive: Allocation ignores the fact that many safety design questions and safety-improvement opportunities cut across the areas that would be given separate allocations, making the achievement of overall safety sub-optimal; that in any given technical area, there might be a tendency to "stop" when the target is reached on the basis that the design is now "safe enough" in that area; that the PRA analyses of the safety level achieved are in any event numerically quite uncertain (and will likely always be!), so that such comparisons are misleading if used as a major basis for a "safe-enough" decision; and that simplifying the risk profile for the public is in fact a major over-simplification of a complicated interconnected technical problem.

On balance, of all the issues raised in this report, this issue of risk allocation seems to be the one with the least prospect for any near-term change from the *status quo*. Its benefits are controversial, its liabilities are difficult to overcome, and it will therefore likely remain in the future as a controversial but unsolved potential possibility.

7 CASE STUDIES

7.1 Introduction

Previous Sections of this report identified attributes of an ideal regulatory Framework for seismic safety and, subsequently, compared and evaluated the current nuclear regulatory Framework against the ideal. This evaluation revealed that several attributes fall short of ideal. In particular, seismic probabilistic risk assessments (SPRAs) have revealed that some NPPs have unbalanced risk profiles, which often translates into inadequate or insufficient levels of defense in depth. Specifically, the seismic failure of a single structure, system, or component (SSC), or of a very small number of SSCs, contributes disproportionately and dominates the seismic risk profile. This situation is inconsistent with the NRC defense-in-depth philosophy that strives to achieve a design in which the risk profile is more “balanced” among a variety of different contributors. In addition, under the current regulatory Framework, SSCs are designed individually for seismic performance. In general, the systems view of the plant does not play a role in how individual SSCs are now designed and regulated against earthquakes, which can result in NPPs with unbalanced seismic risk profiles.

This Section presents two case studies that clearly illustrate these particular shortcomings. Each case study is based on the design and analysis of an operating NPP; however, certain aspects of the design and of the SPRA have been stylized and simplified in order both to protect the plants’ identities and to enable an explanation of the findings without unneeded detail. Therefore, certain technical details about the two plants, which are here called Plant A and Plant B, do not accurately reflect their actual configurations or risk profiles. However, the case studies are representative enough of each plant to illustrate the general points raised in this report.

Each case study considers three different plant designs or configurations. The first configuration represents baseline conditions at the NPP – in other words, the plant as it was at the time of the SPRA analysis. Because the baseline configuration is based on the design of an operating NPP, it is assumed to satisfy all the requirements of the current nuclear regulatory Framework. The second and third configurations that we have studied are adaptations or reconfigurations of the baseline configuration. In particular, the second configuration adjusts the baseline design so that all SSCs satisfy the provisions of ASCE 43-05 (*i.e.*, that each SSC designed according to the standard has an annual probability of failure from earthquakes less than 1×10^{-5}). The third configuration adjusts the baseline design in order to satisfy the suggestions related to defense in depth described in previous Sections of this report (*i.e.*, no single accident sequence contributes more than 25 percent to the seismic CDF at the plant, or the annual probability of occurrence for an individual seismic accident sequence does not exceed 5×10^{-6} , both selected somewhat arbitrarily for the purposes of this case-study analysis). Subsequently, a simplified SPRA is performed for each of the three configurations to determine its risk profile and study the impact of changes made to the baseline configuration.

The discussion that follows is organized into three main Sections. Section 7.2 provides generic background information for both case studies, including lists and descriptions of the SSCs and accident sequences that play important roles in the seismic safety of the plant. This information effectively describes the baseline configuration of the NPP in each case study. Section 7.2 also presents results from a simplified SPRA of each baseline configuration. These results help demonstrate the need for the adjustments made in the second and third configurations. Section 7.3 presents in detail the second and third configurations for both case studies, henceforth referred to as the “ASCE 43-05” configuration and the “enhanced balance” configuration,

respectively. In particular, for each configuration, Section 7.3 describes the adjustments made to the baseline configuration and, using a simplified PRA, analyzes their impact on the seismic risk profile of the plant. Lastly, Section 7.4 discusses the implications of the analyses of the two case studies and three configurations. The main findings, which are described in more detail in Section 7.4, are summarized below:

- The provisions of the current nuclear regulatory Framework do not necessarily ensure that an NPP will have a “balanced” seismic risk profile. In both case studies, one or two accident sequences (or SSCs) dominate the seismic risk profile of the NPP, leaving it more vulnerable to errors in design, analysis, construction, operation, or maintenance, or a failure to characterize the uncertainties appropriately.
- While the provisions of ASCE 43-05 can enhance the safety of individual SSCs, they do not necessarily produce an NPP with a “balanced” seismic risk profile.
- A provision that limits the fractional contribution of a single accident sequence to the seismic risk profile would help produce a more “balanced” plant seismic risk profile, but a balanced risk profile should not take precedence over a reduction in overall risk (i.e., balance is secondary).
- The more effective approach to producing a more “balanced” plant seismic risk profile involves establishing performance requirements for accident sequences rather than for individual SSCs.

Before this discussion can take place, several important limitations need to be acknowledged:

- The SPRAs described herein use core damage as the undesired endpoint of the analysis. Another important endpoint is large early radioactive release; however, this endpoint is not considered or discussed herein, even though as a policy matter it is of great importance. In spite of this, none of the insights and policy imperatives derived from the two case studies is affected by this omission.
- The accident sequences analyzed in the two case studies, and the insights illustrated by them, concentrate on sequences in which all SSC failures are caused by the earthquake. Based on insights from a large number of SPRAs, many of the important accident sequences in SPRAs involve one or more earthquake-induced failures along with one or more non-seismic failures and/or human errors. However, these types of accident sequences are not considered here. Neither are the potential safety insights that might arise from dealing with the non-seismic failures and/or human errors. Specifically, it is often the case that the most effective or efficient way to reduce the frequency or consequences of such an accident sequence involves improving the non-seismic failure probability or the human-error probability rather than improving the seismic capacity of the SSCs involved. These issues, however, are not illustrated in the two case studies examined herein.
- Trying to reduce the risk profile of a plant by making a single plant modification (or a small number of them) to reduce a single accident sequence can be an unwise approach if done in isolation. This is because changes to one part of a plant affect more than a single accident sequence. Therefore, actual decisions regarding such changes need to account for their impact on the entire plant design and operation, which goes far beyond only the

safety aspects captured in the SPRA. Specifically, if several plant changes are contemplated, it is necessary to analyze the entire set of proposed changes as a package, because plant modifications often interact.

- As discussed in Section 3.1 (see Attribute (B)), seismic-initiated accident sequences, by themselves, should not necessarily be the focus when using a PRA to analyze the overall risk profile of an NPP. This is a major limitation to the argumentation here; specifically, the purview for both case studies has been limited to a select set of earthquake-initiated accident sequences within the seismic portion of the overall PRA.

7.2 Background for the Case Studies

This Section provides basic information about the baseline configuration for each of the two case studies. The first case study analyzes Plant A, which is based on an operating NPP located in the eastern United States. The second case study analyzes Plant B, which is based on an operating NPP located in the western United States. Figure 7-1 plots the mean seismic hazard curve for each plant. These two locations were chosen to provide diversity in seismic hazard (*i.e.*, low and high seismicity) and, ultimately, to enhance the robustness of the findings. Section 7.2.1 describes the attributes of Plant A needed for the analysis herein, while Section 7.2.2 describes the attributes for Plant B.

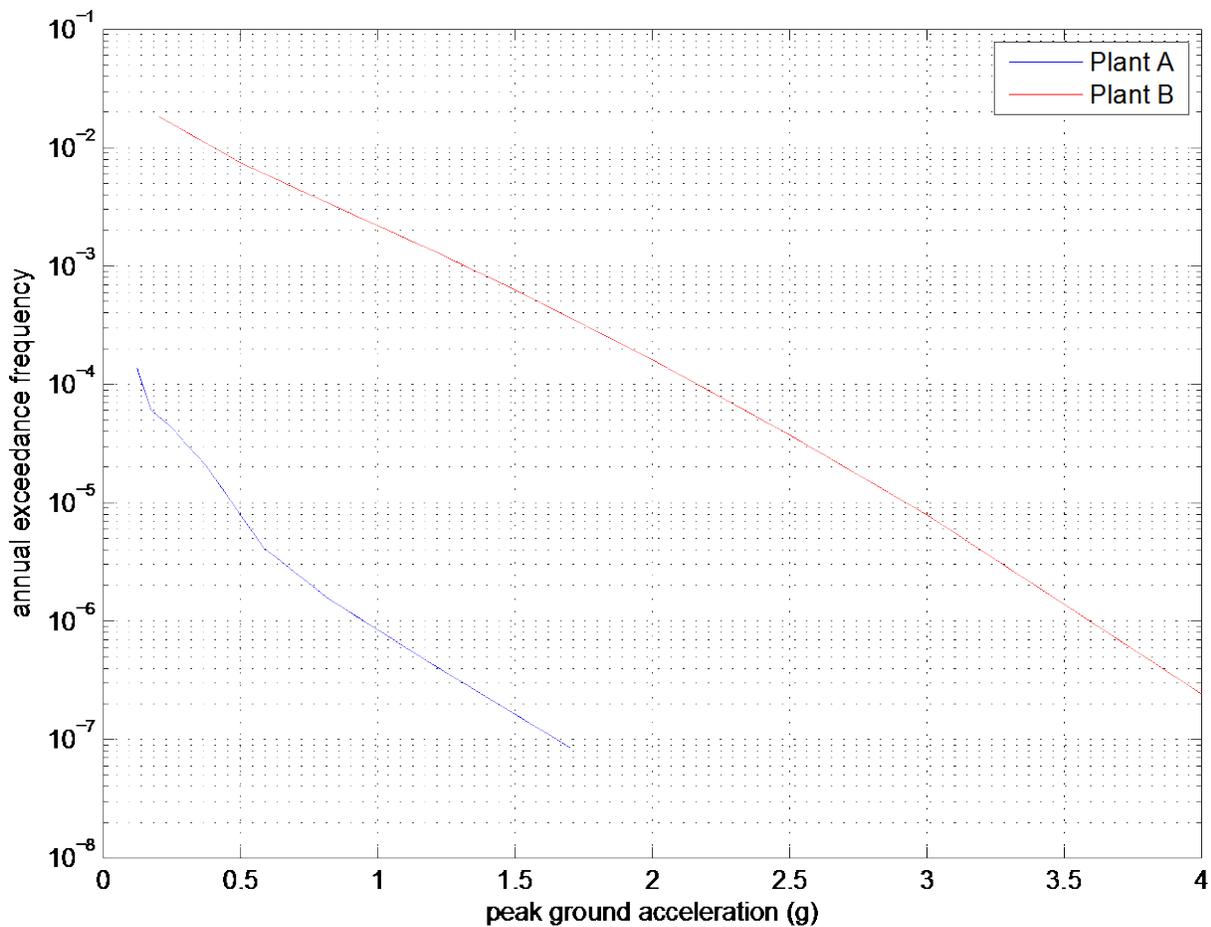


Figure 7-1 Mean seismic hazard curves for both case study sites

7.2.1 Plant A

Table 7-1 lists and describes the SSCs included in the analysis of Plant A. To help ensure anonymity, SSCs are categorized into three general groups: structural, mechanical, and electrical. Table 7-1 lists basic fragility parameters for each SSC, including median seismic capacity, A_m , and uncertainty parameters, β_r and β_u . Table 7-2 lists and describes the dominant seismic accident sequences included in the analysis of Plant A as derived from the SPRA.¹⁰ These sequences are combinations of failures of the SSCs listed in Table 7-1. Together, Table 7-1 and Table 7-2 describe the baseline configuration of Plant A. This baseline configuration represents the NPP as it was originally designed and analyzed, therefore it is assumed to satisfy the provisions of the regulatory Framework that were in effect at that time. In total, there are 8 SSCs and 9 seismic accident sequences in the analysis of Plant A.

Table 7-1 and Table 7-2 also summarize the results of a simplified SPRA of the baseline configuration. In particular, Table 7-1 lists the high confidence of low probability of failure (HCLPF) capacity (in units of g) and the annual probability of failure, P_f , for each SSC, which is calculated by a convolution of the SSC's fragility curve with the relevant seismic hazard curve. Note that several SSCs fail to satisfy the provisions of ASCE 43-05 (*i.e.*, an annual probability of failure less than 1×10^{-5}). Similarly, Table 7-2 lists the HCLPF value, annual probability of occurrence, and contribution to total seismic core-damage frequency (CDF) for each accident sequence. The seismic CDF for the baseline configuration of Plant A is 3.04×10^{-5} per year. This number will be referred to as the baseline seismic CDF for Plant A. Note that a single seismic accident sequence, SEQ1, contributes disproportionately to the plant's seismic risk profile, accounting for approximately 40 percent of the seismic CDF.

Table 7-1 Properties of SSCs in the baseline configuration of Plant A

SSC	A_m (g)	β_r	β_u	HCLPF (g)	P_f
MECH1	0.77	0.25	0.22	0.35	2.54×10^{-6}
MECH2	0.68	0.18	0.32	0.30	4.18×10^{-6}
STRUC1	0.32	0.07	0.20	0.20	2.62×10^{-5}
MECH3	0.13	0.24	0.32	0.05	1.75×10^{-4}
MECH4	0.68	0.30	0.30	0.25	5.04×10^{-6}
ELEC1	0.30	0.27	0.40	0.10	4.62×10^{-5}
ELEC2	0.69	0.23	0.36	0.26	4.88×10^{-6}
STRUC2	0.53	0.23	0.42	0.18	1.20×10^{-5}

¹⁰ The actual SPRA was used as the basis for this work (and for the similar study of Plant B), but modest changes and approximations have been made to simplify the analysis and also to help mask the identities of the two plants.

Table 7-2 Properties of dominant seismic accident sequences in the baseline configuration of Plant A

Sequence	HCLPF (g)	P_f	% of seismic CDF
SEQ1	0.17	1.20×10^{-5}	39.37%
SEQ2	0.34	3.14×10^{-6}	10.32%
SEQ3	0.36	2.85×10^{-6}	9.35%
SEQ4	0.38	2.79×10^{-6}	9.18%
SEQ5	0.34	2.40×10^{-6}	7.88%
SEQ6	0.38	2.00×10^{-6}	6.59%
SEQ7	0.40	1.97×10^{-6}	6.48%
SEQ8	0.42	1.74×10^{-6}	5.73%
SEQ9	0.45	1.55×10^{-6}	5.10%
CDF	--	3.04×10^{-5}	100.00%

Figure 7-2 plots the sensitivity of the seismic CDF to changes in the seismic capacities of individual SSCs. More specifically, it plots the absolute change in seismic CDF produced by changing the median capacity of each individual SSC, one at a time. As can be seen in the figure, some SSCs have more impact on the plant's seismic CDF than others. The impact of a particular SSC is a function of several factors, including its overall strength/capacity (as reflected by its fragility parameters) and also its particular role in the overall system (*i.e.*, the accident sequence or sequences in which it participates).

There are several interesting observations that emerge from Figure 7-2. First, strengthening the most fragile SSC in the analysis, MECH3, has almost no impact on the plant's seismic CDF, implying that simply strengthening the weakest SSCs in the plant may not be the most effective way to improve overall plant safety. Second, the plant's seismic CDF is disproportionately sensitive to changes in the strength of STRUC2. A 30 percent decrease in median capacity produces almost a 60 percent increase in seismic CDF, meaning that an error in the design, analysis, operation, or maintenance of STRUC2 that reduces its capacity can significantly impact the safety of the plant. This heightened sensitivity to a single SSC is inconsistent with the NRC's defense-in-depth philosophy.

