UNITED STATES
NUCLEAR REGULATORY COMMISSION
ADVISORY COMMITTEE ON REACTOR SAFEGUARDS
WASHINGTON, DC 20555 - 0001

March 24, 2004

MEMORANDUM TO: G. Apostolakis, Chairman
Reliability and Probabilistic Risk Assessment Subcommittee

S. Rosen, Chairman
Human Factors Subcommittee

M. Bonaca, Member

Thomas Kress, Member

D. Powers, Member

W. Shack, Member

FROM: B. P. Jain, Senior Staff Engineer ***/RA/***
ACRS/ACNW

SUBJECT: STATUS REPORT FOR THE ACRS JOINT SUBCOMMITTEES (RELIABILITY & PROBABILISTIC RISK ASSESSMENT AND HUMAN FACTORS) MEETING ON STAFF'S GUIDANCE - GOOD PRACTICES FOR IMPLEMENTING HUMAN RELIABILITY ANALYSIS (HRA)

The ACRS Subcommittees on Reliability and Probabilistic Risk Assessment and on Human Factors will be meeting on April 22, 2004, to review the proposed staff's guidance regarding 'Good Practices for Implementing Human Reliability Analysis (HRA)' and data development for Human Event Repository and Analyses (HERA). The Subcommittees will hear presentations by and hold discussions with representatives of the staff and its contractors. The purpose of this meeting is to gather information, analyze relevant issues and facts, and formulate proposed positions and actions, as appropriate, for deliberation by the full Committee at the May 6-8, 2004 ACRS meeting. The staff requests ACRS concurrence for issuing the staff's proposed guidance on 'Good Practices for Implementing Human Reliability Analysis (HRA)' for public comment.

To prepare the Subcommittees for the meeting, the following documents are attached:

A. Proposed Schedule

B. Status Report

C. 'Good Practices for Implementing Human Reliability Analysis (HRA),' Draft Letter Report (JCN W6994) dated March 19, 2004.

D. 'Expert Elicitation Approach for Performing ATHEANA Quantification', a paper by J. Forester et al./Reliability Engineering and System Safety 83 (2004) 207-220.

Those members who are not scheduled to attend the joint Subcommittee meeting on April 22, 2004, should send their comments to Subcommittee Chairman Dr. G. Apostolakis with a copy to me by April 15, 2004.

If you have any questions, please call me (301-415-7270) or E-mail (bpj@nrc.gov).

Attachment: As stated

cc w/attachment: Remaining ACRS Members

cc wo/attachment: J. Larkins
M. Weston
S. Duraiswamy

ADVISORY COMMITTEE ON REACTOR SAFEGUARDS
JOINT MEETING OF SUBCOMMITTEES ON
RELIABILITY & PROBABILISTIC RISK ASSESSMENT AND HUMAN FACTORS
GOOD PRACTICES FOR IMPLEMENTING HUMAN RELIABILITY ANALYSIS (HRA)
APRIL 22, 2004
ROCKVILLE, MARYLAND

**-PROPOSED SCHEDULE-**

ACRS Contact: B.P. Jain (301-415-7270)

**THURSDAY, APRIL 22, 2004, CONFERENCE ROOM T-2B3, TWO WHITE FLINT NORTH, ROCKVILLE, MARYLAND**

| | Topics | Presenters | Time |
|---|---|---|---|
| **I.** | **Opening Remarks** | G. Apostolakis/S. Rosen<br>ACRS | 8:30- 8:40 a.m. |
| **II.** | **Introduction** | D. Lew/E. Lois, RES | 8:40- 8:50 a.m. |
| **III.** | **HRA Good Practices** | A. Kolaczkowski, SAIC | 8:50-10:15 a.m. |
| | | **BREAK** | 10:15-10:30 a.m. |
| **IV.** | **ATHEANA Quantification** | J. Forester, SNL | 10:30-11:45 a.m. |
| **V.** | **Plans for Improving ATHEANA Practices** | S. Cooper, RES | 11:45-12:00 noon |
| | | **Lunch** | 12:00-1:00 p.m. |
| **VI.** | **Human Event Repository and Analyses (HERA)** | B. Hallbert, INEEL | 1:00-1:45 p.m. |
| **VII.** | **Halden HRA Activities** | A. Bye, Halden | 1:45-2:15 p.m. |
| **VIII.** | **Subcommittee Discussion** | | 2:15-2:30 p.m. |
| | **Adjourn** | | 2:30 p.m. |

NOTE: Presentation time should not exceed 50% of the time allocated for a specific item. The remaining 50% of the time is for Subcommittee questions.

ADVISORY COMMITTEE ON REACTOR SAFEGUARDS
JOINT MEETING OF SUBCOMMITTEES ON
RELIABILITY & PROBABILISTIC RISK ASSESSMENT AND HUMAN FACTORS
GOOD PRACTICES FOR IMPLEMENTING HUMAN RELIABILITY ANALYSIS (HRA)
APRIL 22, 2004
ROCKVILLE, MARYLAND

**- STATUS REPORT -**

PURPOSE:

The purpose of the joint meeting of the Reliability and Probabilistic Risk Assessment and Human Factors Subcommittee meeting is to discuss the proposed staff's guidance on 'Good Practices for Implementing Human Reliability Analysis (HRA)' and development of data for Human Event Repository and Analyses (HERA). The Subcommittees will hear presentations by and hold discussions with representatives of the staff and its contractors. The staff is seeking Committee's views on its guidance document and concurrence for publishing it for public comment.

BACKGROUND:

The staff developed the guidance on 'Good Practices for Implementing Human Reliability Analysis (HRA)' for performing and reviewing HRAs as a document supporting Regulatory Guide 1.200. Regulatory Guide 1.200 describes an acceptable approach for determining the technical adequacy of PRA results for risk informed activity, and reflects and endorses guidance provided by the American Society of Mechanical Engineers (ASME) Standard for Probabilistic Risk Assessment for Nuclear Power Plants and the Probabilistic Risk Assessment Peer Review Process Guidance developed by Nuclear Energy Institute (NEI-00-02).

## A. Good Practices for Implementing Human Reliability Analysis

The purpose of the guidance regarding 'Good Practices for Implementing Human Reliability Analysis (HRA)' is to ensure some level of consistency and quality in HRA analyses and their review:

2. It provides guidance for performing a good HRA (whether for the first time or when analyzing a change to current plant practices) when implementing the ASME Standard, and focuses on the attributes of a good HRA regardless of the specific methods or tools that are used. The guidance is specifically for HRAs for reactors operating at full power and internal events applications although most of the guidance may prove to be useful for other applications (e.g., external events, other operating modes...). It does not endorse nor is it meant to suggest that a specific method or tool be used since many exist, and all have strengths and limitations regarding their use and applicability.

3. It provides guidance for assessing the quality of HRAs. In this regard, the practices of a good HRA are provided which should be useful in formulating questions about and measuring the "goodness" of a HRA. Its purpose is not to explicitly provide questions a reviewer should ask, but rather to provide the technical basis for developing questions or a standard review plan for the staff's review of HRA.

The staff's "HRA Good Practices" guidance is being developed in two phases. The first phase is the development of this "HRA Good Practices" document which has been prepared on the basis of the staff's experience and lessons learned from developing HRA methods (e.g., ATHEANA), and performing and reviewing HRAs. The second phase is a review and evaluation of existing HRA approaches for their capability to meet the good practices when employed to address different regulatory applications.

This "HRA Good Practices" document describes the staff's views regarding good practices of an HRA as implemented within a broader PRA framework. This is written in the context of a risk assessment for commercial nuclear power plant operations occurring nominally at full power. However, it is likely that many of the good practices will also be applicable to low power and shutdown operations. As with any evolving technology, both PRA and the implementation of HRA within the PRA framework are continuing to improve. Hence, what is good practice today may be somewhat inferior or outdated tomorrow.

**B. A Technique for Human Event Analysis (ATHEANA) Quantification**

The staff developed AETHENA, a HRA method, to increase the degree to which HRA can represent the kinds of human behaviors seen in accidents and near-miss events at nuclear power plants. An expert elicitation approach has been developed to estimate probabilities for unsafe human actions (UAs) based on error-forcing contexts (EFCs) identified through the ATHEANA search process. The expert elicitation approach integrates the knowledge of informed analysts to quantify UAs and treat uncertainty ('quantification-including-uncertainty'). The analysis focuses on:

(a) the probabilistic risk assessment (PRA) sequence EFCs for which the UAs are being assessed,

(b) the knowledge and experience of analysts (including trainers, operations staff, and PRA/human reliability analysis experts), and

(c) facilitated translation of information into probabilities useful for PRA purposes.

The expert elicitation approach emphasizes asking the analysts what experience and information they have that is relevant to the probability of failure rather than simply asking the analysts their opinion about failure probabilities. The facilitator then leads the group in combining the different kinds of information into a consensus probability distribution. The attached technical paper on the subject (co-authored by the staff) describes the expert elicitation process, presents technical basis, and discusses the controls that are exercised to use it appropriately. The paper also points out the strengths and weaknesses of the approach and how it can be improved.

EXPECTED SUBCOMMITTEE ACTION

The joint Subcommittee Chairman plans to provide a proposed ACRS report to the full Committee at the May 6-8, 2004 ACRS meeting for consideration and approval.

JCN W6994
Draft Letter Report

# Good Practices for Implementing Human Reliability Analysis (HRA)

Date: March 19, 2004

Prepared by
Alan Kolaczkowski and John Forester

Sandia National Laboratories
Albuquerque, NM 87185

USNRC Project Manager: Erasmia Lois

Prepared for
Probabilistic Risk Analysis Branch
Division of Risk Analysis and Applications
Methods Group
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555

# Table of Contents

# Good Practices for Implementing Human Reliability Analysis (HRA)

## 1. INTRODUCTION

### 1.1 Background

In accordance with its policy statement [1] on the use of probabilistic risk assessment (PRA), during the last decade the NRC has been increasingly using PRA technology in "all regulatory matters to the extent supported by the state of the art in PRA methods and data." Examples of risk informed initiatives are: undertaking risk-informed rulemaking activities such as risk-informing 10CFR Part 50[2], generating a risk-informed framework for supporting licensee requests for changes to a plant's licensing basis (Reg Guide 1.174),[3] risk-informing the reactor oversight process, performing risk studies (e.g., for steam generator tube rupture (SGTR), and fire events), and evaluating the significance of events. In addition, the NRC is using PRA in the development of an infrastructure to licence new reactors.