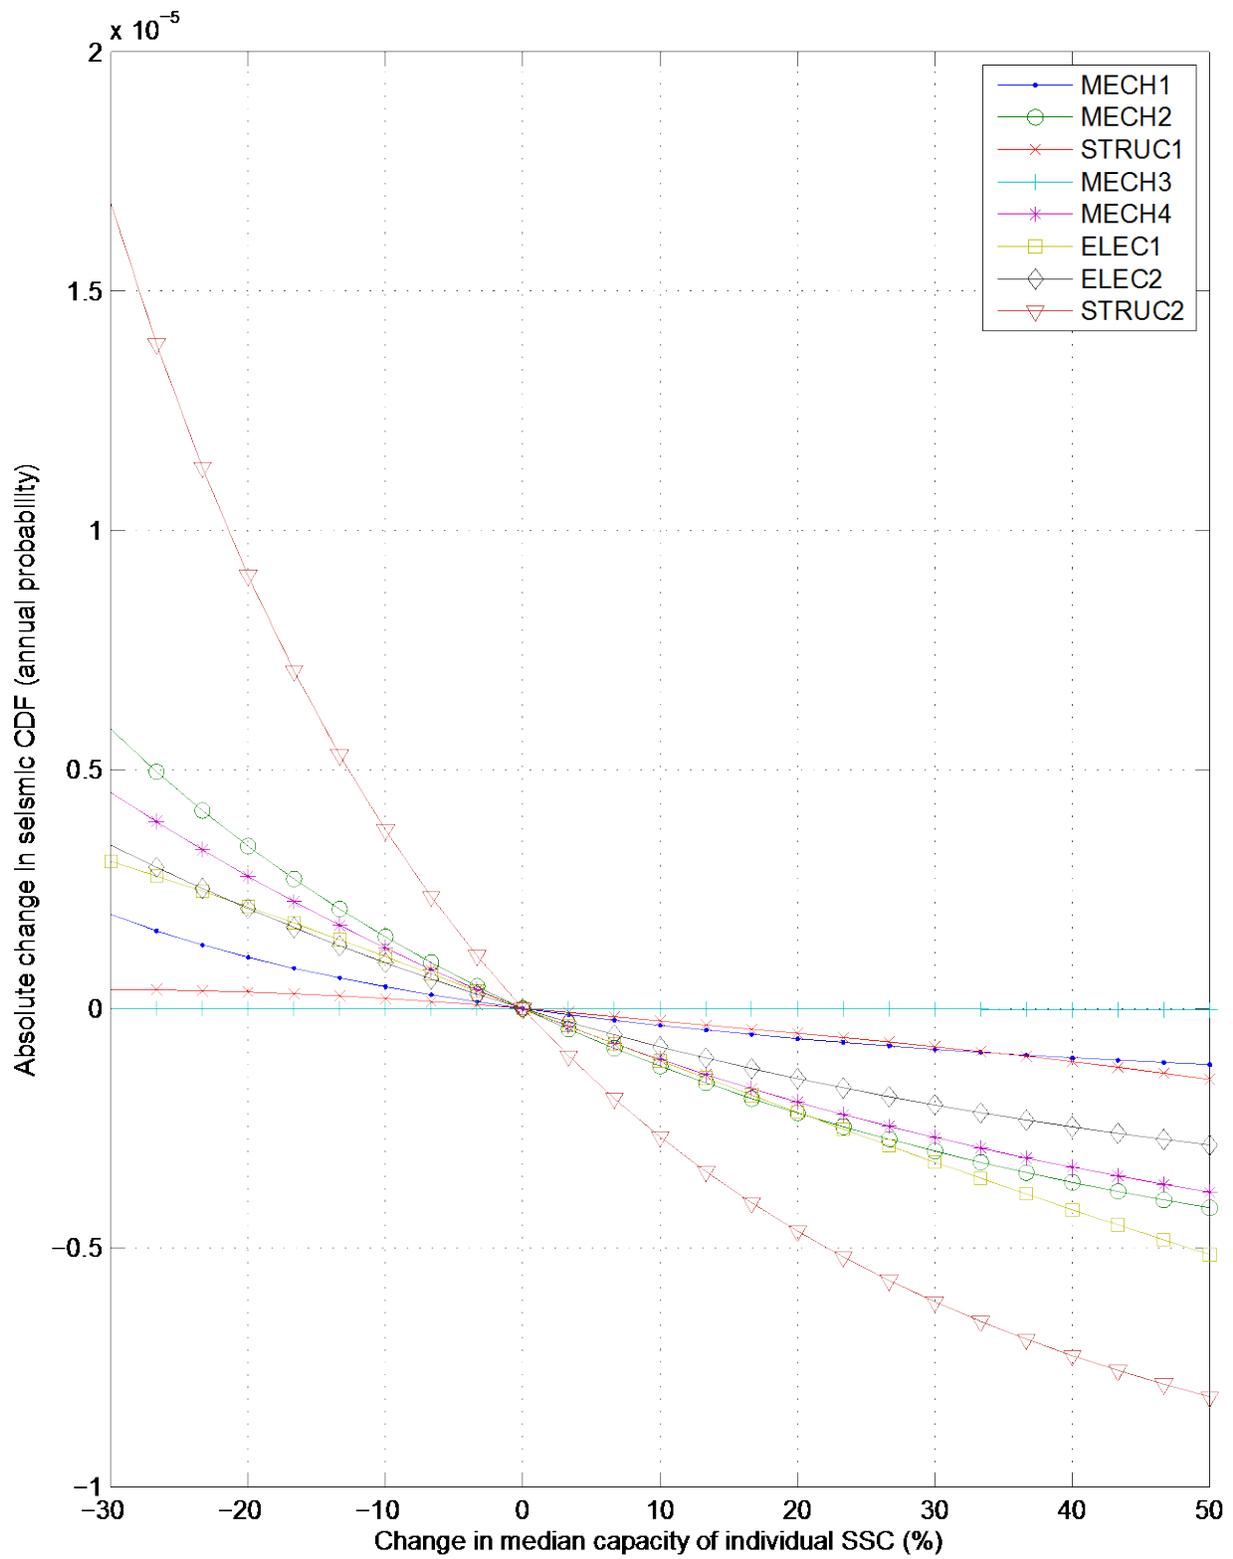


Figure 7-2 Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the baseline configuration of Plant A. The ordinate is a frequency in units of 10^{-5} per year.

Figure 7-3 shows the effect of strengthening ELEC1 on the relative contribution of each accident sequence to the plant's seismic CDF. Overall, a 50 percent increase in the median capacity of ELEC1 results in 17 percent decrease in the plant's seismic CDF (this effect can also be seen in Figure 7-2). However, as shown in Figure 7-3, strengthening ELEC1 does not lead to a more "balanced" plant seismic risk profile. In fact, the relative contribution of the most dominant accident sequence, SEQ1, actually grows from approximately 40 percent to 50 percent as the median capacity of ELEC1 is increased by 50 percent, although it should be emphasized that the plant's seismic CDF decreases by 17 percent.

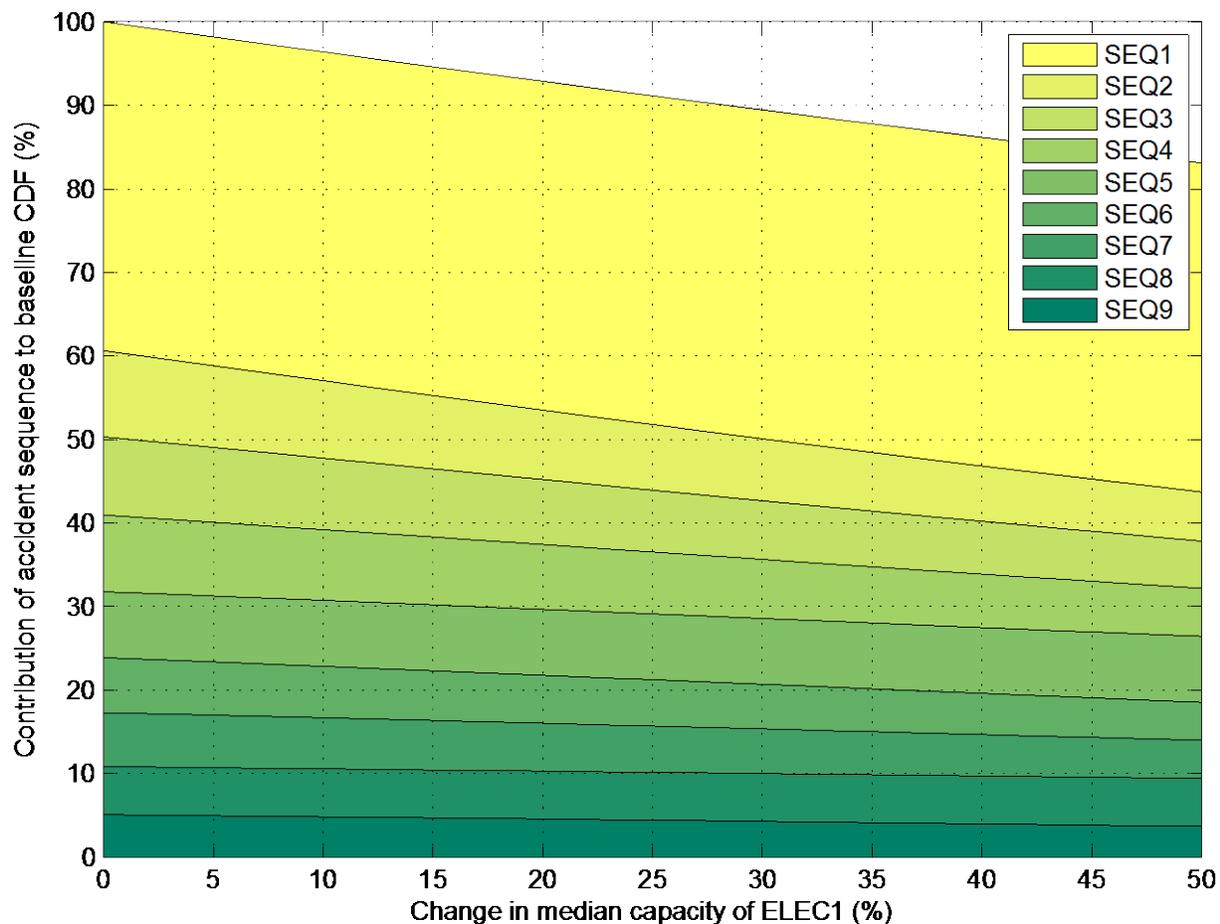


Figure 7-3 Effect of strengthening ELEC1 on the seismic risk profile of baseline configuration of Plant A

Similarly, Figure 7-4 shows the effect of strengthening STRUC2 on the relative contribution of each accident sequence to the plant's seismic CDF. Overall, a 50 percent increase in the median capacity of STRUC2 results in a 27 percent decrease in the seismic CDF (this effect can also be seen in Figure 7-2). Unlike strengthening ELEC1, strengthening STRUC2 helps produce a more "balanced" plant seismic risk profile, because the relative contribution of the most dominant accident sequence (SEQ1) decreases from approximately 40 percent to 18 percent as the median capacity of STRUC2 is increased by 50 percent. So not only does strengthening STRUC2 reduce overall seismic risk, it also produces a more "balanced" plant design in which no single accident sequence or SSC dominates the seismic risk profile.

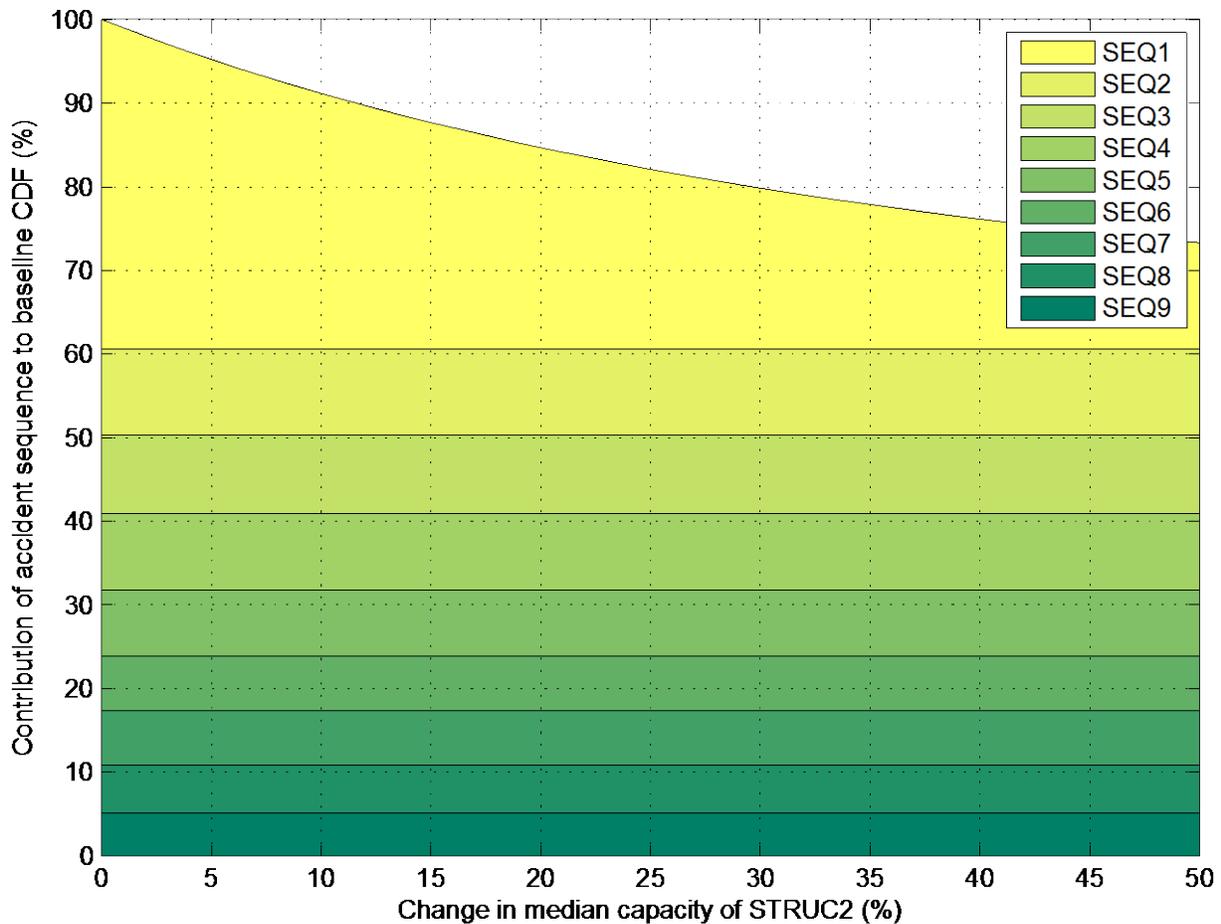


Figure 7-4 Effect of strengthening STRUC2 on the seismic risk profile of baseline configuration of Plant A

7.2.2 Plant B

Table 7-3 lists and describes the SSCs included in the analysis of Plant B, which is a plant with a higher seismic hazard located in the western US. Similar to Plant A, SSCs are categorized into three general groups: structural, mechanical, and electrical. Table 7-3 lists basic fragility parameters for each SSC, including median capacity, A_m , and uncertainty parameters, β_r and β_u . Table 7-4 lists and describes the dominant seismic accident sequences at the plant. These sequences are combinations of failures of the SSCs listed in Table 7-3. Together, Table 7-3 and Table 7-4 describe the baseline configuration of Plant B. This baseline configuration represents the NPP as it was originally analyzed, which is assumed to satisfy the provisions of the regulatory Framework in effect at that time. In total, there are 33 SSCs and 7 seismic accident sequences in the analysis of Plant B.

Table 7-3 and Table 7-4 also summarize the results of a simplified SPRA of the baseline configuration. In particular, Table 7-3 lists the high confidence of low probability of failure (HCLPF) value and the annual probability of failure, P_f , for each SSC. Note that only one SSC fails to satisfy the provisions of ASCE 43-05 (*i.e.*, an annual probability of failure less than 1×10^{-5}). Similarly, Table 7-4 lists the HCLPF value, annual probability of occurrence, and contribution to total seismic core-damage frequency (CDF) for each accident sequence. The seismic CDF for the baseline configuration of Plant B is 2.39×10^{-5} per year. This number will be

referred to as the baseline seismic CDF for Plant B. Note that two seismic accident sequences, SEQ1 and SEQ2, contribute disproportionately to the plant's seismic risk profile: together, SEQ1 and SEQ2 account for approximately 75 percent of the seismic CDF.

Table 7-3 Properties of SSCs in the baseline configuration of Plant B

SSC	A_m (g)	β_r	β_u	HCLPF (g)	P_f
STRUC1	6.91	0.20	0.31	2.98	2.95x10 ⁻⁷
STRUC2	8.55	0.28	0.31	3.23	1.69x10 ⁻⁷
STRUC3	5.79	0.21	0.26	2.67	5.87x10 ⁻⁷
STRUC4	4.87	0.26	0.33	1.84	8.28x10 ⁻⁶
STRUC5	9.23	0.18	0.21	4.85	5.40x10 ⁻¹⁰
STRUC6	8.71	0.25	0.33	3.35	1.33x10 ⁻⁷
MECH1	6.96	0.31	0.29	2.59	8.96x10 ⁻⁷
MECH2	7.62	0.30	0.42	2.32	2.51x10 ⁻⁶
MECH3	8.09	0.24	0.27	3.49	6.13x10 ⁻⁸
MECH4	8.53	0.29	0.21	3.74	3.39x10 ⁻⁸
MECH5	6.31	0.27	0.28	2.55	9.10x10 ⁻⁷
STRUC7	7.22	0.33	0.22	2.91	3.87x10 ⁻⁷
MECH6	8.33	0.27	0.23	3.65	3.84x10 ⁻⁸
MECH7	7.79	0.26	0.20	3.65	3.22x10 ⁻⁸
MECH8	8.78	0.29	0.24	3.66	4.51x10 ⁻⁸
STRUC8	7.40	0.29	0.35	2.57	1.04x10 ⁻⁶
ELEC1	4.55	0.30	0.13	2.24	3.34x10 ⁻⁶
MECH9	8.10	0.31	0.33	2.82	5.46x10 ⁻⁷
ELEC2	7.44	0.31	0.25	2.95	3.19x10 ⁻⁷
ELEC3	10.83	0.31	0.38	3.47	1.62x10 ⁻⁷
ELEC4	10.76	0.34	0.36	3.39	1.89x10 ⁻⁷
ELEC5	6.04	0.30	0.18	2.74	5.78x10 ⁻⁷
ELEC6	9.93	0.34	0.40	2.93	5.97x10 ⁻⁷
ELEC7	6.67	0.35	0.28	2.36	1.83x10 ⁻⁶
ELEC8	6.82	0.31	0.24	2.75	5.43x10 ⁻⁷
ELEC9	5.34	0.28	0.20	2.42	1.36x10 ⁻⁶
ELEC10	7.77	0.31	0.27	2.98	3.04x10 ⁻⁷
ELEC11	7.60	0.27	0.25	3.22	1.29x10 ⁻⁷

Table 7-3 Properties of SSCs in the baseline configuration of Plant B (Cont'd)

SSC	A_m (g)	β_r	β_u	HCLPF (g)	P_f
ELEC12	10.78	0.39	0.28	3.57	1.32x10 ⁻⁷
ELEC13	8.93	0.27	0.20	4.11	1.02x10 ⁻⁸
MECH10	7.09	0.28	0.32	2.63	7.98x10 ⁻⁷
ELEC14	1.69	0.24	0.20	0.82	6.91x10 ⁻⁴
STRUC9	11.22	0.39	0.40	3.05	5.39x10 ⁻⁷

Table 7-4 Properties of dominant seismic accident sequences in the baseline configuration of Plant B

Sequence	HCLPF (g)	P_f	% of seismic CDF
SEQ1	1.80	1.03x10 ⁻⁵	43.18%
SEQ2	1.98	7.60x10 ⁻⁶	31.83%
SEQ3	2.18	2.99x10 ⁻⁶	12.54%
SEQ4	2.42	1.39x10 ⁻⁶	5.82%
SEQ5	2.66	6.85x10 ⁻⁷	2.87%
SEQ6	N/A	4.65x10 ⁻⁷	1.95%
SEQ7	2.81	4.33x10 ⁻⁷	1.81%
CDF	--	2.39x10 ⁻⁵	100.00%

Figure 7-5 plots the sensitivity of the seismic CDF to changes in the seismic capacities of individual SSCs. More specifically, it plots the absolute change in the plant's seismic CDF produced by changing the median capacity of each individual SSC, one at a time. Note that not all 33 SSCs are represented in the figure. Instead, the 8 SSCs with the most impact on the plant's seismic CDF are plotted. Again, several interesting observations emerge from Figure 7-5. First, the most fragile SSCs do not necessarily have the most impact on the plant's seismic CDF, implying that simply strengthening the weakest components in the plant may not be the most effective way to improve overall seismic safety. Second, the plant's seismic CDF is disproportionately sensitive to changes in the strength of STRUC4. A 30 percent decrease in median capacity increases the seismic CDF by a factor of about three, meaning that an error in the design, analysis, operation, or maintenance of STRUC4 that subsequently reduces its capacity can significantly impact the safety of the plant. This heightened sensitivity to a single SSC is inconsistent with the NRC's defense-in-depth philosophy.

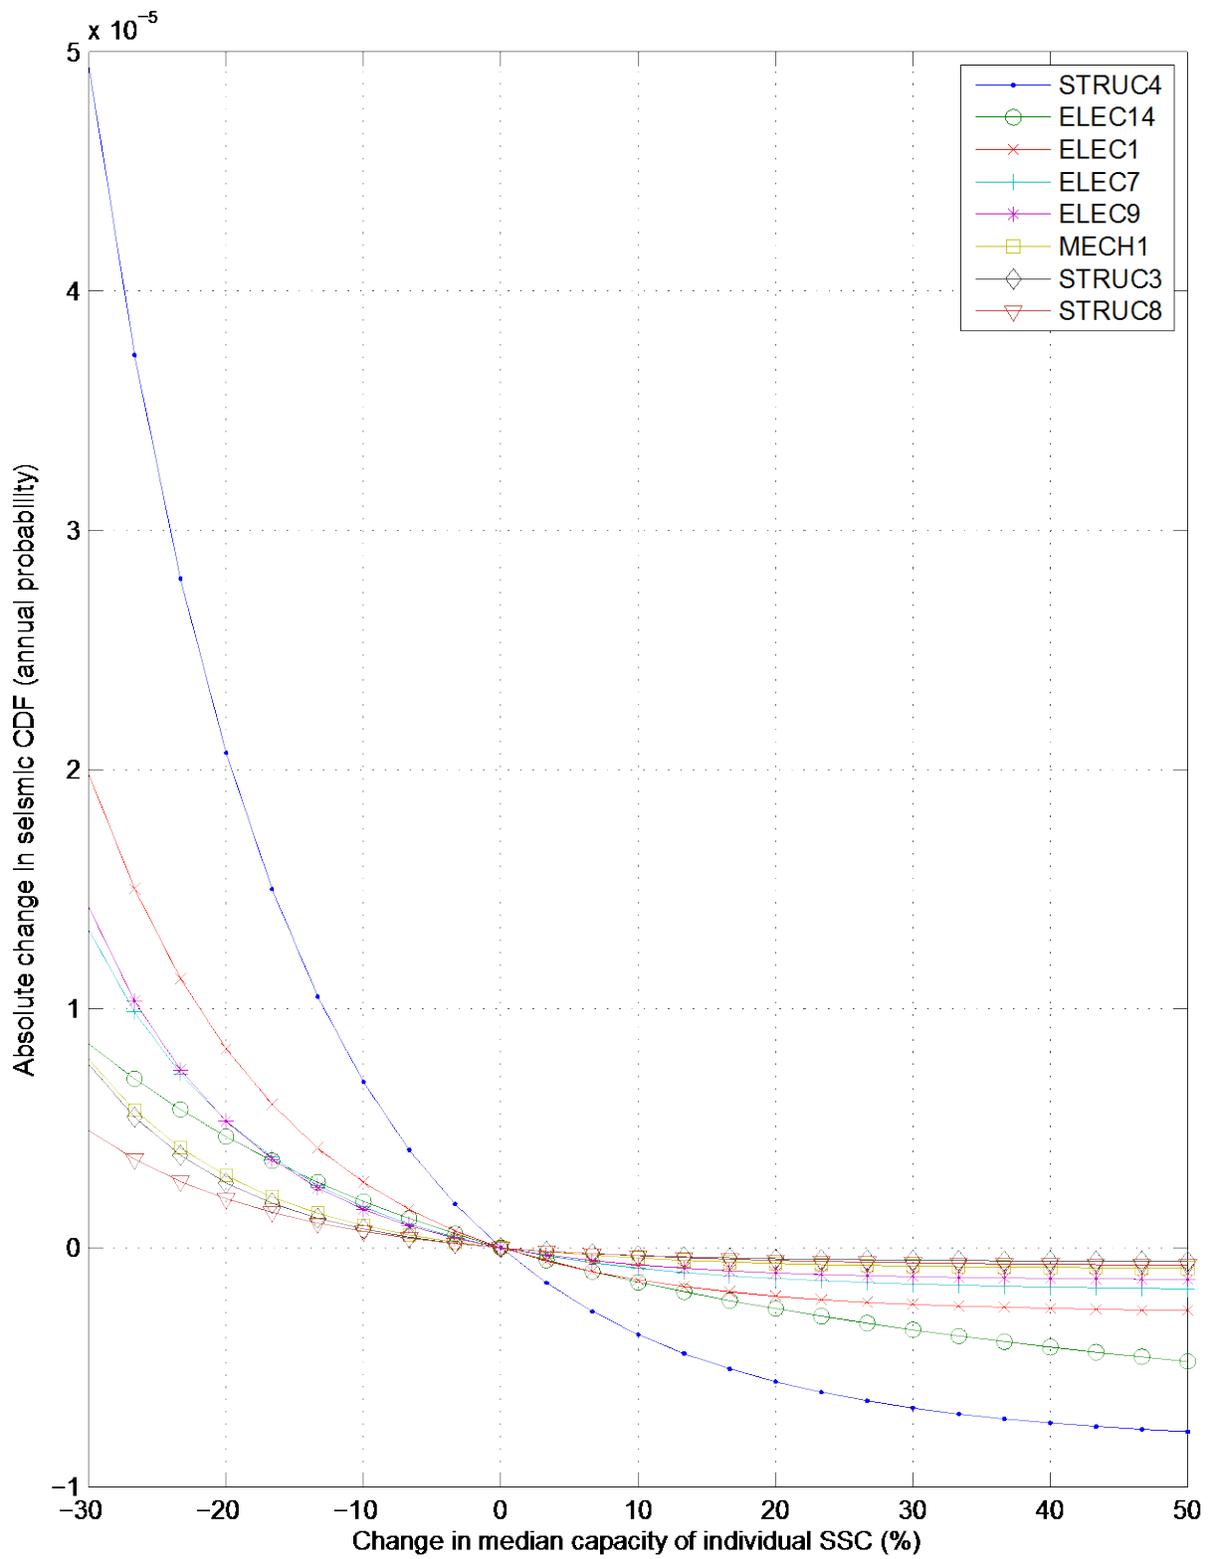


Figure 7-5 Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the baseline configuration of Plant B. The ordinate is a frequency in units of 10^{-5} per year.

Figure 7-6 shows the effect of strengthening ELEC14 on the relative contribution of each accident sequence to the plant's seismic CDF. Overall, a 50 percent increase in the median capacity of ELEC14 results in a 20 percent decrease in seismic CDF (this effect is also captured in Figure 7-5). However, as shown in Figure 7-6, strengthening ELEC14 does not lead to a more "balanced" plant seismic risk profile. In fact, the relative contribution of the most dominant seismic accident sequence, SEQ1, actually grows from approximately 43 percent to 53 percent as the median seismic capacity of ELEC14 is increased by 50 percent, though it is important to emphasize that the plant's seismic CDF decreases by 20 percent.

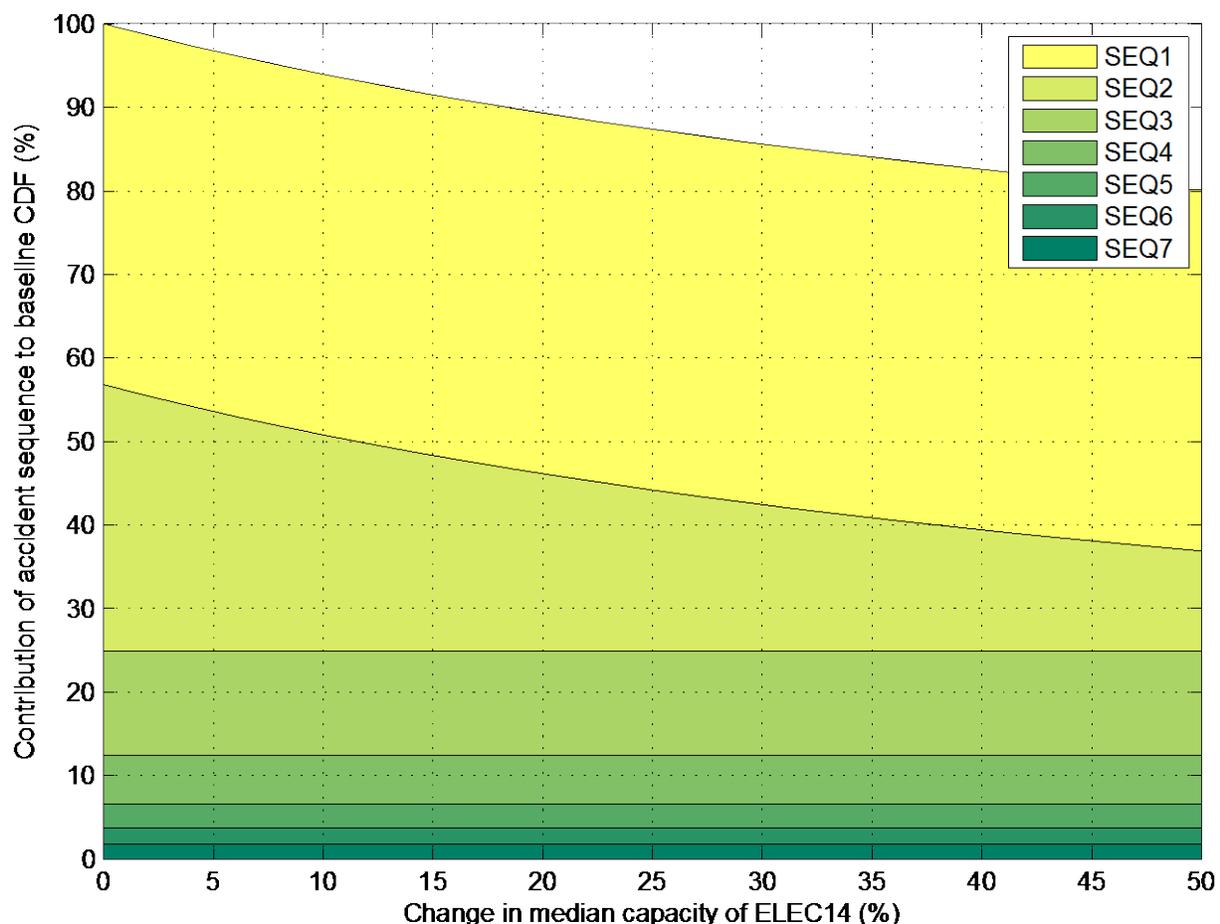


Figure 7-6 Effect of strengthening ELEC14 on the seismic risk profile of baseline configuration of Plant B

Similarly, Figure 7-7 shows the effect of strengthening STRUC4 on the relative contribution of each accident sequence to the plant's seismic CDF. Overall, a 50 percent increase in the median seismic capacity of STRUC4 results in a 32 percent decrease in seismic CDF (this effect can also be seen in Figure 7-5). Unlike strengthening ELEC14, strengthening STRUC4 helps produce a more "balanced" plant seismic risk profile, because the relative contribution of the most dominant seismic accident sequence, SEQ1, decreases from approximately 43 percent to 16 percent as the median seismic capacity of STRUC4 is increased by 50 percent. However, as a result, the relative contribution of the second most dominant seismic accident sequence, SEQ2, increases from approximately 32 percent to 47 percent (although, again, the plant's seismic CDF decreases by 32 percent). Consequently, in Plant B, the seismic capacities

of several SSCs need to be adjusted in order to produce a more “balanced” seismic risk profile. Section 7.3.2 details the specific adjustments required to accomplish this.