Given the increasing importance of the role of PRA in regulatory decision making, it is crucial that decision makers have confidence in the results produced by PRAs. To support this, the NRC has issued Regulatory Guide 1.200[4] that describes an acceptable approach for determining the technical adequacy of PRA results for risk informed activity. Reg Guide 1.200[4] reflects and endorses guidance provided by standards produced by societies and industry organizations. It currently addresses the American Society of Mechanical Engineers (ASME) Standard for Probabilistic Risk Assessment for Nuclear Power Plant Applications[5] which was developed for a full power, internal events (excluding fire) Level 1 PRA and a limited Level 2 PRA, and the Probabilistic Risk Assessment Peer Review Process Guidance (NEI-00-02).[6]

The level of detail provided in the ASME Standard[5] and NEI-00-02[6] is at a high level, addressing what to do, but not how to do it., Consequently, there may be several approaches to address certain analytical elements, which though they may meet the standards, may do so by making different assumptions and approximations, and, therefore, produce different results. This is particularly true of human reliability analysis (HRA) (see section 1.2 for a discussion of HRA). Therefore, the guidance provided by these documents is not sufficient to address the detailed HRA quality issues needed to be considered in regulatory decision making. For example, in section A.8, Modeling of Human Performance, in Standard Review Plan 19,[7] the NRC staff is required to determine if "the modeling of human performance is appropriate." While the ASME Standard[5] and NEI-00-02[6] can address whether the HRA addresses the right issues, they do not give guidance on how they are addressed. Therefore, in order to support the review of human performance issues in the context of PRAs, the NRC is developing this guidance for performing and reviewing HRAs, as a document supporting Reg Guide 1.200[4]. The guidance is being developed in two phases. The first phase is the development of this "HRA Good Practices" document which has been prepared on the basis of the NRC experience and lessons learned from developing HRA methods (e.g., THERP,[8] SLIM,[9] and ATHEANA[10]), performing HRAs (e.g., NUREG-1150[11] studies, and reviewing HRAs (in particular the individual plant examinations [IPEs]). The second phase is a review and evaluation of existing HRA approaches for their capability to meet the good practices when employed to address different regulatory applications.

This volume describes the NRC staff views regarding good practices of an HRA as implemented within a broader PRA. The volume is written in the context of a risk assessment for commercial nuclear power plant (NPP) operations occurring nominally at full power. However, it is likely that many of the good practices will also be applicable to low power and shutdown operations. Similarly, the volume is purposely aimed for applications involving internal initiating events but should generally be appropriate for external initiating events. Additionally, elements of this volume may be of benefit in examining human actions related to nuclear materials and safeguard types of applications.

As with any evolving technology, both PRA and the implementation of HRA within the PRA framework are continuing to improve. Hence, what is good practice today may be somewhat inferior or outdated tomorrow. Much of what is in this volume will always constitute good practice; some of it may be subject to newer technology, methods, and tools. For this reason, this volume must be considered a snapshot of good practices in HRA circa 2004.

With the expectation that PRA will continue to be used in the commercial nuclear industry in assessing current operating risks, in estimating changes in risk as a result of temporary and permanent plant changes to existing plants, and as an adjunct to the design process of newer generation plants, it is important that HRA practitioners perform human reliability analyses in accordance with good practices and that reviewers recognize the implementation of good practices (or failure to do so) in these analyses.

## 1.2 HRA in the Context of PRA

Human reliability analysis in the PRA context is that discipline that identifies and provides probabilities for the human failure events that can negatively impact normal or emergency plant operations. The human failure events modeled in PRAs that are associated with normal plant operation include: 1) events that leave equipment in an unrevealed, unavailable state, such as miscalibration of a level sensor, 2) those that induce an initiating event, such as a human-caused loss of feedwater (typically captured by the initiating event frequency), or 3) those modeled as human events contributing to an initiating event, such as a total loss of service water (e.g., failing to backup the start of service water train B upon loss of train A). The human failure events modeled in PRAs associated with emergency plant operation include events that, if not performed, do not allow the desired function to be achieved, such as failing to initiate feed and bleed. Quantification of the probabilities of the human failure events is based on plant and accident specific conditions, where applicable, including any dependencies among actions and conditions.

This volume provides HRA good practices that when implemented will result in determining the impacts of human actions as *realistically as necessary* in an assessment of risk. Note the emphasis on realistic as necessary rather than as realistic as possible. For example, depending on the purpose for which the PRA is to be used, a conservative (i.e. non-realistic) treatment of human performance may be sufficient to address a PRA application; more realism may not be necessary and could be a waste of resources. However, a conservative approach may not be sufficient when used as the basis for not needing to further investigate the issue at hand. Such an approach could potentially constrain the capability of identifying weaknesses in plant operations and plant practices related to the particular human actions credited in the PRA.

Recognizing that the volume will be used to guide a wide variety of applications, it is not intended that all the practices be met for any specific PRA application; in fact, some may not be applicable or necessary. A practitioner or reviewer should determine the applicable good practices for the PRA application and perform or review the HRA accordingly.

### 1.3 Purpose

This volume serves as a reference guide of good practices in HRA. By good practices we mean those processes and individual analysis tasks and judgments that would be expected of a HRA (considering current knowledge and state-of-the-art) in order for the HRA results to sufficiently represent the anticipated operator performance when making risk-informed decisions. The document is principally focused on the process for performing HRA and does not, for instance, specifically address HRA data or details of specific quantification approaches. As such, it is written in a way that links the prescribed good practices to requirements in the ASME Standard[5] and particularly the HRA section of that document (although nearly all other sections of the standard also have some parallel requirements with regard to operator actions such as in the accident sequence analysis, success criteria, systems analysis, and large early release frequency (LERF) analysis sections).

With this in mind, this volume has at least two primary uses.

1. It provides guidance for performing a good HRA (whether for the first time or when analyzing a change to current plant practices) when implementing the ASME Standard,[5] and focuses on the attributes of a good HRA regardless of the specific methods or tools that are used. The guidance is specifically for HRAs for full power, reactor, and internal events applications although most of the guidance may prove to be useful for other applications (e.g., external events, other operating modes...). It does not endorse nor is it meant to suggest that a specific method or tool be used since many exist, and all have strengths and limitations regarding their use and applicability. Nevertheless, the good practices come from those advocated in such sources as the ASME Standard[5], THERP[8], ASEP[12], SHARP1[13], SPAR-H Method[14], and ATHEANA[10] for example, as well as the experiences of the authors and reviewers of this volume.

2. It supports the review of HRAs in assessing the quality of the analyses. In this regard, the practices of a good HRA are provided which should be useful in formulating questions about and measuring the "goodness" of a HRA. Its purpose is not to explicitly provide questions a reviewer should ask, but rather to provide the technical basis for developing questions or a standard review plan for the staff's review of HRA.

# 2. OVERVIEW OF GOOD PRACTICES FOR HRA

## 2.1 Scope of HRA Good Practices Guidance

The purpose of this document on good practices for implementing HRA is to ensure some level of consistency and quality in HRA analyses and their review. In order to achieve such consistency and quality, the HRA good practices in this document are directed at specific HRA tasks or activities.

The performance of HRA typically involves several tasks or activities. Some of these tasks are dependent on the HRA method or quantification approach that is used. Because this HRA good practices document does not endorse or specify the use of specific HRA methods or quantification approaches, most of the guidance in this document is directed at the process for performing HRA. However, this document does provide some non-method-specific good practices with respect to HRA quantification.

As stated in Section 1, the ASME Standard[5] already addresses these HRA tasks or activities at a high level. In the NRC's judgement, the more detailed guidance given in this document on HRA good practices is necessary to achieving acceptable consistency and quality in HRA.

## 2.2 HRA Good Practices and the State-of-the-Art in HRA

The HRA good practices given in this document are based in part on past experience in performing and reviewing HRAs, including that used to support the IPEs, but also reflect current perspectives on the issues that impact human performance that were gained from developmental projects such as ATHEANA[10]. Consistent with the state of the art in PRAs, errors of commission are not necessarily expected to be included in PRAs, although some guidance is given for identifying characteristics of situations that can facilitate errors of commission. As stated above, these good practices apply to the use of all HRA methods and approaches.

## 2.3 Summary and Organization of HRA Good Practices Guidance

The good practices are presented in a logical analysis approach and linked to the requirements of the ASME Standard.[5] Like the standard, this document specifically addresses pre-initiator (i.e., normal operations) and post-initiator (i.e., emergency operations) human actions since it is assumed that as typical of most PRAs, human actions that cause or contribute to initiating events are already accounted for quantitatively in many initiating event frequencies. Further understanding of specific causes of the initiators is typically not required. It is noted that for support system initiators and other initiators such as those human-induced initiators that may be modeled for other modes (e.g.,shutdown), corresponding initiator fault tree models may specifically include human failure events (HFEs) that have characteristics of either pre- or post-initiating event HFEs. The techniques used to analyze these HFEs are therefore covered by this document and should be followed. For example, see HLR-IE-C high level requirement in the ASME Standard[5] and such supporting requirements as IE-C9 concerning the modeling of recovery actions in an initiator fault tree, and IE-C12 concerning procedural influences on the interfacing system loss of coolant accident (ISLOCA) frequency.

While this document is written in a serial fashion, in practice, it is often desirable to perform or review an HRA in a more holistic manner and address multiple steps of the HRA process simultaneously to achieve greater resource efficiency.

Table 2-1 provides brief summaries of the good practices that are discussed in subsequent sections of this document (to be provide later).