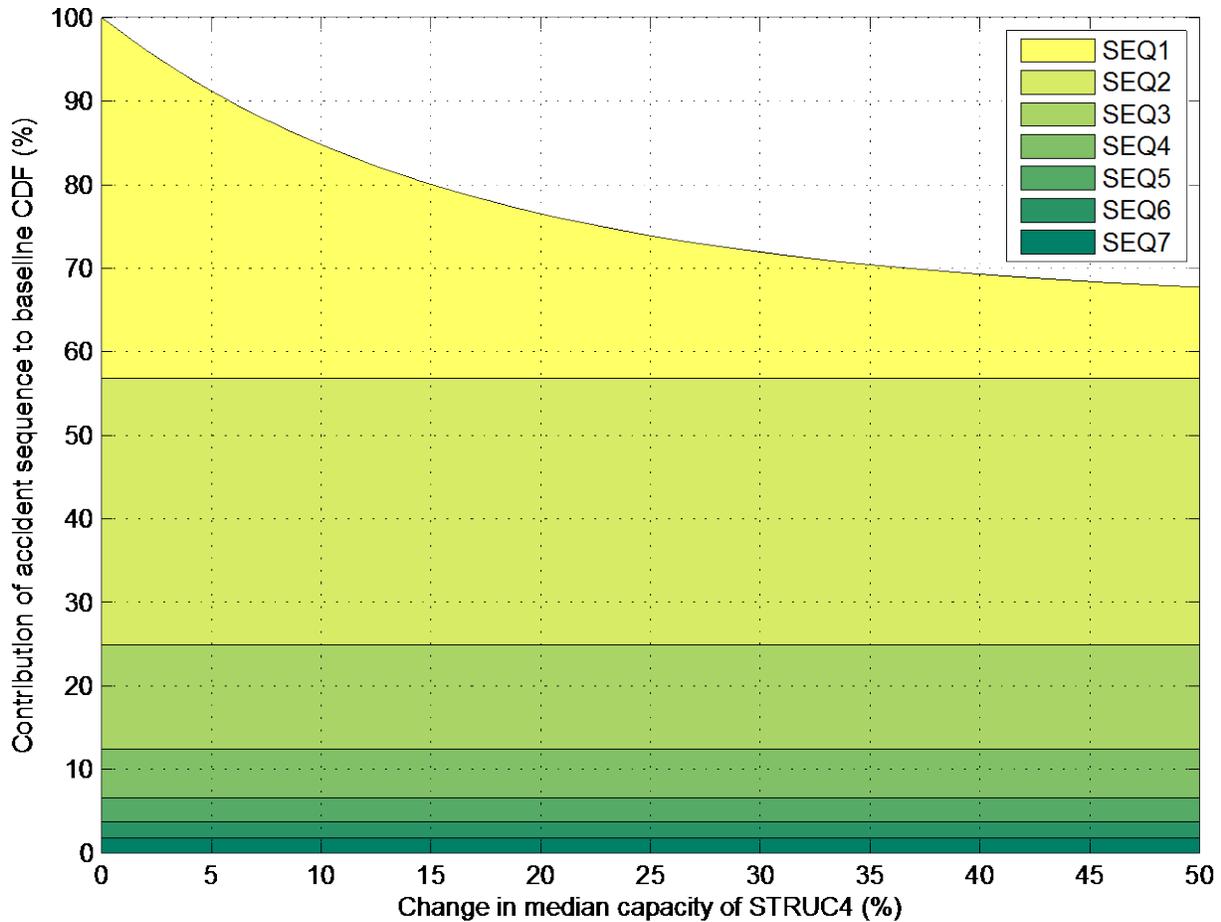


Figure 7-7 Effect of strengthening STRUC4 on the seismic risk profile of baseline configuration of Plant B

7.2.3 Summary

The rationale for a defense-in-depth approach to plant safety is that the overall acceptability of the plant’s safety should not rest on the strength of a singleton, nor rely on a single line of defense, nor be dependent on the correctness of the analysis of a single item. As both case studies have illustrated, NPPs designed in accordance with the provisions and requirements of the current seismic regulatory Framework do not necessarily have adequate defense in depth in the seismic area, as the seismic risk profiles of both plants are dominated by a single SSC and/or a single seismic accident sequence (see Figure 7-2 and Figure 7-5). If, for some reason, this SSC or this sequence is not as strong as anticipated, the bottom-line findings of the entire SPRA are incorrect. This type of situation is exactly why the overall philosophy of defense in depth has been developed and deployed – to avoid the situation where a single error related to a single item (or a single “event” like a common-cause failure of more than one item) means that the NPP is more vulnerable than we think it is.

7.3 Analysis

In this Section we will analyze two different strategies to address the shortcomings highlighted by the two case studies introduced in Section 7.2. These two strategies, which are briefly described in the following paragraphs, will produce two new configurations for each case study: the “ASCE 43-05” configuration and the “enhanced balance” configuration.

The first strategy is one advanced by ASCE 43-05, a consensus industry standard whose intent is to “ensure that nuclear facilities can withstand the effects of earthquake ground shaking with desired performance” (ASCE 2005). ASCE 43-05 achieves this by requiring that individual SSCs attain minimum levels of performance, with the exact level of performance depending on “the severity of adverse radiological and toxicological effects of the hazards that may result from the seismic failure of the SSC on workers, the public, and the environment” (ASCE 2005). For safety-related SSCs in NPPs, this translates into a target performance goal of 1×10^{-5} (or less) annual probability of seismic failure. In Section 7.3.1, the baseline configuration for each of the two case studies presented in Section 7.2 is adjusted in order to satisfy the provisions of ASCE 43-05. In particular, each SSC is redesigned to ensure its annual probability of failure in earthquakes does not exceed 1×10^{-5} . The SPRA of the plant is modified to examine the effect of these changes on the seismic risk profile.

The second strategy is the one described in Section 5 of this report. It aims to improve a plant’s defense in depth by limiting the contribution of single SSCs and accident sequences to the plant’s seismic risk profile. This can be accomplished several ways; for example, by limiting the contribution of a single accident sequence to a specific percentage (e.g., 25 or 50 percent) of the seismic CDF, or by limiting the probability of each accident sequence to an absolute threshold (e.g., 5×10^{-6}), or a combination of the two. In Section 7.3.2, the baseline configuration for each of the two case studies presented in Section 7.2 is adjusted in order to satisfy an arbitrary defense-in-depth provision selected provisionally here for the sake of illustration. In particular, SSCs are redesigned so that no single accident sequence comprises more than 25 percent of the seismic CDF. A revised SPRA analysis of the plant is then performed in order to examine the effect of these changes on the seismic risk profile.

7.3.1 ASCE 43-05 Configuration

As shown in Table 7-1 and Table 7-3, the baseline configuration for both Plant A and Plant B includes a small number of SSCs that have an annual probability of failure in earthquakes greater than 1×10^{-5} . In order to bring the baseline configuration into compliance with ASCE 43-05, the median seismic capacity of each noncompliant SSC from the baseline configuration is increased (strengthened) until its annual probability of failure in earthquakes is less than 1×10^{-5} . The resulting configuration is referred to as the “ASCE 43-05” configuration. This configuration represents the following hypothetical situation. An engineer uses current industry codes and standards to design and evaluate the SSCs in an NPP. This initial design is equivalent to the baseline configuration described in Section 7.2. After performing a SPRA, the engineer discovers that several SSCs do not satisfy the provisions of ASCE 43-05. Subsequently, the engineer redesigns and strengthens those noncompliant SSCs, ultimately producing the “ASCE 43-05” configuration. The following subsections describe in detail this configuration for both Plant A and Plant B.

Plant A

As shown in Table 7-1 for Plant A, four SSCs (STRUCT1, MECH3, ELEC1, and STRUCT2) require adjustment in order to satisfy the provisions of ASCE 43-05. In order to decrease the annual probability of failure for these SSCs, the median capacity of each SSC is increased until its annual probability of failure drops below 1×10^{-5} . Table 7-5 summarizes these adjustments. Note that no other properties were changed (e.g., β_u and β_r), and also that SSCs that already satisfy the provisions of ASCE 43-05 in the baseline configuration were not altered. Note that the median capacity of one component, MECH3, needed to be increased by a factor of four (a 300 percent increase) in order to comply with the requirements of ASCE 43-05. Although strengthening MECH3 reduces its annual probability of failure to acceptable levels (as defined by ASCE 43-05), this strengthening has almost no impact on the plant’s seismic CDF (see Figure 7-2).

Table 7-5 Changes in seismic capacities of SSCs in the “ASCE 43-05” configuration of Plant A (see Table 7-1 for SSCs not listed)

SSC	baseline A_m (g)	updated A_m (g)	% change in A_m	baseline P_f	updated P_f	% change in P_f
STRUC1	0.32	0.45	40.6%	2.62×10^{-5}	9.91×10^{-6}	-62.2%
MECH3	0.13	0.53	307.7%	1.75×10^{-4}	9.58×10^{-6}	-94.5%
ELEC1	0.30	0.57	90.0%	4.62×10^{-5}	9.99×10^{-6}	-78.4%
STRUC2	0.53	0.57	7.5%	1.20×10^{-5}	9.88×10^{-6}	-17.5%

Table 7-6 shows how the changes in Table 7-5 impact each of the nine accident sequences and also the plant’s seismic CDF. Overall, as a result of these changes, the plant’s seismic CDF decreases 35 percent from 3.04×10^{-5} to 1.98×10^{-5} . Note, however, that a single seismic accident sequence, SEQ1, continues to dominate the risk profile, with its contribution growing from 39 percent (see Table 7-2) to 50 percent, although it is important to emphasize that the plant’s seismic CDF has decreased by 35 percent.

Table 7-6 Changes in properties of dominant seismic accident sequences in the “ASCE 43-05” configuration of Plant A

Sequence	updated HCLPF (g)	updated P_f	% change in P_f	% of updated seismic CDF	% of baseline seismic CDF
SEQ1	0.19	9.88×10^{-6}	-17.55%	49.98%	32.46%
SEQ2	0.52	9.43×10^{-7}	-69.97%	4.77%	3.10%
SEQ3	0.52	1.03×10^{-6}	-63.63%	5.24%	3.40%
SEQ4	0.52	1.07×10^{-6}	-61.62%	5.43%	3.52%
SEQ5	0.34	2.40×10^{-6}	0.00%	12.13%	7.88%
SEQ6	0.54	9.53×10^{-7}	-52.45%	4.82%	3.13%
SEQ7	0.56	9.47×10^{-7}	-51.99%	4.79%	3.11%
SEQ8	0.42	1.74×10^{-6}	0.00%	8.83%	5.73%
SEQ9	0.55	7.93×10^{-7}	-48.85%	4.01%	2.61%
CDF	--	1.98×10^{-5}	-35.05%	100.00%	64.95%

Together, Table 7-5 and Table 7-6 summarize the “ASCE 43-05” configuration of Plant A. Figure 7-8 shows the sensitivity of this configuration to changes in the seismic capacity of individual SSCs. Similar to Figure 7-2, which shows the sensitivity of the baseline configuration, Figure 7-8 demonstrates that the “ASCE 43-05” configuration of Plant A is also highly sensitive to changes in the seismic capacity of a single SSC, STRUC2 (note that the scale of the left-hand vertical axis on both Figure 7-2 and Figure 7-8 is the same). A 30 percent decrease in the median seismic capacity of STRUC2 results in a 75 percent increase in the plant’s seismic CDF, meaning that an error in the design, analysis, operation, or maintenance of STRUC2 that subsequently reduces its seismic capacity can significantly impact the safety of the plant. In summary, probably because most of the important SSCs already meet ASCE 43-05 in the baseline configuration, the provisions of ASCE 43-05 do little to address the unbalanced seismic risk profile of Plant A.

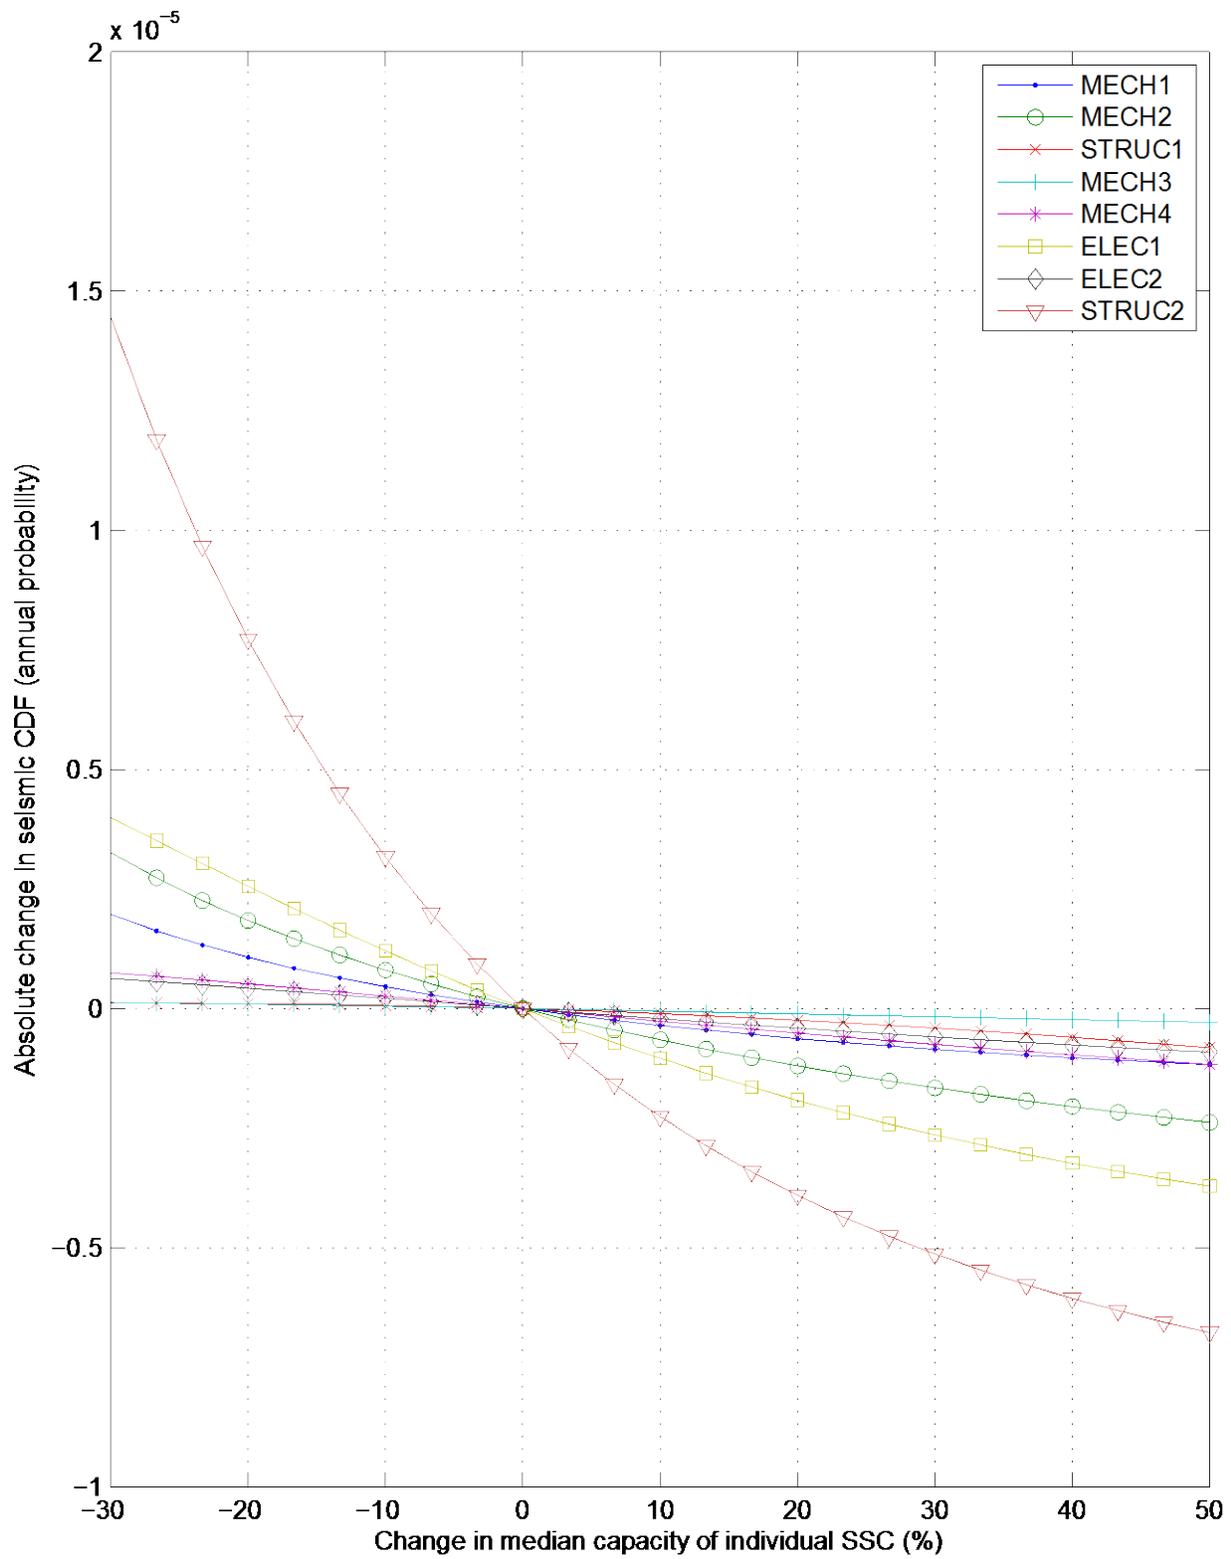


Figure 7-8 Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the “ASCE 43-05” configuration of Plant A. The ordinate is a frequency in units of 10^{-5} per year.

Plant B

As shown in Table 7-3, for Plant B only one SSC (ELEC14) requires adjustment in order to satisfy the provisions of ASCE 43-05. Table 7-7 summarizes this adjustment. Note that no other properties were changed (e.g., β_u and β_r), and also that SSCs that already satisfied the provisions of ASCE 43-05 in the baseline configuration were not altered.

Table 7-7 Changes in seismic capacities of SSCs in the “ASCE 43-05” configuration of Plant B (see Table 7-3 for SSCs not listed)

SSC	baseline A_m (g)	updated A_m (g)	% change in A_m	baseline P_f	updated P_f	% change in P_f
ELEC14	1.69	3.78	123.7%	6.91×10^{-4}	9.92×10^{-6}	-98.6%

Table 7-8 shows how the change in Table 7-7 impacts each of the seven seismic accident sequences and also the plant’s seismic CDF. Overall, as a result of the change (which is a strengthening of a single SSC), the plant’s seismic CDF decreases 29 percent from 2.39×10^{-5} to 1.69×10^{-5} . Recall that in the baseline configuration, two seismic accident sequences, SEQ1 and SEQ2, dominated the seismic risk profile of the plant (see Table 7-4). In the ASCE 43-05 configuration, the contribution of SEQ2 has decreased from 32 percent of the baseline seismic CDF to 4 percent of the updated CDF. Consequently, the contribution of SEQ1 has increased from 43 percent of the baseline seismic CDF to 61 percent of the updated CDF, meaning that a single seismic accident sequence now dominates the seismic risk profile of the plant, even though there has been no change in the overall CDF from SEQ1, although again it should be noted that the plant’s seismic CDF has decreased by 29 percent due to the decreased contribution from SEQ2.

Table 7-8 Changes in properties of dominant seismic accident sequences in the “ASCE 43-05” configuration of Plant B

Sequence	updated HCLPF (g)	updated P_f	% change in P_f	% of updated seismic CDF	% of baseline seismic CDF
SEQ1	1.80	1.03×10^{-5}	0.00%	60.97%	43.18%
SEQ2	2.64	6.34×10^{-7}	-91.66%	3.75%	2.65%
SEQ3	2.18	2.99×10^{-6}	0.00%	17.70%	12.54%
SEQ4	2.42	1.39×10^{-6}	0.00%	8.21%	5.82%
SEQ5	2.66	6.85×10^{-7}	0.00%	4.05%	2.87%
SEQ6	N/A	4.65×10^{-7}	0.00%	2.75%	1.95%
SEQ7	2.81	4.33×10^{-7}	0.00%	2.56%	1.81%
CDF	--	1.69×10^{-5}	-29.17%	100.00%	70.83%

Together, Table 7-7 and Table 7-8 summarize the ASCE 43-05 configuration of Plant B. Figure 7-9 shows the sensitivity of this configuration to changes in the seismic capacities of individual SSCs. Similar to Figure 7-5, which shows the sensitivity of the baseline configuration, Figure

7-9 shows that the “ASCE 43-05” configuration of Plant B is also highly sensitive to changes in the seismic capacity of a single SSC, STRUC4 (note that the scale of the left-hand vertical axis on both Figure 7-5 and Figure 7-9 is the same). A 30 percent decrease in the median seismic capacity of STRUC4 increases the plant’s seismic CDF by a factor of four, meaning that an error in the design, analysis, operation, or maintenance of STRUC4 that reduces its seismic capacity can significantly impact the seismic safety of the plant. In summary, the provisions of ASCE 43-05 do little to address the unbalanced seismic risk profile of Plant B, probably because in the baseline configuration all but one of the important SSCs already meet the performance-target provisions of ASCE 43-05.

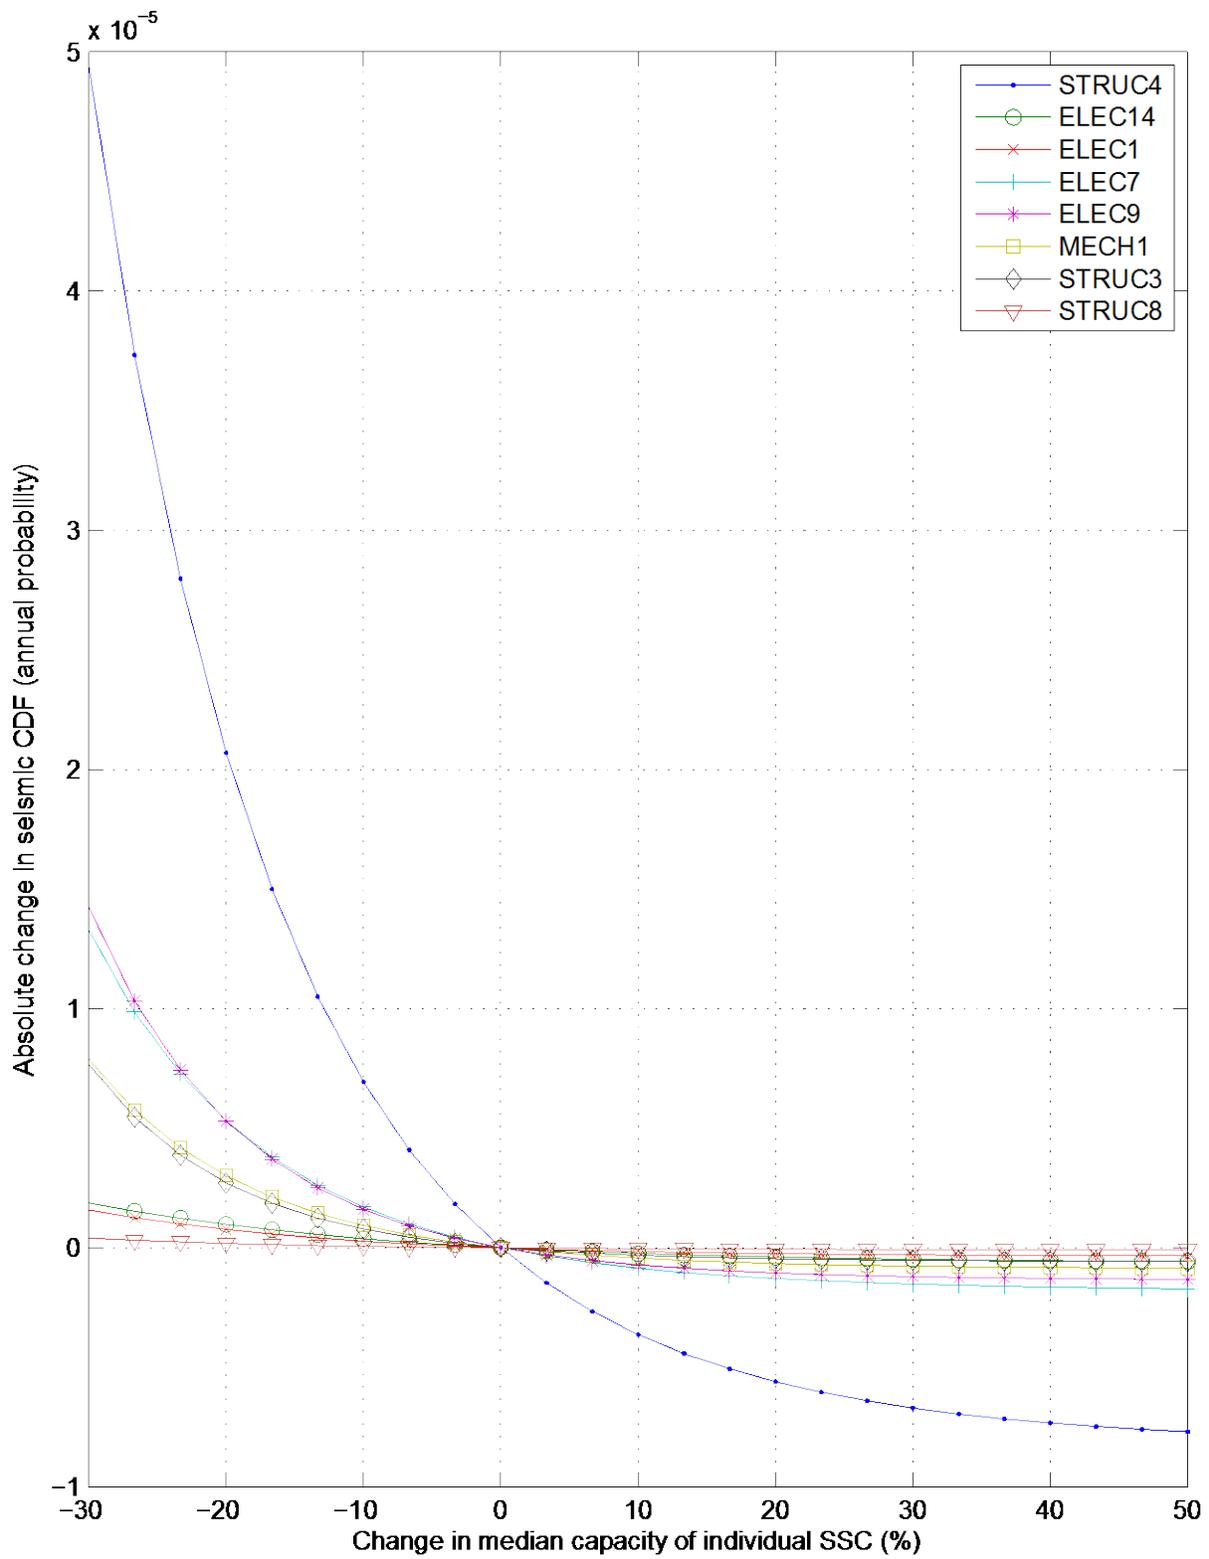


Figure 7-9 Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the “ASCE 43-05” configuration of Plant B. The ordinate is a frequency in units of 10^{-5} per year.

7.3.2 Enhanced Balance Configuration

As shown in Table 7-2 and Table 7-4, the seismic risk profile of the baseline configuration for both Plant A and Plant B is dominated by a small number of accident sequences. In Plant A, a single seismic accident sequence comprises 40 percent of the seismic CDF, while in Plant B two accident sequences comprise 75 percent. Furthermore, as Figure 7-2 and Figure 7-5 capture, these dominant seismic accident sequences themselves are dominated by the contribution of a single SSC. The resulting plant seismic design is unbalanced in that an error in the design or analysis of the dominant SSC or SSCs, or a large variability in expected performance due to mis-characterized uncertainty, can potentially affect the overall safety of the plant.

In order to produce a more balanced seismic risk profile, the following risk-balance requirement will be enforced: *no single seismic accident sequence can account for more than 25 percent of the plant's seismic CDF*. Note that this rule is somewhat arbitrary in that the target percentage could be some other value, say 10 or 50 percent. Alternatively, the rule could establish an absolute limit on the annual probability of occurrence of an accident sequence (e.g., less than 5×10^{-6}), or be a combination of relative and absolute performance requirements. In order to bring the baseline configuration of each case study into compliance with the arbitrary rule chosen in this example (i.e., 25 percent), the median seismic capacity of each relevant SSC will be adjusted until the relative contribution of the corresponding seismic accident sequence or sequences drops below 25 percent of the seismic CDF. The resulting configuration is referred to as the “enhanced balance” configuration. This configuration represents the following hypothetical situation. An engineer uses current industry seismic codes and standards to design and evaluate the SSCs in an NPP. This initial design is equivalent to the baseline configuration described in Section 7.2. After performing a SPRA, the engineer discovers that the seismic risk profile of the plant does not satisfy the above risk-balance requirement. Subsequently, the engineer redesigns and strengthens SSCs until the required risk profile is obtained, ultimately producing the “enhanced balance” configuration. The following subsections describe this configuration in detail for both Plant A and Plant B.

Plant A

As shown in Table 7-2, there is only one seismic accident sequence (SEQ1) that comprises more than 25 percent of the plant's seismic CDF. This accident sequence, in turn, is dominated by a single SSC, STRUC2. Therefore, in order to reduce the contribution of SEQ1 to the seismic risk profile of the plant, the median seismic capacity of STRUC2 needs to be increased (strengthened) until SEQ1 comprises less than 25 percent of the plant's seismic CDF. Table 7-9 summarizes the change to STRUC2 required to bring about the desired goal. Note that no other properties of STRUC2 were changed (e.g., β_u and β_r), and also that SSCs not listed in Table 7-9 were not altered from the baseline configuration described in Table 7-1.

Table 7-9 Changes in seismic capacities of SSCs in the “enhanced balance” configuration of Plant A (see Table 7-1 for SSCs not listed)

SSC	baseline A_m (g)	updated A_m (g)	% change in A_m	baseline P_f	updated P_f	% change in P_f
STRUC2	0.53	0.68	28.3%	1.20×10^{-5}	6.07×10^{-6}	-49.3%

Table 7-10 shows how the change in Table 7-9 impacts each of the nine accident sequences and also the plant's seismic CDF. Overall, as a result of these changes, the plant's seismic CDF decreases 20 percent from 3.04×10^{-5} to 2.45×10^{-5} . At the same time, the contribution of SEQ1 drops from 39 percent of the baseline seismic CDF (see Table 7-2) to 25 percent of the updated seismic CDF, bringing it in compliance with the risk-balance requirement chosen for this particular example (*i.e.*, 25 percent limit on seismic CDF contribution).