**Table 2-1 Summary of Good Practices**

*[Table 2-1 to be inserted later]*

## 3. HRA TEAM FORMATION AND OVERALL GUIDANCE

If human actions are going to be included realistically in the PRA, the modeling of human interactions must consider each action evaluated in the context of a complete accident scenario or sequence of events. To do this, HRA has evolved from the days when PRA analysts provided the human events of interest to a HRA specialist who then assigned human error probabilities (HEPs) to the human events, often in isolation. Such a process is no longer considered good practice. Understanding an accident sequence context is a complex, multi-faceted process. The interaction of plant hardware response and the response of plant operators must be investigated and modeled accordingly. Such characteristics as the following need to be understood and reflected, as necessary, in the model of a specific human action or group of actions:

- plant behavior and conditions,
- timing of events and the occurrence of human action cues,
- the parameter indications used by the operators and changes in those parameters as the scenario proceeds,
- the time available and locations necessary to take the human actions,
- the equipment available for use by the operators based on the sequence ,
- the environmental conditions under which the decision to act must be made and the actual response must be performed,
- the degree of training guidance and procedure applicability, among many other characteristics.

Much of the guidance in this volume is aimed at good practices for understanding the context associated with each modeled human action, and how that context affects both the definition of human failure events and an assessment of their probabilities.

This emphasis on the need to adequately understand and address context in order to more realistically address human performance is based on advances in our understanding of the factors that can influence human performance. These advances come from recent reviews of operational events involving serious accidents (e.g., ATHEANA[10]) and from other international efforts and recent research in the cognitive sciences that together have provided a clearer picture of the ways in which various factors and situations can interact to influence the occurrence of inappropriate human actions (e.g., Reason[15], Woods[16], Endsley[17]...). Improvements have been made for how to address the broad range of potential influences on human performance, for both the identification of the human actions to be modeled in the PRA as well as what to consider during screening and detailed quantification of the actions. The guidance in this volume provides good practices that reflect these improvements and ensures the proper treatment of context in performing a reasonably realistic HRA.

Hence, the modeling of human actions in the PRA should involve an integrated effort among PRA modelers, HRA and human factors practitioners, thermal-hydraulic analysts, operations and maintenance personnel, and sometimes other disciplines depending on the accident sequence (e.g., structural engineers such as if the timing of an action is dependent on when and how the containment might fail). Each discipline provides a portion of the context knowledge. When the context is sufficiently understood, only then can human failure events be realistically modeled and quantified. In addition, as good practice in HRA, it is encouraged that there be the use of walkdowns of areas where the action needs to take place, talk-throughs of the scenarios and actions of interest with plant operators or maintenance personnel, field observations, and at least for the more important actions, simulations of the human actions to be credited. Finally, the HRA should be performed consistently for both core damage prevention/mitigation and large early release prevention/mitigation since both measures are considered in making risk-informed decisions as addressed in Regulatory Guide 1.174[3].

Therefore, in summary and as the first measure of a good HRA, it should be clear that an HRA assessment has utilized an integrated team and tools as summarized in Table 3-1 to the extent necessary and practical for the PRA application and the specific issue being addressed. This is an important aspect that should lead to HRA results that are credible.

**Table 3-1 Overall HRA Good Practices**

| |
|---|
| 1. The HRA is an integral part of the PRA (not performed as an isolated task in the PRA process) whereby the inputs from the following types of disciplines are used together to define the PRA structure including which human events need to be modeled, how they are defined and modeled in the PRA, and the considerations used to quantify the associated HEPs:<br>• PRA modelers<br>• HRA practitioners<br>• Thermal-hydraulic analysts<br>• Operations and maintenance personnel<br>• Other disciplines (e.g., structural engineers, system engineers...) as necessary |
| 2. Besides the review of plant documents, the HRA is performed using the insights gained from the following to confirm judgments and assumptions made from the document review:<br>• Walkdowns of areas where decisions and actions are to take place<br>• Talk-throughs of scenarios and actions of interest<br>• Field observations<br>• Simulator exercises |
| 3. As part of the integrated effort, the HRA is performed consistently for both core damage and large early release outcomes, since both are equally important in risk-informed applications. |

# 4. PRE-INITIATOR HRA

The ASME Standard[5] separates its requirements into two broad classifications; those that address the modeling of failures of pre-initiator human actions and those that address the modeling of failures of post-initiator human actions. This section provides good practices for implementing the requirements for addressing pre-initiator human failure events in a PRA.

Pre-initiator human failure events are events that represent the impact of human failures committed during actions performed prior to the initiation of an accident sequence (e.g., during test or maintenance or the use of calibration procedures). They are important to model because plant personnel can make the equipment needed to mitigate a particular accident sequence unavailable, thus reducing the overall capability to respond to the initiating event. Hence, depending on the issue being addressed, this impact may need to be included in a PRA if a realistic assessment of risk is required.

The following good practices are categorized under four major analysis activities for doing pre-initiator HRA. These analysis activities are:

1. Identifying activities that have the potential to result in pre-initiator human failures
2. Screening out the activities for which human failures do not need to be modeled
3. Modeling specific human failure events (HFEs) corresponding to the unscreened activities
4. Quantifying the corresponding human error probabilities (HEPs) for the specific HFEs.

## 4.1 Identifying potential pre-initiator human failures

4.1.1 OBJECTIVE: To identify from routine plant actions, those pre-initiator actions whose failure to perform correctly could result in the human-induced unavailability of PRA-modeled equipment that is credited in the PRA accident sequences. This is important since these actions represent other potential modes of unavailability of the credited equipment (besides the equipment simply failing to start or other failure modes in the PRA) that contribute to overall plant risk. Note that not all the identified actions will be modeled since some may be screened from further analysis in the following analysis activity (screening). The following provides good practices for identifying potential pre-initiator human failures while implementing the related Standard requirements.

4.1.2 CORRESPONDING ASME STANDARD REQUIREMENTS:

The Standard calls for a systematic process to be used to identify routine activities that if not completed correctly, may impact the availability of equipment. There are multiple supporting requirements in the Standard that address the need to consider test and maintenance activities, calibration activities, and actions that could affect multiple equipment.

4.1.3 GOOD PRACTICES:

4.1.3.1 Good Practice #1:

The HRA process should include a review of the following:

- All routine (scheduled) test and maintenance as well as calibration procedures that affect equipment to be credited in the PRA (for core damage frequency (CDF) and LERF) should be identified and reviewed.

- Actions specified in the above procedures that realign equipment outside their normal operation or standby status, or otherwise could detrimentally affect the functionality of credited equipment if not performed correctly (e.g., miscalibration) should be identified.

- "Affected" equipment should include (if routinely acted on and credited in the PRA):

  - the primary systems, structures, and components (SSCs) (e.g., emergency core cooling systems' components, containment cooling systems' components...),

  - support systems (e.g., power, air, cooling water...),

  - cascading effects among the equipment (e.g., if the realignment of an equipment item in one procedure such as an air-operated valve would implicitly require the subsequent realignment of another equipment item such as isolation of an air line that would then disable a portion of the air system), and

  - instrumentation (e.g., indicators, alarms, sensors, logic devices...) and controls (e.g., hand switches...) that (a) affect automatic operation of the above primary and support system equipment and/or (b) *at least singularly* are relied upon (as opposed to multiple, redundant items) to credit post-initiator human actions to be included in the model (e.g., a single subcooling indication relied upon to meet an emergency core cooling termination criteria which if miscalibrated could induce failure of the appropriate post-initiator operator action).

4.1.3.2 Good Practice #2:

The identification process should identify pre-initiator human actions even if they may be potentially covered by the affected equipment failure data (see section 4.1.4 for additional information).

4.1.3.3 Good Practice #3:

If applicable and credited in the analysis, the identification process should address other operational modes and routine actions affecting barriers and other structures such as fire doors, block walls, drains, seismic restraints, etc.

4.1.3.4 Good Practice #4:

The identification process needs to include possible pre-initiator actions *at least within each system* where redundant or multiple diverse equipment can be affected by (a) a single act (e.g., misalignment of a valve affecting multiple system trains or even multiple systems) or (b) through a common failure with similar multiple acts (e.g., mis-calibrating multiple sensors due to incorrect implementation of the same calibration procedure or use of the same mis-calibrated standard). For the latter case, the analyst should not duplicate that already covered under the common cause failure modeling of the equipment, but should include consideration of possible commonalities such as:

- same crew, same shift performing the actions (common "who" mechanism),
- common incorrect calibration source (common "what" mechanism),
- common incorrect tool, process, or procedure/training, or inadequate material (e.g., wrong grease) (common "what/how" mechanisms), and
- close proximity in time and/or space/location of similar multiple acts (common "when/where" mechanisms).

The more these commonalities co-exist, the more the identification process should consider the act as a potentially important pre-initiator action to be included.

4.1.4 POSSIBLE IMPACTS OF NOT PERFORMING GOOD PRACTICES AND ADDITIONAL REMARKS:

Besides the obvious issues associated with incompleteness and inaccuracy and thereby potentially missing a risk-significant pre-initiator action, the following observations are noted.

- Missing or unnecessarily including an action is often not a serious mistake (i.e., would not significantly affect the overall risk) unless the action can affect multiple equipment items. This is because with common nuclear plant practices and designs, typically those actions that could affect multiple trains of equipment tend to be the more significant pre-initiator human failures. Those affecting just one equipment item are usually not important unless the equipment item has a high operating reliability (e.g., failure to start or run is in the 1E-4 or lower probability range) and so the pre-initiator failure probability could be a significant contributor to the unavailability of the equipment.
- One should include the possible failures associated with routine test and maintenance or calibration procedures that could affect critical instrumentation, diagnostic devices, or specific items like pushbuttons, etc. that have no redundancy or diverse means of function. While typically such situations do not exist in nuclear power plants, changes to the plant could conceivably and unintentionally create such a situation. Affecting the operator's ability to take the desired action is similar, functionally, to affecting the equipment item itself which is to be activated. Hence, it at least should be ensured that such situations, from a possible pre-initiator perspective, do not exist or if they do, they are addressed.

- In practice, it is best to include pre-initiator actions even if the associated failure may already be included in the failure data for the affected equipment item (e.g., in the failure-to-start data). This is because it is often hard to determine if the failure data bases include such human failures since data bases are typically insufficiently documented to know if the potential pre-initiator failure is already included. Generally, unless the failure can affect multiple equipment, such failings tend to not be important since missing them or double-counting them tend not to be serious PRA problems. Potential double-counting is the most conservative thing to do, and yet typically not a serious over-estimation of the failure's significance. In addition, including all identified pre-initiators gives analysts the opportunity to identify potentially problematic actions such as those with procedural or training problems, those that do not require appropriate checks, etc.

- If applicable, one should include the possible failures associated with routine test and maintenance or calibration procedures that could affect equipment critical to external events such as fire barriers (e.g., opening a fire door and failing to restore its closed position), seismic restraints, floor drains and barriers, wind barriers, etc. While typically such situations do not exist in nuclear power plants since such equipment items often do not have routine test, maintenance, or calibration activities that would adversely affect their function, changes to the plant or plant practices, for instance, could conceivably and unintentionally create such a situation. To the extent the analysis assumes the functionality of these normally highly reliable devices, pre-initiator failures that could affect these devices could be potentially important. Hence, it at least should be ensured that such situations, from a possible pre-initiator perspective, do not exist or if they do, they are addressed.

- Considering the potential importance of acts that affect multiple equipment, the identification process should search for acts that affect multiple equipment items *at least within a system* (e.g., auxiliary feedwater system, reactor core injection system...) as this represents the current state of the art in PRA. A search across multiple systems (e.g., auxiliary feedwater and high pressure injection) is an expansion of the current state of the art and should not be expected except for those cases where the same instrumentation or equipment (e.g., pressure signals, same tank level equipment) activates or affects multiple systems.

## 4.2 Screening those activities for which human failure events do not need to be modeled

4.2.1 OBJECTIVE: To screen out those activities for which associated failures do not need to be analyzed because they should be probabilistically unimportant. The screening process, though largely qualitative, is based on the belief that certain design or operational practices make some pre-initiator failures sufficiently unlikely that they will not be risk significant failures and therefore do not need to be modeled. The following provides good practices for screening out pre-initiator human actions and associated human failures while implementing the related Standard requirements.

4.2.2 CORRESPONDING ASME STANDARD REQUIREMENTS:

The Standard addresses allowable screening of activities based on practices that limit the likelihood of errors in those activities. There are multiple

supporting requirements in the Standard that address screening rules or criteria, as well as the requirement to not screen actions that could affect multiple equipment.

4.2.3 GOOD PRACTICES:

4.2.3.1 Good Practice #1:

A candidate pre-initiator action can be screened out (i.e., not to be modeled) if the nature of the associated action meets any of the following criteria and the reason for screening is documented (see exception under Good Practice #2 below):

- the affected equipment will receive an automatic realignment signal and it can respond (i.e., is not disabled) if demanded, or
- there is a valid post-maintenance/test functional check after the original manipulation which will reveal misalignment or incorrect status (e.g., faulty position, improper calibration), or
- following the original action(s), an independent second verification of equipment status using a written checklist that will verify incorrect status is performed, or
- a valid check, at least once per shift, of equipment status that will reveal misalignment or incorrect status, is used, or
- there is a compelling signal (e.g., annunciator or indication) of improper equipment status or inoperability in the control room, it is checked at least shiftly or daily, and realignment can be easily accomplished, or
- other criteria as long as it can be demonstrated that the resulting human error probabilities would be low compared with the failure probabilities (e.g., failure to open) of the equipment.

4.2.3.2 Good Practice #2:

Do not screen out those actions and possible pre-initiator failures that simultaneously affect multiple (redundant or diverse) equipment items (see Good Practice #4 under Section 4.1.3).

4.2.3.3 Good Practice #3 (application-specific):

For a specific PRA application and depending on the issue being addressed (e.g., examination of a specific procedure change), revisit the original PRA screening process to ensure issue-relevant human actions have not been deleted from the PRA prior to its use to assess the new issue.

4.2.4 POSSIBLE IMPACTS OF NOT PERFORMING GOOD PRACTICES AND ADDITIONAL REMARKS:

Besides the obvious issues associated with inappropriate screening and thereby potentially missing a risk-significant pre-initiator action, the following observations are noted.

- Generally, screening out pre-initiator failures (i.e., don't have to be modeled) is acceptable based on experience with past PRAs and the types of pre-initiator failures that are typically found to be unimportant. This is done to simplify the model and not expend resources addressing unimportant pre-initiator actions. It should be clear that an appropriate level of investigation has been performed to ensure the above criteria have been met and if these or other criteria are used, their justification is documented for outside review. It is advisable that a record of all screened actions be kept for later reference when performing specific applications (see Good Practice #3). When in doubt, it is recommended the pre-initiator action not be screened out but the corresponding failure modeled in the PRA for further analysis.

- Since pre-initiator actions and related failures affecting multiple equipment items can sometimes be risk important, none of these should be screened out but should be modeled and examined in more detail in the PRA because of the potential consequences of the failure.

- There can be a tendency to want to use an existing PRA model to address issues such as changes to the plant, without spending the appropriate time to revisit some of the underlying assumptions and modeling choices made to create the original PRA. Such a review should be done to see if these assumptions and choices still apply for the issue being addressed. In this case, some pre-initiator failures may not have been included in the original PRA (i.e., screened out) that in light of the new issue being addressed, should now be included in the model (i.e., could be important for addressing the issue). Hence it is good practice to implement a process that ensures that some of the formerly screened out pre-initiator failures do not have to be added back-in to the model in order to appropriately address the issue.

## 4.3 Modeling specific human failure events (HFEs) corresponding to the human failures

4.3.1 OBJECTIVE: To define how the specific pre-initiator HFE is to be modeled in the PRA to accurately represent the failure of each action identified and not screened out from the above analysis activities. The HFE needs to be linked to the affected equipment (single or multiple) and needs to appropriately define the failure mode of that equipment that makes the equipment unavailable. The following provides good practices for modeling pre-initiator human failure events while implementing the related Standard requirements.

4.3.2 CORRESPONDING ASME STANDARD REQUIREMENTS:

The Standard calls for the modeling of pre-initiator HFEs based on the impact of the failure in the PRA. There are multiple supporting requirements in the Standard that address the modeling level of detail for each HFE and the modes of failure to be considered.

### 4.3.3 GOOD PRACTICES:

#### 4.3.3.1 Good Practice #1:

Define each specific pre-initiator HFE to be modeled in the PRA as a basic event that describes the human-induced failure mode and is located in the model such that it is linked to the unavailability of the affected component, train, system, or overall function (i.e., level of modeling) depending on the effect(s) of the HFE (e.g., a single valve will not close, a train will be isolated, the automatic start signal for an entire system will be disabled). The following attributes, as a minimum, should be used to define the pre-initiator failure level properly in the PRA:

- the nature of the manipulation affects a whole train, system, etc. so it makes more sense to define the HFE at that level,

- multiple individual acts affecting multiple equipment (e.g., different components) can be combined as a single pre-initiator HFE affecting a higher level of equipment resolution (e.g., the train containing the different components) as long as (a) the acts and effects are related, (b) how the single HFE will be quantified (i.e., the performance-shaping factors that would affect quantification as discussed later) is not significantly different or will be conservatively bounding than if the individual acts were to be modeled and quantified separately, and (c) there are no potential commonalities/dependencies with other pre-initiator acts elsewhere in the model so that potential common failures among similar individual acts might be missed (e.g., miscalibration of multiple signal channels), and

- consideration of the level of detail already modeled in the PRA (e.g., train, system) for failures of the associated equipment (less important factor).

The failure modes (fail to close, fail to start, etc.) should be a direct result of considering the equipment affected and the effects of the human-induced failure (refer to all the Good Practices under Section 4.1.3) and stem from failure to restore equipment and/or otherwise correct the adverse effect (such as miscalibration) so that the equipment is again operable. The failure modes should clearly describe the HFE effect to ensure proper interpretation of the HFE in the model (e.g., only two of three redundant sensors need to be disabled to make the actuation signal unavailable, and not all three sensors have to be disabled).

As an aid to ensure appropriate modeling, it is recommended practice (but not necessary) that the pre-initiator failure be placed in close proximity, in the PRA model, to the equipment affected by the human failure. In this way, a quick comparison can be made between the equipment failure and the pre-initiator human failure to ensure they are consistent.

### 4.3.4 POSSIBLE IMPACTS OF NOT PERFORMING GOOD PRACTICES AND ADDITIONAL REMARKS:

The precise definition of the pre-initiator basic events and their placement in the model (from both a logic and failure mode standpoints) ultimately define how the model addresses the effects of the human failures. This needs to be done accurately if the model is going to logically represent the

real effects of each human failure and if the corresponding HFE is going to be correctly quantified (as discussed later).

## 4.4 Quantifying the corresponding human error probabilities (HEPs) for the specific HFEs

4.4.1 OBJECTIVE: To address how the human error probabilities (HEPs) for the modeled HFEs from the previous analysis activity are to be quantified. This section provides good practices guidance on an attribute or criteria level and does not endorse a specific tool or technique (although THERP[3] or its ASEP[4] simplification are among those often used). Ultimately, it is these probabilities along with the other equipment failure and post-initiator human error probabilities as well as initiating event frequencies that are all combined to determine such risk metrics as CDF, LERF, $\triangle$CDF, $\triangle$LERF, etc. as addressed in Regulatory Guide 1.174[11]. The following provides good practices for quantifying pre-initiator human failure events while implementing the related Standard requirements.

4.4.2 CORRESPONDING ASME STANDARD REQUIREMENTS:

The Standard calls for a systematic process for assessing the pre-initiator HEPs that addresses plant-specific and activity-specific influences. There are multiple supporting requirements in the Standard that address many factors associated with quantifying the HEPs. These include when screening vs. detailed estimates are appropriate, performance-shaping factors considered in the evaluations, treatment of recovery, consideration of dependencies among HFEs, uncertainty, and reasonableness of the HRA results.