Table 7-10 Changes in properties of dominant accident sequences in the “enhanced balance” configuration of Plant A

Sequence	updated HCLPF (g)	updated P_f	% change in P_f	% of updated seismic CDF	% of baseline seismic CDF
SEQ1	0.22	6.07×10^{-6}	-49.32%	24.76%	19.95%
SEQ2	0.34	3.14×10^{-6}	0.00%	12.81%	10.32%
SEQ3	0.36	2.85×10^{-6}	0.00%	11.61%	9.35%
SEQ4	0.38	2.79×10^{-6}	0.00%	11.39%	9.18%
SEQ5	0.34	2.40×10^{-6}	0.00%	9.78%	7.88%
SEQ6	0.38	2.00×10^{-6}	0.00%	8.18%	6.59%
SEQ7	0.40	1.97×10^{-6}	0.00%	8.04%	6.48%
SEQ8	0.42	1.74×10^{-6}	0.00%	7.11%	5.73%
SEQ9	0.45	1.55×10^{-6}	0.00%	6.32%	5.10%
CDF	--	2.45×10^{-5}	-19.42%	100.00%	80.58%

Together, Table 7-9 and Table 7-10 summarize the “enhanced balance” configuration of Plant A. Figure 7-10 shows the sensitivity of this configuration to changes in the seismic capacities of individual SSCs. In contrast with Figure 7-2, which shows the sensitivity of the baseline configuration, Figure 7-10 demonstrates that the “enhanced balance” configuration of Plant A is not as sensitive to changes in the properties of individual SSCs and, in particular, changes to STRUC2 (note that the scale of the left-hand vertical axis on both Figure 7-2 and Figure 7-10 is the same). A 30 percent decrease in the median seismic capacity of STRUC2 results in a 40 percent increase in seismic CDF, which is more in line with other SSCs in the plant. In summary, the proposed risk-balance requirement produces a more balanced plant seismic design in that no single SSC or seismic accident sequence dominates the seismic risk profile.

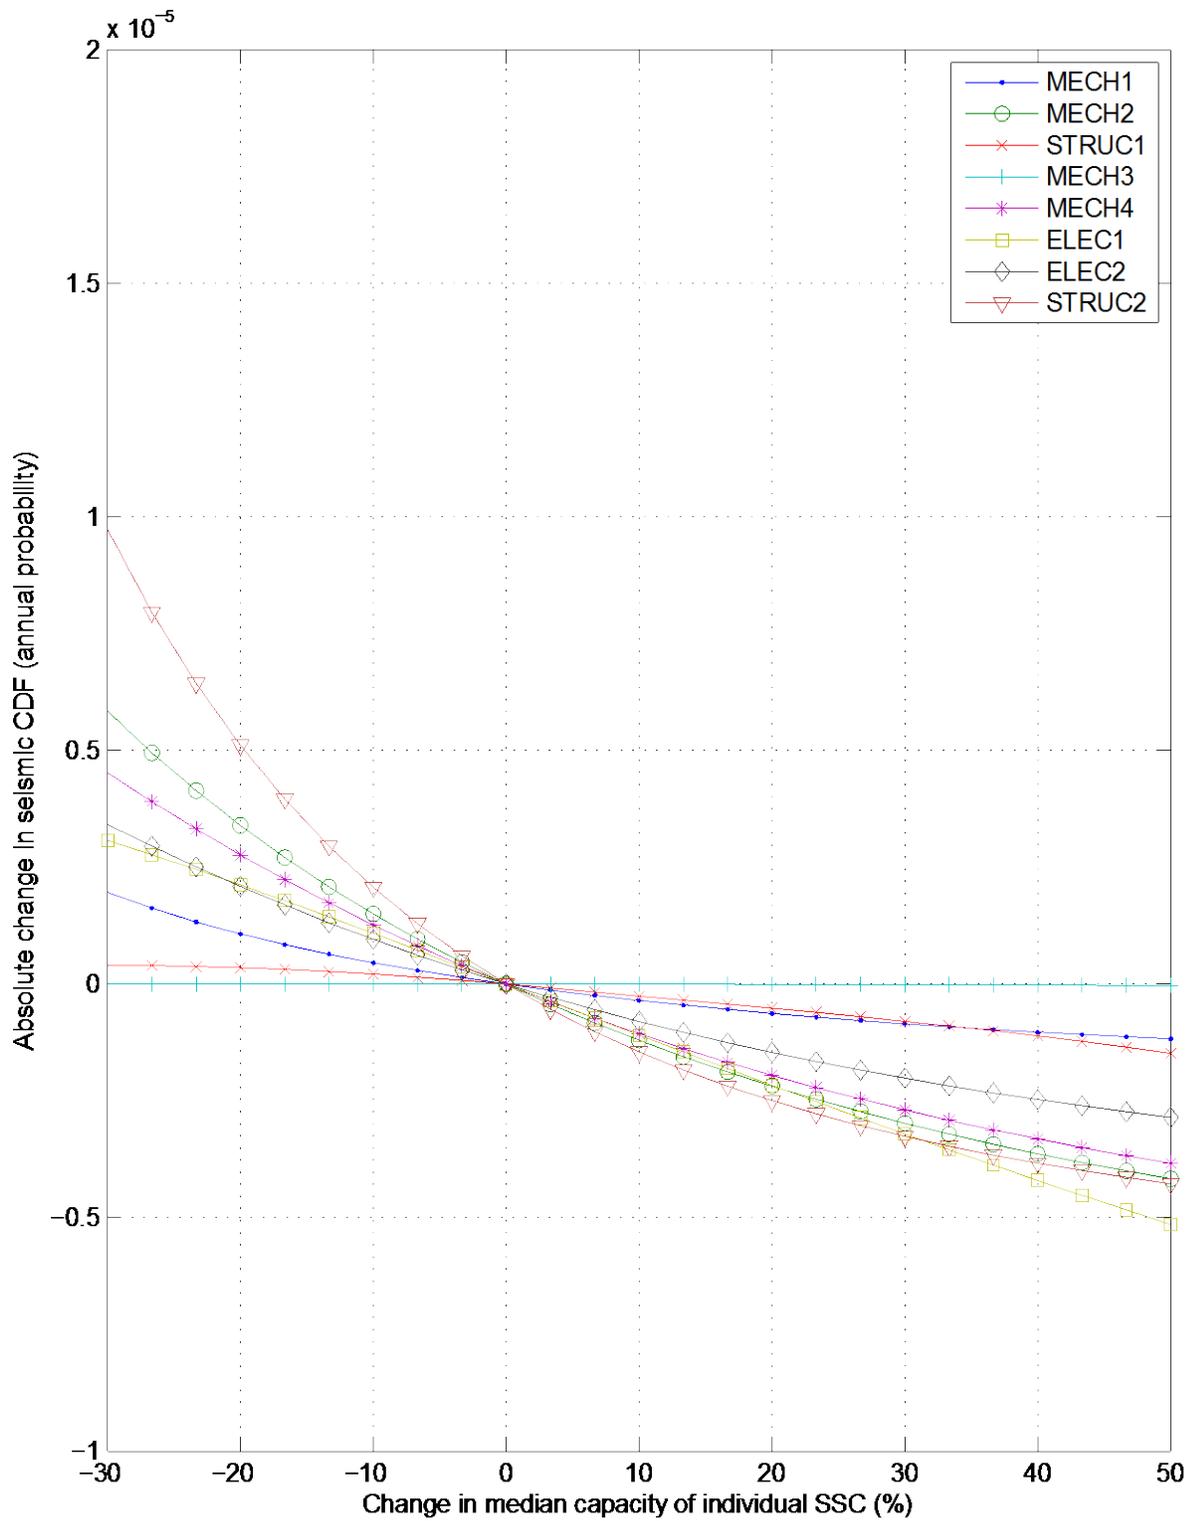


Figure 7-10 Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the “enhanced balance” configuration of Plant A. The ordinate is a frequency in units of 10^{-5} per year.

Plant B

As shown in Table 7-4, there are two seismic accident sequences (SEQ1 and SEQ2) that comprise more than 25 percent of the plant's seismic CDF: SEQ1 comprises 43 percent while SEQ2 comprises 32 percent. Both accident sequences, in turn, are dominated by a single SSC: for SEQ1 it is STRUC4 while for SEQ2 it is ELEC14. Therefore, in order to reduce the contribution of both accident sequences to the seismic risk profile of the plant, the median seismic capacities of STRUC4 and ELEC14 need to be increased (strengthened) until each sequence comprises less than 25 percent of the plant's seismic CDF. Table 7-11 summarizes the changes to STRUC4 and ELEC14 required to bring about the desired goal. Note that no other properties of STRUC4 or ELEC14 were changed (*e.g.*, β_u and β_r), and also that SSCs not listed in Table 7-11 were not altered from the baseline configuration described in Table 7-3.

Table 7-11 Changes in seismic capacities of SSCs in the “enhanced balance” configuration of Plant B (see Table 7-3 for SSCs not listed)

SSC	baseline A_m (g)	updated A_m (g)	% change in A_m	baseline P_f	updated P_f	% change in P_f
STRUC4	4.87	6.75	38.6%	8.28×10^{-6}	1.02×10^{-6}	-87.7%
ELEC14	1.69	2.45	45.0%	6.91×10^{-4}	1.24×10^{-4}	-82.1%

Table 7-12 shows how the changes in Table 7-11 impact each of the seven accident sequences and also the plant's seismic CDF. Overall, as a result of these changes, the plant's seismic CDF decreases 49 percent from 2.39×10^{-5} to 1.21×10^{-5} . At the same time, the contribution of SEQ1 drops from 43 percent of the baseline seismic CDF (see Table 7-2) to 25 percent of the updated seismic CDF, while the contribution of SEQ2 decreases from 32 percent of the baseline seismic CDF to 25 percent of the updated CDF. Note, however, that the contribution of SEQ3 increases from 12 percent to 25 percent, although its absolute CDF contribution is unchanged, which is consistent with the risk-balance requirement established for this example. To a certain extent this side effect is expected because the relative contribution of the targeted seismic accident sequences (*e.g.*, SEQ1 and SEQ2) has decreased. An iteration or two may be required before a configuration that satisfies the risk-balance requirement can be reached.

Table 7-12 Changes in properties of dominant seismic accident sequences in the “enhanced balance” configuration of Plant B

Sequence	updated HCLPF (g)	updated P_f	% change in P_f	% of updated seismic CDF	% of baseline seismic CDF
SEQ1	2.17	3.05×10^{-6}	-70.43%	25.10%	12.77%
SEQ2	2.21	3.13×10^{-6}	-58.76%	25.79%	13.13%
SEQ3	2.18	2.99×10^{-6}	0.00%	24.64%	12.54%
SEQ4	2.42	1.39×10^{-6}	0.00%	11.43%	5.82%
SEQ5	2.66	6.85×10^{-7}	0.00%	5.64%	2.87%
SEQ6	N/A	4.65×10^{-7}	0.00%	3.83%	1.95%
SEQ7	2.81	4.33×10^{-7}	0.00%	3.57%	1.81%
CDF	--	1.21×10^{-5}	-49.12%	100.00%	50.88%

Together, Table 7-11 and Table 7-12 summarize the “enhanced balance” configuration of Plant B. Figure 7-11 shows the sensitivity of this configuration to changes in the seismic capacities of individual SSCs. In contrast with Figure 7-5, which shows the sensitivity of the baseline configuration, Figure 7-11 demonstrates that the “enhanced balance” configuration of Plant B is not as sensitive to changes in the capacities of individual SSCs and, in particular, STRUC4 (note that the scale of the left-hand vertical axis on both Figure 7-5 and Figure 7-11 is the same). A 30 percent decrease in the median seismic capacity of STRUC4 results in an 80 percent increase in the plant’s seismic CDF, which is approximately similar to the increase in seismic CDF produced by other SSCs in the plant (*i.e.*, STRUC4 is no longer an outlier). In summary, the proposed risk-balance requirement produces a more balanced plant design in which no single SSC or seismic accident sequence dominates the seismic risk profile.

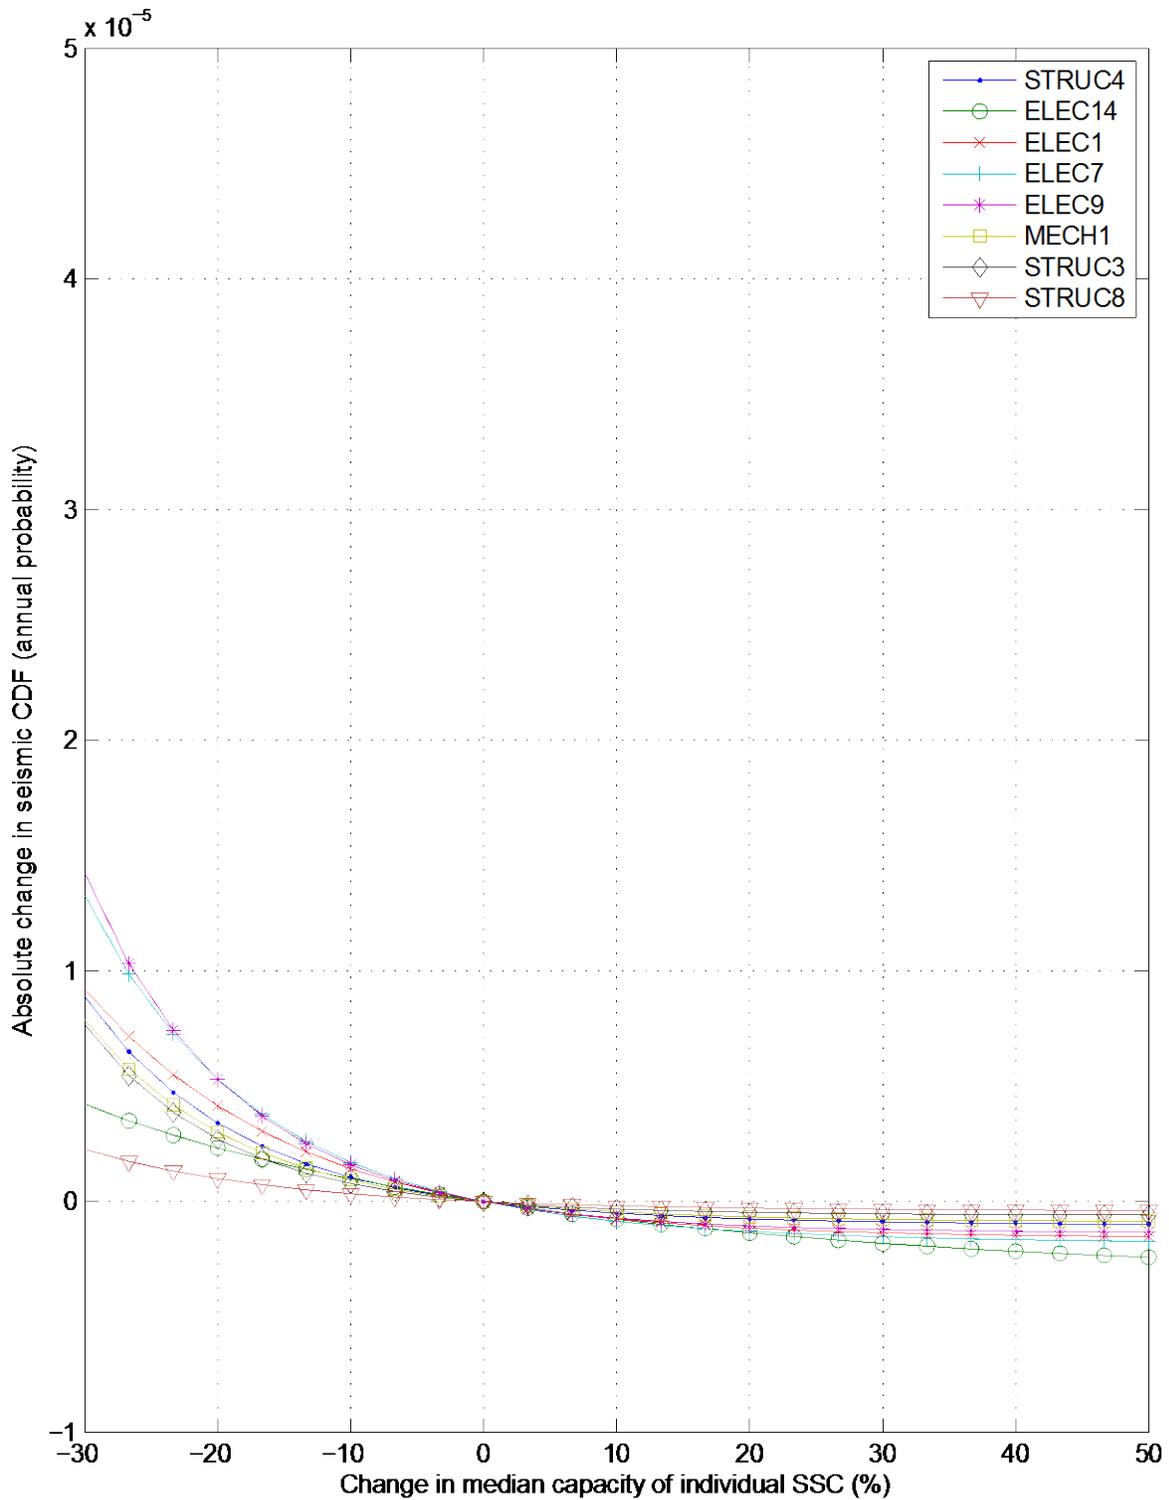


Figure 7-11 Sensitivity of the seismic CDF to one-at-a-time changes in the median seismic capacities of individual SSCs in the “enhanced balance” configuration of Plant B. The ordinate is a frequency in units of 10^{-5} per year.