4.4.3 GOOD PRACTICES:

4.4.3.1 Good Practice #1:

The use of screening-level human error probability (HEP) estimates is virtually necessary during the early stages of PRA development and quantification. This is acceptable (and almost necessary since not all the potential dependencies among human events can be pre-known) provided (a) it is clear that the individual values used are over-estimations of the probabilities if detailed assessments were to be performed AND (b) dependencies among multiple human failure events appearing in an accident sequence are conservatively accounted for. These screening values should be set so as to be able to make the PRA quantification process more efficient (by not having to perform detailed analysis on every human failure event), but not so low that later detailed analysis would actually result in higher HEPs. The screening estimates should consider both individual HEPs and the potential for multiple and possibly dependent human failure events for a given accident sequence (scenario). To meet these conditions, it is recommended that (unless a more detailed assessment is performed of the individual or combination events to justify lower values):

- no individual pre-initiator HEP screening value should be lower than 1E-2 (typical of highest pre-initiator values in PRAs), and

- multiple HEPs in the same sequence should not have a collective value lower than 5E-3 (accounts for a 0.5 high dependency factor) at this stage.

4.4.3.2 Good Practice #2:

As needed for the issue being addressed to produce a more realistic assessment of risk, detailed assessments (not just screening estimates) of at least the significant human failure event contributors should be performed. The PRA analyst can define the significant contributors by use of typical PRA criteria (not addressed here) such as importance measure thresholds as well as other qualitative and quantitative considerations. While the use of screening-level values (supposedly purposely conservative) may, at first, seem to be a "safe" analysis process, it can have negative impacts. Screening values can focus the risk on inappropriate human actions or related accident sequences and equipment failures because of the intentionally high HEPs. Such incorrect conclusions need to be avoided by ensuring a sufficient set of more realistic, detailed HEPs are included in the model.

4.4.3.3 Good Practice #3 (application-specific):

For a specific PRA application and depending on the issue being addressed (e.g., examination of a specific procedure change), revisit the use of screening vs. detail-assessed HEPs to ensure issue-relevant human actions have not been prematurely deleted from the PRA or there is an inappropriate use of screening vs. detailed values to properly assess the issue and the corresponding risk.

4.4.3.4 Good Practice #4:

HEP assessments should account for the most relevant plant-specific and activity-specific performance-shaping factors in the analysis of each pre-initiator HFE. There is not one consensus list of appropriate contextual factors (e.g., plant conditions, PSFs, activity characteristics, etc.) to be considered in the evaluation of the pre-initiator HEPs. Additionally, for a specific action, what factors are most relevant may be different (e.g., perhaps one act is time-sensitive because it is done in a high radiation area while another is most affected by the complexity of steps with many opportunities to make undetected mistakes). It should be qualitatively apparent that the factors seemingly most relevant to the act (based on an understanding of the act) have been considered in the corresponding HEP estimate.

Factors that are typically important to address because they tend to be variable and not almost always optimal based on typical nuclear plant practices, include:

- whether written work plans, job briefs, and related procedures (positive influences tending to lower the HEP), or verbal guidance and/or memory (more negative influences tending to raise the HEP) are used, as well as the quality of the information (e.g., look for ambiguities, incompleteness, inconsistencies, etc. that are negative influences and thus tend to raise the HEP),

- complexity (e.g., multiple and/or repetitive steps that are hard to track, use (or not) of checklists, several variables involved and calculations required...), and

- ergonomic issues (e.g., layout, available information [instruments, alarms, computer readouts, etc.], labeling, readability, highly physical...).

Other factors that tend to not be as important either because of typical nuclear plant practices or because the factors are typically less relevant include (it should still be ensured that the typical practice or irrelevancy is not compromised):

- skill level/experience/training of crew (typically adequate in nuclear plants for the jobs each crew member is to perform),

- stress level (not usually relevant in pre-initiator failures unless special situations such as potential personal harm, the need for fast sequential responses, etc. play a role),

- environmental factors such as temperature, humidity, radiation, noise, lighting, etc. (typically the environment is sufficiently benign except for special circumstances such as a high radiation environment and thus the desire to hurry the actions), and

- availability of time (not usually a strong factor in pre-initiator failures).

If the large majority of these factors affect the human performance negatively or if even just one or two is an overwhelming negative influence, the HEP will tend to be higher (e.g., 0.01 to 0.1 or even higher, not accounting for recovery addressed under Good Practice #5 below). Conversely, mostly positive influences should yield lower HEPs (e.g., <1E-3, with additional recovery factors still to be applied as addressed under Good Practice #5 below).

4.4.3.5 Good Practice #5:

Applicable recoveries applied to the HEP evaluations for the HFEs being analyzed should be used (multiple recoveries may be acceptable) where appropriate, but any dependencies among the initial failure and the recoveries, and among the recoveries themselves, must be considered (see Good Practice #6 below). Typical considerations in applying recovery include:

- post-maintenance or post-calibration tests are required and performed by procedure,

- independent verification, using a written check-off list, which verifies component status following maintenance/testing/calibration is used, and its practice has been verified by walk-throughs and examination of plant experience,

- the original performer, using a written check-off list, makes a separate check of component status at a later time,

- work shift or daily checks are performed of component status, using a written check-off list,

- there is a compelling feedback (e.g., alarm) that will enhance the original failure being detected and can be quickly recovered, or

- combinations of the above.

The more of these are applicable for a given pre-initiator HFE, the more the situation tends to increase the recovery potential (i.e., decrease the HEP) since each recovery, to the extent they are independent, result in a multiplier (e.g., 0.1) on the original HEP estimate thereby reducing its overall value.

Basic HEPs for pre-initiator HFEs for nuclear plant applications (including recovery) are typically expected in the 0.01 (among the highest) to 0.0001 range. Any values below the 0.0001 to 0.00001 range should be considered suspect unless justified.

4.4.3.6 Good Practice #6:

Dependencies among the pre-initiator HFEs and hence the corresponding HEPs in an accident sequence should be quantitatively accounted for in the PRA model. This is particularly important so that combined probabilities are not inadvertently too optimistic, resulting in the inappropriate decrease in the risk significance of human actions and related accident sequences and equipment failures. In the extreme, this could result in the inappropriate screening out of accident sequences from the model because the combined probability of occurrence of the events making up an accident sequence drops below a threshold value used in the PRA to drop sequences from the final risk results.

To address these dependencies, usually a level or degree of dependence among the HFEs in an accident sequence is determined, at first qualitatively (e.g., low, high, complete), and then combined HEPs are assessed accordingly. Once the first HEP has been estimated, subsequent quantitative factors for dependent human failures or recoveries of the original failure are typically expected to be:

- 0.01 to 0.1 for low dependence
- 0.1 to 0.5 for high dependence
- >0.5 for very high or 1.0 for complete dependence

Note that specific tools/techniques may use somewhat different probabilities than provided here based on specific considerations.

In establishing the level of dependence, Good Practice #4 under Section 4.1.3 addresses typical commonalities that tend to make HEPs more dependent (i.e., an HFE is not independent of another HFE and so once the first human failure occurs, there is a high likelihood that a similar second or third, etc. human failure will also occur such as the failure to restore the lineup of one train of equipment after a test and then failing to similarly restore the second train of equipment after a similar test). Good Practice #5 just above addresses recovery characteristics that tend to break-up these commonalities because they "recover" any initial error, making the individual HFEs more independent. The more the types of commonalities addressed under Good Practice #4 under Section 4.1.3 exist and the less corresponding recoveries under Good Practice #5 above exist, the higher should be the assessed level of dependence among the HFEs. To the extent the converse is true, low or even no dependence should be assessed.

4.4.3.7 Good Practice #7:

Point estimates should be mean values for each HEP (excluding screening HEPs) and an assessment of the uncertainty in the mean values should be performed at least for the dominant HEPs to the extent that these uncertainties need to be understood and addressed in order to make appropriate risk-related decisions. Assessments of uncertainty are typically performed by:

- assigning uncertainty distributions for the HEPs and propagating them thru the quantitative analysis of the entire PRA such as by a Monte Carlo technique, and/or

- performing sensitivity analyses that demonstrate the effects on the risk results for extreme estimates in the HEPs based on at least the expected uncertainty range about the mean value.

Note, in some cases, it may be sufficient to address the uncertainties by just qualitative arguments without the need to specifically quantify them (e.g., justifying why the HEP cannot be very uncertain or why a change in the HEP has little relevancy to the risk-related decision to be made).

In assessing the uncertainties, and particularly when assigning specific uncertainty distributions, the uncertainties should include (a) those epistemic uncertainties because of lack of knowledge of the true expected performance of the human for a given context and associated set of performance-shaping factors, and (b) consideration of the combined effect of the relevant aleatory (i.e., random) factors to the extent they are not specifically modeled in the PRA and to the extent that they could alter the context and performance-shaping factors for the HFE. For pre-initiator HFEs, there should be few or no aleatory factors worthy of consideration, since typically the procedure used, the environment experienced, etc. do not randomly change. But, for example, if different and significant crew experience levels are known to exist, it is random as to which crew will perform the pre-initiator act at any given time. In such a case, the mean should represent the average crew experience level and the uncertainty should reflect the possible range in those levels. Again, aleatory factors are typically not very relevant to pre-initiator HEPs and so typically are not important to address.

Whatever uncertainty distributions are used, the shape of the distributions (e.g., log-normal, normal, beta...) are typically unimportant to the overall risk results (i.e., the results are usually not sensitive to specific distributions). Further, typical uncertainties include values for the HEP that represent a factor of 10 to 100 between the lower bound value and the upper bound value that encompass the mean value.

4.4.3.8 Good Practice #8:

The pre-initiator HEPs (excluding the screening HEPs) should be reasonable from two standpoints:

- first and foremost, relative to each other (i.e., the probabilistic ranking of the failures when compared one to another), and

- in absolute terms (i.e., each HEP value) to the extent that the sensitivity of the risk-related decision is not important as to the absolute values for the HEPs.

This reasonableness should be checked based on consideration of actual plant experience and history, against other evaluations (such as for similar acts at other plants), and the qualitative understanding of the actions and the relevant contexts and performance-shaping factors under which the acts are performed. It is suggested that a rank-ordered list of the pre-initiator HFEs by probability be used as an aid for checking reasonableness. For example, simple, procedure-guided, independently checked actions should have lower HEPs than complex, memorized, not checked actions, all other factors being the same. Typical expectations of pre-initiator HEPs can be wide-spread (~0.01 to 0.0001) and depend particularly on the relevant contextual factors, applicable recoveries, and proper consideration of dependencies as discussed under many of the Good Practices covered above.

### 4.4.4 POSSIBLE IMPACTS OF NOT PERFORMING GOOD PRACTICES AND ADDITIONAL REMARKS:

Besides the obvious concerns about inaccuracies in the HEP quantification and thus whether the HEPs "make sense", as well as the resulting potential misinformation about the dominant risk contributors if quantification is not done well, the following observations are noted.

- Screening is a useful and most often, necessary part of HRA so as to avoid the expenditure of resources on unimportant human events and accident sequences. The above guidance is aimed at allowing a level of useful screening without inadvertently and inappropriately allowing the analytical phenomenon of, for instance, multiplying three human events in the same sequence each at a screening value of 1E-2 to yield a 1E-6 combined probability, without checking for dependencies among the human events. In such a case some human failure events and combinations of events, or even whole accident sequences, may inappropriately screen out of the PRA model entirely because the accident sequence frequency drops below a model threshold. Hence some of the dominant individual or combination contributors may be missed. This is why the screening values both individually and for combined events should not be too low during the screening stage. Further, if screening values are left permanently assigned to some human failure events that should be assessed with more detail to obtain a more realistic assessment of risk (supposedly lowering the probability), the risk significance of these human failure events and related equipment failures are likely to be over-emphasized at the expense of improperly lessening the relative importance of other events and failures.

- It is important to be sure that dependencies among the various modeled HFEs including the associated recoveries, have been investigated (e.g., the same person as the originator of the action performing the recovery may be more prone to fail to detect the original failure than an independent checker). Treating HFEs and any corresponding recoveries as independent acts without checking for dependencies (thereby being able to multiply the individual HEPs) can inappropriately lessen the risk significance of those HFEs and related equipment failures in accident sequences. This can cause the inappropriate dropping out of accident sequences because the sequences quantitatively drop below a model threshold value as discussed above under screening. Proper consideration of the dependencies among the human actions in the model is necessary to reach the best possible evaluation of both the relative and absolute importance of the human events and related accident sequence equipment failures.

- The use of mean values and addressing uncertainties are a part of the Regulatory Guide 1.174[11] guidance and to the extent addressed therein, the HRA quantification needs to be consistent with that guidance when making risk-informed decisions.

- There can be a tendency for analysts to want to use an existing PRA model to address issues such as changes to the plant, without spending the appropriate time to revisit some of the underlying assumptions and modeling choices made to create the original PRA. A review should be done to see if these assumptions and choices still apply for the issue being addressed. In this case, some pre-initiator human failure events may be quantified in the original model using a set of screening estimates and detailed failure probabilities that may not be appropriate for the new issue being addressed. As an example, where higher screening values may have been acceptable for purposes of the original PRA, these supposedly conservative values may over-estimate the contribution of these human failure events for the issue being addressed. Further, the relative risk contribution of equipment and associated accident sequences with which the human failure events appear, may be artificially too high (and therefore other events too low) because of the screening values. Hence it is good practice to revisit the use of screening and detailed human failure event probabilities in order to appropriately address the issue.

## V. POST-INITIATOR HRA

The ASME Standard[5] separates its requirements into two broad classifications; those that address the modeling of failures of pre-initiator human actions and those that address the modeling of failures of post-initiator human actions. This section provides good practices for implementing the requirements for addressing post-initiator human failure events (HFEs) in a PRA.

Post-initiator human failure events are events that represent the impact of human failures committed during actions performed in response to the initiation of an accident sequence (e.g., while following post-trip procedures or performing other recovery actions). They are important to model because humans can have a direct influence on the mitigation or exacerbation of undesired plant conditions after the initial plant upset. Hence, depending on the issue being addressed, this impact may need to be included in a PRA if a realistic assessment of risk is required.

The following good practices are categorized under four major analysis activities for doing post-initiator HRA. These analysis activities include:

1. Identifying potential post-initiator human failures
2. Modeling specific human failure events (HFEs) corresponding to the human failures
3. Quantifying the corresponding human error probabilities (HEPs) for the specific HFEs
4. Adding recovery actions to the PRA.

### 5.1 Identifying potential post-initiator human failures

5.1.1 OBJECTIVE: To identify the key human response actions that may need to be taken by the operators in response to a variety of possible accident sequence s and that will therefore need to be modeled in the PRA. This is important since failures associated with these actions (e.g., failure to start standby liquid control, failure to initiate feed and bleed, failure

to properly control steam generator feed flow, failure to align containment/suppression pool cooling) are represented in the PRA such that in combination with equipment failures, are expected to lead to core damage and/or large early releases. Such failures contribute to the overall risk and thus a systematic process needs to be followed to identify these response actions. The following provides good practices for identifying post-initiator human failures while implementing the related Standard requirements.

5.1.2 CORRESPONDING ASME STANDARD REQUIREMENTS:

The Standard calls for a systematic review to identify operator responses required for each of the accident sequences. There are multiple supporting requirements in the Standard that address what to review as well as the types of actions to be included. Use of talk throughs and simulator observations are also addressed as part of the supporting requirements.

5.1.3 GOOD PRACTICES:

5.1.3.1 Good Practice #1:

Reviews of the following form the primary bases for identifying the post-initiator actions.

- Review plant-specific emergency operating procedures (EOPs), abnormal operating procedures (AOPs), annunciator procedures, system operating procedures, severe accident management guidelines (SAMGs), and other special procedures (e.g., fire emergency procedures) as appropriate. The review is done to identify ways operators are intended to interact with the plant equipment after an initiator as a function of the various conditions that can occur as defined by the development of the PRA accident sequences and equipment unavailabilities and failure modes. Particularly note where operator actions are called out in these procedures and under what plant conditions and indications (cues) such actions are carried out. It will also be useful at this time to examine whether there are any potential accident conditions under which the procedures might not match the situation as well as would be desired, e.g., potentially ambiguous decision points or incorrect guidance provided under some conditions. Information about such potential vulnerabilities will be useful later during quantification and may help identify actions that need to be modeled.

While not necessary at this stage of the analysis (probably more beneficial during the modeling and quantification phases, but could be started at this stage on a selective basis of likely importance), the results of the following additional reviews may add to the list of actions and/or help interpret how procedural actions should be defined based on how they are actually carried out.

- Review of training material including, where possible, talk-throughs or walkdowns of the actions with operations or training staff to ensure consistency with training policies and teachings, and to identify likely operator response tendencies for various conditions that may not be evident in the procedures (although it is not the intent to perform numerous or detailed talk-throughs, walkdowns, or simulations at this phase of the analysis - the use of these techniques is more relevant later under the HFE modeling and quantification phases). For example, operators may cite a reluctance to restart reactor coolant pumps in spite of the

procedure direction based on their training and perceived adverse effects, or they may have a preference to use condensate as a BWR injection source before using lower pressure emergency core cooling system. These added "interpretations" of the procedures can help complete and/or clarify the identified actions and ensure that later modeling and quantification of the actions will reflect the "as-operated" plant.

- Observations of simulated accidents since these can provide valuable insights with regard to how the actions are actually carried out, by whom, and particularly how procedure steps are interpreted by plant crews especially where the procedure is ambiguous or leaves room for flexibility in the crew response (although it is not the intent to perform numerous or detailed talk-throughs, walkdowns, or simulations at this phase of the analysis - the use of these techniques is more relevant later under the HFE modeling and quantification phases). For example, through simulation it may be observed that a "single action" in the procedure (e.g., align recirculation) is actually carried out by a series of numerous and sequential individual actions (e.g., involving the use of many handswitches in a certain sequence). Again these observed "interpretations" of the procedures can help complete and/or clarify the identified actions and ensure that later modeling and quantification of the actions will reflect the "as-operated" plant.

5.1.3.2 Good Practice #2:

The review process should involve the following:

- Knowledge of the functions and associated systems and equipment to be modeled in the PRA for both CDF and LERF.

- Identifying whether the function is needed (e.g., injection) or undesired (e.g., stuck-open safety relief valve) recognizing these may vary with different initiators and sequences.

- Identifying the systems/equipment that can contribute to performing the function or cause the undesired condition including structures and barriers where appropriate (e.g., fire door, floor drains) especially for external event analyses.

- Identifying ways the equipment can functionally succeed (i.e., the success criteria) and fail.

- Based on the above, identifying ways the operators are (a) intended/required to interact with the equipment credited to perform the functions modeled for the accident sequences modeled in the PRA and/or (b) to respond to equipment and failure modes that can cause undesired conditions per the PRA. During the identification process, it is helpful to use action words such as actuate, initiate, isolate, terminate, control, change, etc. so that the desired actions are clear.

5.1.3.3 Good Practice #3:

While the specific actions to be identified may be plant-specific, in general, the following types of actions are expected to be identified. Note that actions that are heroic (e.g., must enter an extreme high radiation environment) or without any procedure guidance or not trained on, should not be

included or credited in the analysis (exceptions may be able to be justified, but this should not be normal practice).

- Include necessary and desired/expected actions (e.g., initiate RHR, control vessel level, isolate a faulted steam generator, attempt to reclose a stuck-open relief valve).

- Include backup actions to failed automatic responses (e.g., manually start a diesel generator that should have auto started) but be sure the action can be credited to recover the auto failure mode.

- Include anticipated procedure-guided or skill-of-the-craft recovery actions (e.g., restore offsite power, align firewater backup) although these may best be defined later as the PRA quantification begins and important possible recovery actions become more apparent.