7.4 Implications

Figure 7-12 displays the relative contribution of each accident sequence to the seismic risk profile for the three configurations of Plant A. The following observations can be made:

- The baseline configuration of Plant A has an unbalanced seismic risk profile in that a single accident sequence, SEQ1, accounts for approximately 40 percent of the seismic CDF.
- Both the “ASCE 43-05” and “enhanced balance” configurations have a smaller overall seismic CDF than the baseline configuration.
- Though the “ASCE 43-05” configuration reduces the absolute seismic CDF associated with SEQ1, the relative contribution of SEQ1 to the updated seismic CDF actually increases from 40 to 50 percent, resulting in an even less balanced seismic risk profile, although it must be emphasized that the plant’s seismic CDF decreases by 35 percent.
- The “enhanced balance” configuration reduces both the absolute and relative seismic CDF associated with SEQ1, resulting in a more balanced plant design where no single seismic accident sequence comprises more than 25 percent of the seismic CDF.
- Though the “ASCE 43-05” configuration has lower absolute seismic CDF than the “enhanced balance” configuration, it requires strengthening four SSCs (STRUCT1 by a factor of approximately 1.4, MECH3 by 4.0, ELEC1 by 1.9, and STRUC2 by 1.1), whereas the “enhanced balance” configuration requires strengthening only one SSC (STRUC2 by a factor of approximately 1.3).

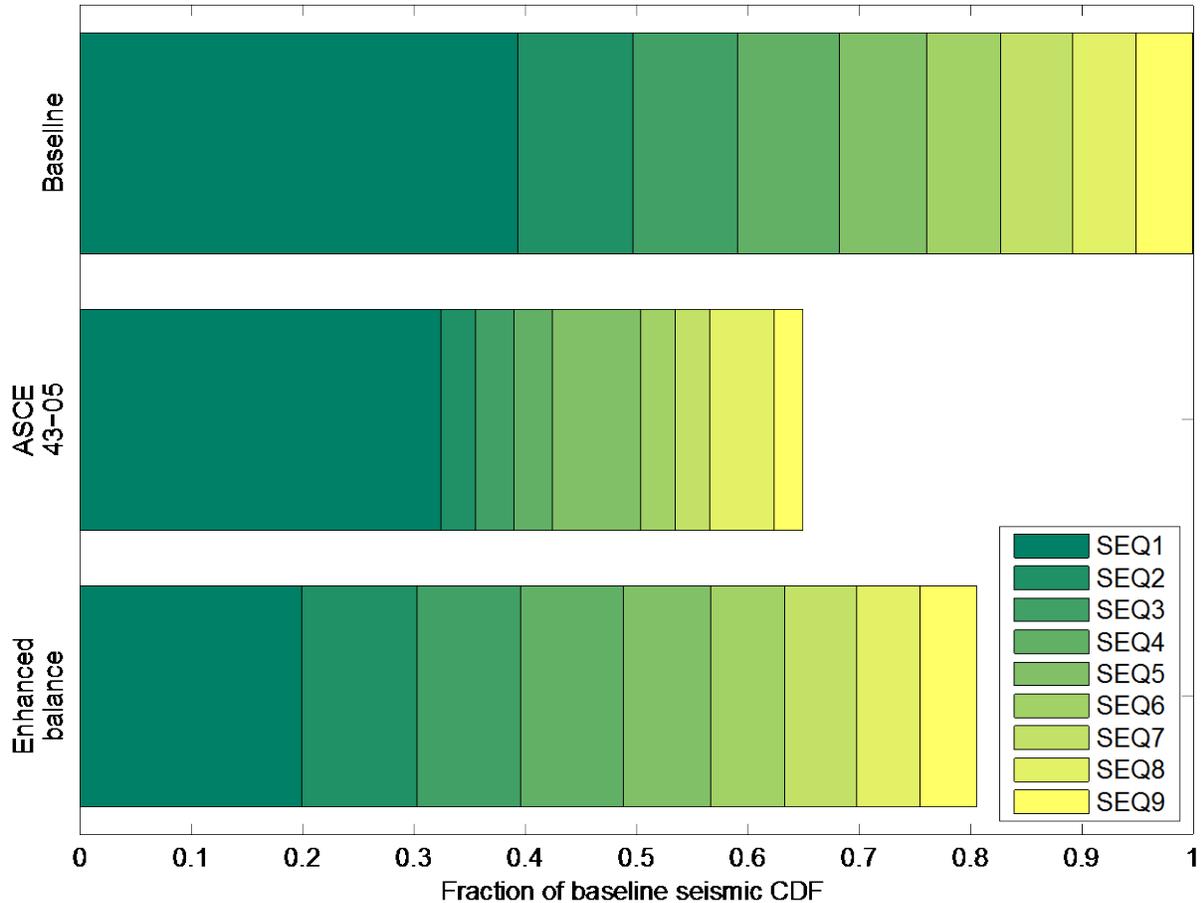


Figure 7-12 Seismic risk profile for each of the three configurations of Plant A

Figure 7-13 displays the relative contribution of each seismic accident sequence to the seismic risk profile for the three configurations of Plant B. The following observations can be made:

- The baseline configuration of Plant B has an unbalanced seismic risk profile in that two accident sequences, SEQ1 and SEQ2, account for approximately 75 percent of the seismic CDF.
- Both the “ASCE 43-05” and “enhanced balance” configurations have a smaller overall seismic CDF than the baseline configuration.
- The “ASCE 43-05” configuration reduces both the absolute and relative seismic CDF associated with SEQ2, but does not reduce the absolute seismic CDF associated with SEQ1, resulting in an unbalanced seismic risk profile in which SEQ1 comprises more than 60 percent of the updated seismic CDF. Again, it should be emphasized that the seismic CDF for the plant decreases 30 percent.
- The “enhanced balance” configuration reduces both the absolute and relative seismic CDF associated with SEQ1 and SEQ2, resulting in a more balanced plant design where no single seismic accident sequence comprises more than 25 percent of the seismic CDF.

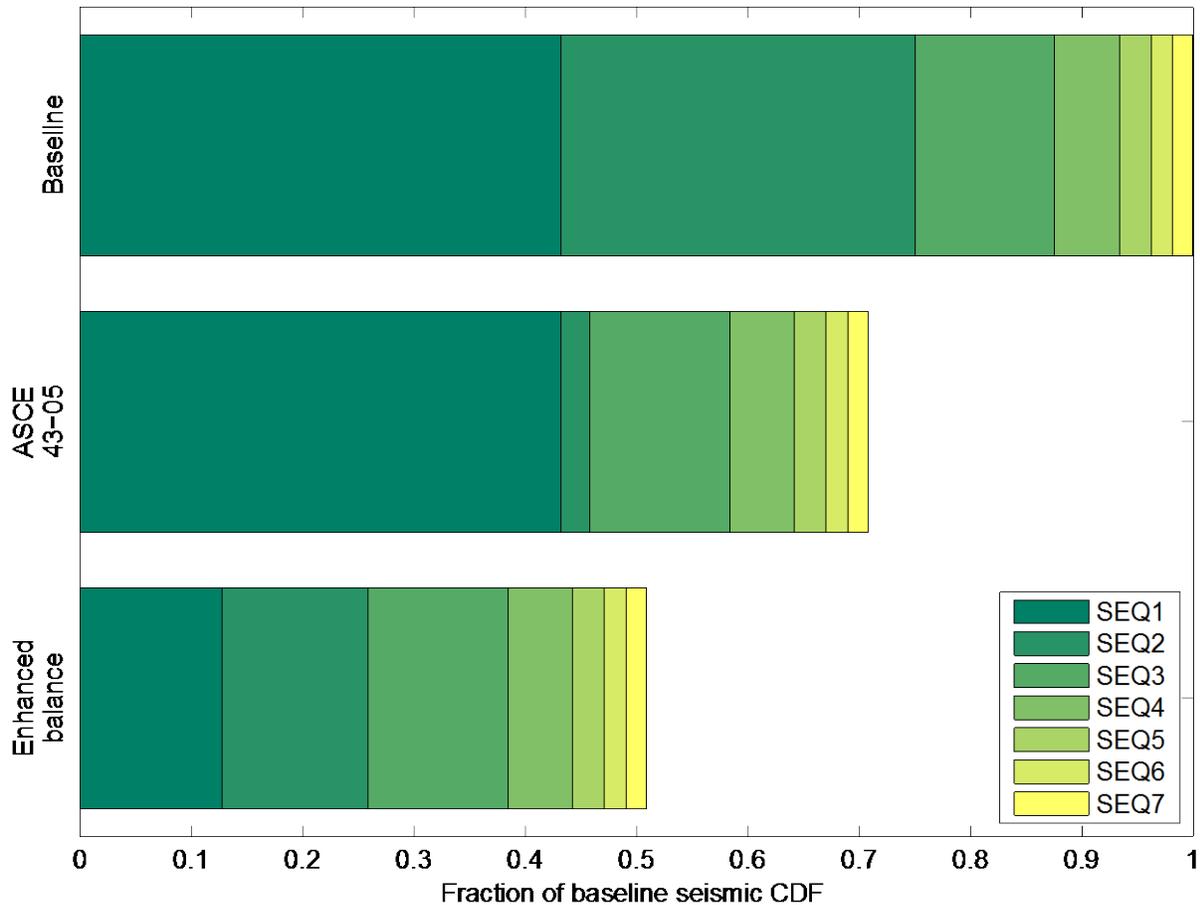


Figure 7-13 Seismic risk profile for each of the three configurations of Plant B

As the baseline configurations of the two case studies demonstrate, the provisions of the current nuclear regulatory Framework do not necessarily ensure that an NPP will have a “balanced” seismic risk profile. The whole rationale for a defense-in-depth approach to plant safety is that the overall acceptability of the plant’s safety should not rest on the strength of a singleton, nor rely on a single line of defense, nor be dependent on the correctness of the analysis of a single item. In both case studies, one or two accident sequences (or SSCs) dominate the seismic risk profile of the NPP, leaving it more vulnerable to errors in design, analysis, construction, operation, or maintenance. This vulnerability is captured in Figure 7-2 and Figure 7-5, both of which show that a reduction in the seismic capacities of certain SSCs can significantly reduce the plant’s seismic CDF, which changes the bottom-line findings of the SPRA.

While the provisions of ASCE 43-05 enhance the seismic capacities of individual SSCs, they do not necessarily produce a more balanced seismic risk profile. In contrast, a provision limiting the relative contribution of a single seismic accident sequence to the risk profile could help produce a more balanced plant seismic design. Figure 7-12 and Figure 7-13 support these conclusions, suggesting that the more effective approach to improving the plant’s seismic risk

balance involves establishing performance requirements for seismic accident sequences instead of individual SSCs. It should be noted, however, that a balanced design should not take precedence over a reduction in overall risk (*i.e.*, balance is secondary). In other words, despite its less balanced seismic risk profile, the “ASCE 43-05” configuration is preferable to the baseline configuration because its seismic CDF is smaller.

8 SUGGESTIONS FOR A PATH FORWARD

The broad scope of this report has been an analysis of the current “Framework” under which large nuclear power plants (NPPs) are designed, constructed, maintained, operated, analyzed, and regulated to achieve adequate safety against earthquakes. As noted in Section 1, “although the current approach has produced a fleet of operating NPPs and also a group of new designs that are demonstrably safe against earthquakes with adequate margin, the approach has generally failed to take advantage of several modern approaches to designing, analyzing, and regulating NPP seismic safety that could make the NPPs demonstrably safer, as well as more efficient.” The main thrust of the report is the development of the basis for a set of suggestions, including several new policy proposals, which if implemented would substantially advance the current Framework for achieving seismic safety at NPPs. These proposals could make the Framework more performance-based and risk-informed, and also make the elements of the Framework work together more effectively. The most salient of the various issues and proposals are summarized next. Please note that the body of the report contains discussions of several other technical and policy issues not summarized here.

8.1 Summary of Issues and Potential Paths Forward

8.1.1 A Balanced Risk Profile and Defense in Depth

This is discussed in Section 4.2.1(F) and Section 5. The issue or problem can be described as follows: based on a survey and review of a large number of SPRAs of existing plants (NRC 2001), the current design approach has produced operating NPPs whose seismic “risk profile” is at many plants characterized by a single “leading” or “dominant” SSC whose seismic failure is the most important in contributing to overall seismic risk. Figure 7-2 and Figure 7-5 graphically capture this vulnerability. This is in addition to the seismic vulnerability of offsite power, which is assumed to fail in almost any earthquake large enough to cause other failures. This dominance of a single SSC failure, or the failure of a very small number of them, or the failure of a group of SSCs due to a common cause, is inconsistent with the overall NRC philosophy of defense in depth, which seeks a more balanced “risk profile.” This is true even though a plant’s overall seismic design has otherwise apparently achieved an acceptable safety level and otherwise meets all applicable current regulations.

This is perhaps where one of the greatest opportunities exists for a near-term advance, as discussed in Section 5, where a detailed proposal is presented that we believe could be accomplished in the short term (in a very few years), and probably more rapidly than any of the other issues discussed in this report.

For a new NPP design, the “interim proposal” encompasses 5 “Steps”, as described in the next paragraph.

Step A involves completing a new NPP design using existing codes and standards, existing engineering practice, and existing regulations. Step B involves performing a probabilistic risk assessment (PRA) that includes a full SPRA, or that includes a Seismic Margin Assessment (SMA) using PRA-type systems-analysis methods. Step C involves using the SPRA or SMA to identify the “leading” accident sequences and “leading” SSCs arising due to potential large earthquakes. Step D involves comparing the risk profile and the seismic part of it to an overall “safety target.” This analysis is done using standard analytical techniques. Of course, if any shortfalls are found compared to this requirement, a design change is required. Step E involves

evaluating the adequacy of the design vis-à-vis an imbalance in the risk profile that goes against the fundamental ideas of defense in depth. This Step cannot now be accomplished using existing regulatory policies, because no such policies exist. In today's NRC regulatory scheme, there are no criteria beyond a general qualitative exhortation to seek out and avoid imbalances when they are identified.