Consistent with present day state-of-the art, acts whose failure involve an error of omission (EOO) should be included when identifying post-initiator acts of concern. These involve failure to take the appropriate actions as called out in the procedures and/or trained on or expected as common practice. For example, failure to initiate feed and bleed or failure to start standby liquid control, are EOOs. Possible acts whose failure would involve an error of commission (EOC) are generally beyond current PRA practice. These involve performing expected acts incorrectly or performing extraneous and detrimental acts such as shutting down safety injection when it is not appropriate. These are not necessarily expected to be identified but see Section 7 of this document for more on this subject.

Finally, it should be recognized that iterations as well as refinement and review of the PRA model construction may (and often do) provide additional opportunities to identify any potentially important missed actions as the PRA model evolves.

5.1.4 POSSIBLE IMPACTS OF NOT PERFORMING GOOD PRACTICES AND ADDITIONAL REMARKS:

While not all the post-initiator actions will be important in the final assessment of risk, unlike the pre-initiator actions, it is difficult to predetermine (at this stage) a set of actions that do not have to be included as part of the identification process. Ways the operators interact with the plant and affect the outcome of any accident sequence need to be assessed in order to determine their relative significance. Hence the good practices herein are aimed at ensuring potentially risk-significant post-initiator actions (based on the procedures as well as the ways the procedures are interpreted and carried out) are identified at this stage of the analysis. Otherwise, the model could be incomplete and/or inaccurate, potentially resulting in misinformation as to the risk dominant plant features (including the important human actions).

## 5.2 Modeling specific human failure events (HFEs) corresponding to the human failures

5.2.1 OBJECTIVE: To define how each specific post-initiator HFE is to be modeled in the PRA to accurately represent the failure of each action identified. This involves the modeling of the HFEs as human-induced unavailabilities of functions, systems, or components consistent with the level of detail in the PRA accident sequences and system models, possible grouping of responses into one HFE, and ensuring the modeling reflects certain plant-specific and

accident sequence-specific considerations. The following provides good practices for modeling post-initiator human failure events while implementing the related Standard requirements.

5.2.2 CORRESPONDING ASME STANDARD REQUIREMENTS:

The Standard calls for the HFEs to be defined so that they represent the impact of not properly performing the required responses, consistent with the structure and level of detail of the accident sequences. There are multiple supporting requirements in the Standard that address the modeling level of detail for each HFE and how to complete the definition of each HFE.

5.2.3 GOOD PRACTICES:

5.2.3.1 Good Practice #1:

Define each specific post-initiator HFE to be modeled in the PRA as a basic event that describes the human failure of not properly performing the required response and is located in the model such that it is linked to the unavailability of the affected component, train, system, or overall function (i.e., level of modeling) depending on the effect(s) of the HFE (e.g., failure to manually depressurize using the safety relief valves, failure to manually scram, failure to align the backup train of service water). The following considerations should be used to define the post-initiator failure level properly in the PRA:

- the nature of the action is performed on a train, system, etc. level so it makes more sense to define the HFE at that level,

- the consequences of the failure and what would be affected by the failure (just a component is affected, a whole train, a system, multiple systems, an entire function),
- multiple individual acts/responses such as at a system or component level (e.g., starting high pressure injection and then subsequently opening a power-operated pressurizer relief valve) can be combined as a single post-initiator HFE affecting a higher level of equipment resolution such as at a system or a function level (e.g., initiating feed and bleed) as long as (a) the acts and effects are related, (b) how the single HFE will be quantified (i.e., the performance-shaping factors that would affect quantification as discussed later) is not significantly different or will be conservatively bounding than if the individual acts were to be modeled and quantified separately, and (c) there are no potential commonalities/dependencies with other post-initiator acts elsewhere in the model so that potential common failures among similar individual acts might be missed (see the discussion presented below),

- the level of detail already modeled in the PRA (train, system, etc.) for failures of the associated equipment (less important factor).

As an example of how human responses may be grouped and modeled as one or more HFEs, consider the case in a boiling water reactor (BWR) of a desired response to control reactivity in an anticipated transient without scram scenario. Failure to control reactivity could be defined as one

HFE, or as several HFEs based on the subtasks involving inhibiting the automatic depressurization system, lowering reactor water level, and initiating the standby liquid control system.

For situations such as the above example, if failure to perform the subtasks (a) have different effects, (b) may individually be impacted by very different performance-shaping factors (e.g., in-control room actions vs. local actions in a high steam environment area, a subtask performed early in the scenario vs. another subtask performed much later in the scenario), or (c) involves an action that has a dependency with some other action to be modeled in the PRA (e.g., failure to trip two reactor coolant pumps followed by subsequent failure to trip the remaining reactor coolant pumps when conditions warrant), the failures are best modeled as separate HFEs. An alternative is to model them all as one HFE and model the bounding consequence (such as the failure to control reactivity example cited above) as long as the most limiting performance-shaping factors are used (e.g., the shortest time that any of the subtasks must be performed, the most complex of the subtasks, etc.) and any subtask dependencies with other HFEs are identified, treated in the model, and properly quantified.

The failure effects as depicted in the PRA model should be a direct result of considering the equipment affected and the effects of the human-induced failure (refer to the Good Practices under Section 5.1.3) and stem from failure to properly perform the correct responses. The failures should sufficiently describe the HFE and its effect to ensure proper interpretation of the HFE in the model (e.g., fail to initiate feed and bleed within 5 minutes of the reactor pressure achieving 2400 psig).

As an aid to ensure appropriate modeling, it is recommended practice (but not necessary) that the post-initiator failure be placed in close proximity, in the PRA model, to the component, train, system, or function affected by the human failure. In this way, a quick examination of the model can reveal the modeled effect of the human failure.

5.2.3.2 Good Practice #2:

Each of the modeled post-initiator HFEs should be defined such that they are plant- and accident sequence-specific. Where helpful to fully understand the nature of the act(s) (e.g., who performs it, what is done, how long does it take, are there special tools needed, are there environmental issues or special physical needs, etc.), use of talk-throughs, walkdowns, field observations, and simulations are particularly encouraged.

In order for the act to occur, the operator must diagnose the need to take the act and then execute the act. While many performance-shaping factors are used to quantify the probability for failing to perform the act correctly (as discussed later under quantification), all of which should be evaluated based on plant and accident sequence-specifics, the following requirements are particularly germane to a basic understanding of the HFE and should be met to complete the definition of each HFE:

- to the extent possible, the time by which the act needs to be performed (e.g., fail to initiate feed and bleed by 2 minutes after primary pressure reaches 2400 psig), and the time necessary to diagnose the need for and to perform the act (1 minute) should be based on plant and accident sequence-specific timing and nature of the complexity and/or subtasks involved in implementing the act (i.e., not another plant analysis or a general analysis for the “average”

plant since the number and nature of the specific manipulations could be different, the plant thermal hydraulic response could be different, the location for local actions may require different travel times, some sequences require a fast response while others may require a much quicker response for the same act, etc.),

- similar to the above, the availability and timing of plant and accident sequence-specific cues (i.e., indications, alarms, visual observations, etc. and when they will be manifested) should be used as these can be different from plant-to-plant and different in a variety of accident sequences (e.g., such as a DC bus failure causing loss of some indications or alarms), and will affect the likelihood and timing of diagnosing the need for the action,

- plant-specific procedure and training guidance should be used based on the reviews under the Good Practices in Section 5.1.3,

- where the act is performed (e.g., in the control room, locally in the auxiliary building) should be noted, and

- the use of walkdowns, talk-throughs, and field or simulator observations are encouraged when defining the HFE as mentioned under Good Practice #1 under Section 5.1.3. See more about the benefits of these techniques in Appendix A.

5.2.4 POSSIBLE IMPACTS OF NOT PERFORMING GOOD PRACTICES AND ADDITIONAL REMARKS:

The precise definition of the post-initiator basic events and their placement in the model (from both a logic and failure mode standpoints) ultimately define how the model addresses the effects of the human failures. This needs to be done accurately if the model is going to logically represent the real effects of each human failure and if the corresponding HFE is going to be correctly quantified (as discussed later). This accuracy is best obtained if plant-specific and accident sequence-specific information is used. Nevertheless, the following observation is noted.

- Not using plant/accident sequence-specific thermal hydraulic information for timing may or may not be critical based on the relevancy and thus appropriateness of the non-specific (i.e., "general") timing information that is used. It is better to use plant and accident-specific information, though it is recognized that in some areas (e.g., containment response for LERF), from a practical standpoint, modified "general" information may be all that is readily available. Further, as long as the timing considerations used are reasonable and accurate to within the resolution of the HRA quantification tool to be used, differences between plant and accident-specific vs. more "general" timing considerations may not be a significant issue. Analysts should ensure that if non-specific timing information is used, it is reasonable to expect it to be appropriate for the plant and accident sequence being analyzed.

## 5.3 Quantifying the corresponding human error probabilities (HEPs) for the specific HFEs

5.3.1 OBJECTIVE: To address how the human error probabilities (HEPs) for the modeled HFEs from the previous analysis activity are to be quantified. This section provides good practices guidance on an attribute or criteria level and does not endorse a specific tool or

technique. Ultimately, it is these probabilities along with the other equipment failure and pre-initiator human error probabilities as well as initiating event frequencies that are all combined to determine such risk metrics as CDF, LERF, △CDF, △LERF, etc. as addressed in Regulatory Guide 1.174[11]. The following provides good practices for quantifying post-initiator HEPs while implementing the related Standard requirements.

5.3.2 CORRESPONDING ASME STANDARD REQUIREMENTS:

The Standard requires that a well-defined and self-consistent process be used to quantify the post-initiator HEPs. There are multiple supporting requirements in the Standard that address many factors associated with quantifying the HEPs. These include when conservative vs. detailed estimates are appropriate, consideration of cognitive and execution failures, performance-shaping factors considered in the evaluations, consideration of dependencies among HFEs, uncertainty, and reasonableness of the HRA results.