The solution, as a practical matter, is for the NRC to initiate a technical debate on this subject, ultimately leading to some sort of policy guidance – perhaps a Commission policy position, perhaps a rulemaking, perhaps more general or less binding guidance.

8.1.2 Core Damage vs. Large Early Release

This is discussed in Section 3.1, Section 4.2.1(A), and Section 6.1. The issue is that the NRC's approach to seismic-safety regulation has not given as much emphasis to accident sequences involving a large early release as it has to core-damage accidents. This is a gap, and filling that gap would be an advance. If promulgated in more-or-less the same way as the criterion for core damage, in terms of a plant-level HCLPF seismic capacity, it would provide a second plant-level design target, to be used in parallel with the existing one for core damage. An advantage is that the ability to analyze such a HCLPF capacity is well within the capabilities of today's SPRA analysis experts.

This is an area ripe for debate followed by the development of an NRC regulatory position.

8.1.3 Confidence Level

This is discussed in Section 3.1, Section 4.2.1(E), and Section 6.2. There is no explicit regulatory declaration of what confidence level an analysis must meet vis-à-vis the regulatory requirement(s) in order for the design being analyzed to be in regulatory compliance. An ideal Framework, as noted in Section 3, would provide an explicit confidence level as an integral part of the regulatory acceptance criterion. As noted in Section 6.2, this is another area where a new regulatory position could be developed now or soon, although some debate in the affected community would be necessary before a final position is reached.

8.1.4 Providing an Option to Use a Performance Target

This is discussed in Section 4.2.3(K), Section 4.2.3(M), and Section 6.3. There is no provision today allowing the designer the option of designing explicitly against a performance target instead of using the design codes. A provision along these lines pre-supposes the existence of a regulatory "performance target." That is perhaps the more difficult issue, but again seems amenable to both debate (extended if necessary) and then promulgation of a regulatory position. This is an area where it should not be difficult for the NRC staff to develop a proposal for a regulatory position.

8.1.5 Systems-analysis Issues

This is discussed in Section 4.2.4 (R), Section 4.2.4 (S), and Section 6.4. In trying to optimize the design, while meeting applicable regulations, a designer has no explicit way in the current regulatory Framework to take into account the fact that non-seismic failures and human errors contribute to many important seismic accident sequences. Also, the fact that seismic failures of identical or similar co-located SSCs may fail in a correlated way is not accounted for in either the design guidance or any regulatory positions. Furthermore, the current regulatory

Framework neither asks explicitly for an analysis of the effects of these issues, nor has any context for how to deal with the information if it were made available. An advanced Framework would not only recognize the existence of these issues as contributing to the seismic “risk profile” but would provide guidance to the designer about how to use the information to improve seismic safety, or at least to demonstrate the achieved level of safety with greater confidence. This is an area where the technical basis for more advanced NRC guidance exists. A debate will be necessary with the affected community, but reaching consensus should not be difficult.

8.1.6 SSCs Are Designed Individually for Seismic Performance

This is discussed in Section 4.2.4(Q) and Section 6.5. Currently, each SSC within an NPP is designed individually to achieve an adequate seismic capacity. To advance beyond this to account for the design and the seismic safety of the “plant as a whole” will likely require a consensus code committee to address head-on how to use the systems insights from SPRA to modify the seismic design guidance to account for the broader issues. As noted in the body of the report, there are several possible approaches to address this. Perhaps an affected party could convene a “workshop” in which interested individuals can explore the issues thoroughly. Perhaps the NRC or EPRI could support a research project along these lines. It seems unlikely that progress can occur absent one or another of these precursor activities.

8.1.7 Variations in Margin Among Industry Design Codes

This is discussed in Section 2.4 and Section 6.6, where it is noted that there are significant differences in the margins embedded in the consensus design codes promulgated over the years by code committees organized under different standards-development organizations (e.g., the American Nuclear Society, American Concrete Institute, American Society of Civil Engineers, Institute of Electrical and Electronic Engineers, and others.) That there are differences in embedded margin is not surprising, given that the committees have generally worked independently and that the consensus codes represent different philosophies of design related to the different fields of engineering. To advance beyond this, so as to achieve greater harmony across the various code committees leading to greater uniformity in the margins achieved, will likely require intervention by one of two catalysts: either an overarching body like the NRC could enforce some consistency of approach, or a consortium of the major code committees could somehow come together to bring this about. Absent some intervention along these lines, a consistency of approach is unlikely to develop anytime soon.

8.1.8 Broad Institutional Barriers

This is discussed in Section 2.5 and Section 6.7. There are two broad institutional barriers to making sweeping progress on the agenda put forth in this report. One involves the NRC and the other involves the code committees.

NRC regulations

How ripe is the NRC for a reassessment and revision to embed more risk-informed and performance-based requirements into the seismic regulatory Framework, along the lines of this report? As noted earlier, recent progress bodes well for the future. Specifically, there have been several recent initiatives along these general lines. However, a good deal of debate is required, and appropriate, before a consensus can emerge on which policy advance(s) might be best.

Code committees

How ready are code committees for revisions to their deterministic approaches to seismic design and analysis along the lines of this report? Throughout the community of experts, the philosophy has been changing, slowly but inexorably, for years. Now this issue is ripe for even more movement toward performance-based design and analysis, if not risk-informed too, which is more difficult. It will take some leadership to follow up, which can probably only come from the NRC, although it is possible that it might emerge from a consensus of industry experts instead, or in addition.

NRC and industry working together

How ready are these two major stakeholders to work together? As discussed in the body of the report, the industry consensus committees and the NRC do not yet “work together” to achieve a prescribed safety target, in major part because the NRC’s current safety targets (see the discussion in Section 5.4) are not written in a form amenable for direct use by the code committees, and in part because working together has not historically been the pattern: the various code committees themselves have seldom seen either the need or the motivation to give major consideration to achieving a measure of uniformity in the safety achieved. And old habits die hard. To advance beyond this will require an explicit decision, probably by the NRC staff with policy input from management (including the Commissioners), that achieving the desired uniformity is important enough to give priority to the effort required. Recent progress at the NRC to embed more risk-informed and performance-based approaches into their regulations and other regulatory positions bodes well for the future. However, a good deal of debate is required, and appropriate, before a consensus can emerge on which policy advance(s) might be best.

8.2 Conclusion

As noted above, the broad scope of this report has been an analysis of the current “Framework” under which large NPPs are designed, constructed, maintained, operated, analyzed, and regulated to achieve adequate safety against earthquakes. The main thrust of the report is the development of the basis for a set of suggestions, including several new policy proposals, which if implemented would substantially advance the current Framework for achieving seismic safety at NPPs. These proposals could make the Framework more performance-based and risk-informed, and also make the elements of the Framework work together more effectively. The report presents a case to support the proposition that in each case, for each issue discussed, there is already (now) not only an adequate technical basis for proceeding in the “right direction” but also a well-developed basis for an advance on the policy side.

Some of the policy initiatives suggested in this report overlap with or coincide with other industry and NRC initiatives that are currently in the works or under active debate, that taken together are moving toward a more performance-based approach that accounts for risk information wherever it can. What is apparent from the discussion in this report is that in the technical arena concerned with the safety of large NPPs against earthquakes, the Framework is already more advanced than in many other safety areas, and is therefore closer to being “ripe” for even further advances – advances that can pave the way to the broader advances in other technical areas that are now under discussion.

9 REFERENCES

ASCE (2005): Standard ASCE/SEI 43-05, "Seismic Design Criteria for Structures, Systems and Components in Nuclear Facilities," American Society of Civil Engineers, Reston VA (2005)

ASCE (2013): Standard ASCE 4, "Dynamic Analysis of Nuclear Structures," American Society of Civil Engineers, Reston VA (revision of ASCE 4-98, but basically a new standard, planned for release in late 2013)

ASME-ANS (2013): American Society of Mechanical Engineers (ASME) and American Nuclear Society (ANS), "ASME-ANS RA-Sb-2013, Addendum b to RA-S-2008, "Standard for Level 1- Large Early Release Frequency Probabilistic Risk Assessment for Nuclear Power Plant Applications," American Society of Mechanical Engineers, New York NY (2013)

ASME-ANS (2015): American Society of Mechanical Engineers (ASME) and American Nuclear Society (ANS), ASME-ANS Joint Committee on Nuclear Risk Management, "General Requirements for a Level 1 PRA, Including Large Early Release Frequency, for Advanced Light-Water Reactors in the Preoperational Phases," American Society of Mechanical Engineers, New York NY (planned for release in 2015)

Budnitz (2013): Robert J. Budnitz, private communication with M. Salmon about progress of the ASCE 4 committee on a new edition. Salmon chairs that code committee.

EPRI (1991): Electric Power Research Institute, "Nuclear Plant Seismic Margin R-1," EPRI Report NP-6041, Palo Alto, CA (1991)

EPRI (2002): Electric Power Research Institute, "Seismic Fragility Application Guide," EPRI Report 1002988, Palo Alto, CA (December 2002)

Kennedy and Ravindra (1984): R.P. Kennedy and M.K. Ravindra, "Seismic Fragilities for Nuclear Power Plant Risk Studies," Nuclear Engineering and Design, Vol. 79, No. 1, pp. 47–68 (May 1984)

Reed and Kennedy (1994): J.W. Reed and R.P. Kennedy, "Methodology for Developing Seismic Fragilities," EPRI Report TR-103959, Electric Power Research Institute, Palo Alto, CA (1994)

NRC (1977): US Nuclear Regulatory Commission, Code of Federal Regulations, Title 10, Appendix A to Part 100, "Seismic and Geological Siting Criteria for Nuclear Power Plants," USNRC, Washington, DC (1977)

NRC (1985): US Nuclear Regulatory Commission, NUREG/CR 4334, "An Approach to the Quantification of Seismic Margins in Nuclear Power Plants," Expert Panel (R.J. Budnitz, P.J. Amico, C.A. Cornell, W.J. Hall, R.P. Kennedy, J.W. Reed, and M. Shinozuka), USNRC, Washington, DC (1985)

NRC (1986): US Nuclear Regulatory Commission, "Safety Goals for the Operation of Nuclear Power Plants: Policy Statement," Federal Register 51, 30028, USNRC, Washington, DC (August 21, 1986)

NRC (1990): US Nuclear Regulatory Commission, NUREG 1150, "Severe Accident Risks: An Assessment of Five U.S. Nuclear Plants," USNRC, Washington, DC (1990)

NRC (1993): US Nuclear Regulatory Commission, SECY-93-087, "Policy, Technical, and Licensing Issues Pertaining to Evolutionary and Advanced Light Water Reactor (ALWR) Designs" (April 2, 1993) and the Staff Requirements Memo SRM-93-087 that accompanied it, USNRC, Washington, DC (July 21, 1993)

NRC (1996): US Nuclear Regulatory Commission, Code of Federal Regulations, Title 10, Appendix S to Part 50, "Earthquake Engineering Criteria for Nuclear Power Plants," USNRC, Washington, DC (1996)

NRC (1997): US Nuclear Regulatory Commission, Regulatory Guide 1.165, "Identification and Characterization of Seismic Sources and Determination of Safe Shutdown Earthquake Ground Motion," USNRC, Washington, DC (1997)

NRC (2001): US Nuclear Regulatory Commission, NUREG-1742, "Perspectives Gained from the Individual Examination of External Events (IPEEE) Program," in two volumes, USNRC, Washington, DC (2001)

NRC (2004): US Nuclear Regulatory Commission, Code of Federal Regulations, Title 10, Part 50, Section 50.69, "Risk-Informed Categorization and Treatment of Structures, Systems and Components for Nuclear Power Reactors," USNRC, Washington, DC (2004)

NRC (2004a): US Nuclear Regulatory Commission, NUREG-BR-0058, "Regulatory Analysis Guidelines of the US Nuclear Regulatory Commission, Revision 4," USNRC, Washington, DC (2004)

NRC (2007): US Nuclear Regulatory Commission, Regulatory Guide 1.208, "A Performance-based Approach to Define the Site-specific Earthquake Ground Motion," USNRC, Washington, DC (2007)

NRC (2010): US Nuclear Regulatory Commission, DC/COL-ISG 20, "Interim Staff Guidance on Implementation of a Probabilistic Risk Assessment-Based Seismic Margin Analysis for New Reactors," USNRC, Washington, DC (2010)

NRC (2012): US Nuclear Regulatory Commission, NUREG-1614, Volume 5, "Strategic Plan, Fiscal Years 2012-2016," USNRC, Washington, DC (2012)

NRC (2012a): US Nuclear Regulatory Commission, NUREG-2150, "A Proposed Risk Management Regulatory Framework," USNRC, Washington, DC (2012)

NRC (2012b): US Nuclear Regulatory Commission, Letter to All Power Reactor Licensees *et al.*, "Request for Information Pursuant to Title 10 of the Code of Federal Regulations 50.54(f) Regarding Recommendations 2.1, 2.3 and 9.3 of the Near-Term Task Force Review of Insights from the Fukushima Dai-Ichi Accident," USNRC, Washington, DC (March 12, 2012)

NRC (2012c): US Nuclear Regulatory Commission, JLD-ISG-2012-04, "Interim Staff Guidance on Performing a Seismic Margin Assessment in Response to the March 2012 Request for Information Letter," USNRC, Washington, DC (2012)

BIBLIOGRAPHIC DATA SHEET

(See instructions on the reverse)

2. TITLE AND SUBTITLE

Toward a More Risk-Informed and Performance-Based Framework for the Regulation of the Seismic Safety of Nuclear Power Plants

3. DATE REPORT PUBLISHED

MONTH

YEAR

May

2016

4. FIN OR GRANT NUMBER

5. AUTHOR(S)

Robert J. Budnitz and Michael W. Mieler

6. TYPE OF REPORT

Technical

7. PERIOD COVERED (Inclusive Dates)

September, 2010 to July, 2016

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.)

Lawrence Berkeley National Laboratory
Energy Geosciences Division, Mail Stop 74-316C
University of California
Berkeley, CA 94720

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
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001

10. SUPPLEMENTARY NOTES

11. ABSTRACT (200 words or less)

Dozens of seismic probabilistic risk assessments (SPRAs) exist that have analyzed large nuclear power plants (NPPs). A major insight from these SPRAs is that, although generally the plants are adequately safe against earthquake threats, the way the industry and the United States Nuclear Regulatory Commission (NRC) currently design, build, operate, analyze, and regulate seismic NPP safety may not be optimal. This suboptimal situation means that for both operating NPPs and new NPP designs not yet built, they may fail to take advantage of possible additional safety insights and improvements, may be more difficult to analyze and to regulate than they need to be, and may cost more than they otherwise would. There is room for improvement in several areas. The major topics covered herein are the variations in residual seismic risk from plant to plant; the unbalanced risk profiles and incomplete defense in depth achieved at many plants; the impact of structures and components being designed individually for seismic performance rather than taking a systems approach; the variations in margin that exist among the industry consensus codes and standards for seismic design and analysis; and the observation that the design codes and the NRC regulations do not work together well. Each of these topics is analyzed in detail, and case-study examples based on the seismic risk profiles of two operating plants are used to illustrate the issues. Suggestions are presented that would improve the seismic framework in each of these areas.

12. KEY WORDS/DESCRIPTORS (List words or phrases that will assist researchers in locating the report.)

Earthquake Resistant Design; Seismic Design Criteria; Performance-Based Design; Performance Targets; Risk-Informed Design; Nuclear Power Plants; Seismic Probabilistic Risk Assessment; Core Damage Frequency; Large Early Release; Risk Profiles; Seismic Fragility; Regulatory Framework; Design Codes and Standards; American Society of Civil Engineers; Institute of Electrical and Electronic Engineers; American Society of Mechanical Engineers; American Nuclear Society

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



NUREG/CR-7214

**Toward a More Risk-Informed and Performance-Based Framework for the
Regulation of the Seismic Safety of Nuclear Power Plants**

May 2016