5.3.3 GOOD PRACTICES:

5.3.3.1 Good Practice #1:

Whether using conservative or detailed estimation of the post-initiator HEPs, the evaluation should include both cognitive (i.e., "thinking) as well as execution failures. For example, incorrectly interpreting a cue or not seeing a cue and thus not performing the act can be one mode of failure. Or, the operator can intend to take the act based on the proper and recognized cues but still otherwise fail to take the act or perform it correctly. Both need to be part of the HEP evaluations.

5.3.3.2 Good Practice #2:

The use of conservative human error probability (HEP) estimates is virtually necessary during the early stages of PRA development and quantification. This is acceptable (and almost necessary since not all the potential dependencies among human events can be pre-known) provided (a) it is clear that the individual values used are over-estimations of the probabilities if detailed assessments were to be performed AND (b) dependencies among multiple human failure events appearing in an accident sequence are conservatively accounted for. These conservative values should be set so as to be able to make the PRA quantification process more efficient (by not having to perform detailed analysis on every human failure event), but not so low that later detailed analysis would actually result in higher HEPs. The conservative estimates should consider both individual HEPs and the potential for multiple and possibly dependent human failure events for a given accident sequence (scenario). To meet these conditions, it is recommended that (unless a more detailed assessment is performed of the individual or combination events to justify lower values):

- no individual post-initiator HEP conservative value should be lower than the worse case anticipated detailed value and generally not lower than 0.1 (typical of high post-initiator values in PRAs), and

- multiple HEPs in the same sequence should not have a joint probability value lower than the worse case anticipated detailed joint probability value and generally not lower than 5E-2 (accounts for a 0.5 high dependency factor) at this stage.

5.3.3.3 Good Practice #3:

As needed for the issue being addressed to produce a more realistic assessment of risk, detailed assessments (not just conservative estimates) of at least the dominant human failure event contributors should be performed. The PRA analyst can define the dominant contributors by use of typical PRA criteria (not addressed here) such as importance measure thresholds as well as other qualitative and quantitative considerations. While the use of conservative values may, at first, seem to be a "safe" analysis process, it can have negative impacts. Conservative values can focus the risk on inappropriate human actions or related accident sequences and equipment failures because of the intentionally high HEPs. Such incorrect conclusions need to be avoided by ensuring a sufficient set of more realistic, detailed HEPs are included in the model.

5.3.3.4 Good Practice #4 (application-specific):

For a specific PRA application and depending on the issue being addressed (e.g., examination of a specific procedure change), revisit the use of conservative vs. detail-assessed HEPs to ensure issue-relevant human actions have not been prematurely deleted from the PRA or there is an inappropriate use of conservative vs. detailed values to properly assess the issue and the corresponding risk.

5.3.3.5 Good Practice #5:

As "good practice," the following table of performance-shaping factors (Table 5-1) for both in-control room and ex-control room (local) actions should be treated in the evaluation of each HEP per the table guidance. The guidance should fit most cases, but it should be recognized that for specific actions, some of the factors may not apply while others may be so important, the others do not matter (e.g., time available is so short, the act almost assuredly cannot be done regardless of the other factors). Further, if a specific situation warrants treatment of unique factors that are not, and cannot be

**Table 5-1 Post-Initiator PSFs To Be Considered**

| In-Control Room Actions | Ex-Control Room Actions |
|---|---|
| **Always Consider the Following PSFs** | **Always Consider the Following PSFs** |
| Applicability and suitability of training and experience | Applicability and suitability of training and experience |
| Suitability of relevant procedures and administrative controls | Suitability of relevant procedures and administrative controls |
| Availability and clarity of instrumentation (cues to take actions as well as confirm expected plant response | Availability and clarity of instrumentation (cues to take actions as well as confirm expected plant response |
| Time available and time required to complete the act, including the impact of concurrent and competing activities | Time available and time required to complete the act, including the impact of concurrent and competing activities |

| In-Control Room Actions | | Ex-Control Room Actions | |
|---|---|---|---|
| Complexity of required response along with workload, time pressure, the need for special sequencing, and familiarity | | Complexity of required response along with workload, time pressure, the need for special sequencing, and familiarity | |
| Team/crew dynamics and crew characteristics (degree of independence among individuals, operator attitudes - biases - rules, use of status checks, approach for implementing procedures, e.g., aggressive vs. slow and methodical...) | | | |
| Consideration of 'realistic' accident sequence diversions and deviations (e.g., extraneous alarms, failed instruments, outside discussions, sequence evolution not exactly like that trained on..) (Better Practice) | | | |
| **Additional PSFs to Consider** | **Conditions When Particularly Relevant** | **Additional PSFs to Consider** | **Conditions When Particularly Relevant** |
| Available staffing and resources | If typical CR staff is expected to be decreased or impacted so others must perform more than their typical tasks (not usually an issue) | Available staffing and resources | Particularly when many or complex actions need to occur concurrently or in a short time, and staffing needs may be stretched |
| Human-machine interface | If could be problematic, or not easily accessed or used (not usually an issue but consider, for instance, the need to use backboards, deal with common workarounds...) | Human-machine interface | If could be problematic (e.g., poor labeling) or not easily accessed or used |
| **Additional PSFs to Consider** | **Conditions When Particularly Relevant** | **Additional PSFs to Consider** | **Conditions When Particularly Relevant** |
| Environment in which the act needs to be performed | Potentially adverse or threatening situations such as fire, flood, seismic, loss of ventilation...(not usually an issue) | Environment in which the act needs to be performed | Potentially adverse situations such as high radiation, high temperature, high humidity, smoke, toxic gas, noise, poor lighting, weather, flooding, seismic... |

| In-Control Room Actions | | Ex-Control Room Actions | |
|---|---|---|---|
| Accessability and operability of equipment to be manipulated | If could be problematic, or not easily accessed or used such as the need to use backboards, or when indications/controls could be affected by the initiating event or other failures (e.g., loss of DC) | Accessability and operability of equipment to be manipulated | If could be problematic, or not easily accessed or used such as when the equipment could be affected by the initiating event or the environment (e.g., fire, flood, weather) |
| The need for special tools (keys, ladders, hoses, clothing such as to enter a radiation area...) | Not usually an issue but consider, for instance, accessability of keys for keylock switches | The need for special tools (keys, ladders, hoses, clothing such as to enter a radiation area...) | For situations where other than simple switch or similar type operations are necessary, or when needed to be able to access the equipment |
| Communications (strategy and coordination) as well as whether one can be easily heard | Not usually an issue - simply ensure that communication strategy allows crisp direction and proper feedback; otherwise only in special situations such as needing to communicate with SCBAs on | Communications (strategy and coordination) as well as whether one can be easily heard | For situations where communication among crew members (locally and/or with CR) are likely to be needed and there could be a threat such as too much noise, failure of the communication equipment, availability and location issues associated with the communication equipment... |
| **Additional PSFs to Consider** | **Conditions When Particularly Relevant** | **Additional PSFs to Consider** | **Conditions When Particularly Relevant** |

| In-Control Room Actions | | Ex-Control Room Actions | |
|---|---|---|---|
| Time of day | Special sequences or events such as involving numerous failures where task workloads may be extremely high and preferred additional in-CR staffing needs may be difficult to obtain such as during graveyard shift (typically not an issue) | Time of day | Particularly when many or complex actions need to occur concurrently or in a short time, and staffing needs may be stretched such as during graveyard shift |
| | | Special fitness needs | For special situations expected to involve the use of heavy or awkward tools/equipment, carrying hoses, climbing... |
| | | Team/crew dynamics and crew characteristics (degree of independence among individuals, operator attitudes - biases - rules, use of status checks, approach for implementing procedures, e.g., aggressive vs. slow and methodical...) | To the extent that the timing and the appropriateness of the directions from the CR, and the subsequent carrying out of the ex-CR action(s) could be affected |
| | | Consideration of 'realistic' accident sequence diversions and deviations (e.g., extraneous alarms, outside discussions, sequence not exactly like that trained on...) | To the extent that these could affect the timing, specific directions, or successful performance of the ex-CR action(s) |

addressed by the following list of factors, identification of other performance-shaping factors should complement the list below. Consideration of the impact of the factors on the HEPs should be as plant- and accident sequence-specific as necessary to address the issue and confirmed, where useful, by such techniques as talkthroughs, walkdowns, field observations, simulations, and examination of past events in order to be realistic. Appendix A provides more specific guidance and discussion of the PSFs presented below, as well as why some are considered generally more important than others.

It should be apparent that the factors seemingly most relevant to the act (either as positive or negative influences) and having the most impact on the HEP, have been considered quantitatively. Further, the more the impacts of the factors have been determined based on talkthroughs, walkdowns, field observations, and simulations vs. simple assumptions or judgements, the better the quality of the HEP evaluations.

5.3.3.6 Good Practice #6:

Dependencies among the post-initiator HFEs and hence the corresponding HEPs in an accident sequence should be quantitatively accounted for in the PRA model by virtue of the joint probability used for the HEPs. This is to account for the evaluation of each sequence holistically, considering the performance of the operators throughout the sequence response and recognizing that early operator successes or failures can influence later operator judgments and subsequent actions. This is particularly important so that too optimistic combined probabilities are not inadvertently assigned potentially resulting in the inappropriate decrease in the risk significance of human actions and related accident sequences and equipment failures. In the extreme, this could result in the inappropriate screening out of accident sequences from the model because the combined probability of occurrence of the events making up an accident sequence drops below a threshold value used in the PRA to drop sequences from the final risk results.

In analyzing for possible dependencies among the HFEs in an accident sequence, look for links among the acts including:

- the same crew member(s) is responsible for the acts,

- the actions take place relatively close in time in the sense that a crew "mindset" or interpretation of the situation might carryover from one event to the next,

- the procedures and cues used along with the plant conditions related to performing the acts are identical (or nearly so) or related, and the applicable steps in the procedures have few or no other steps in between the applicable steps,

- there are similar performance shaping factors for performing the acts,

- how the acts are performed is similar and they are performed in or near the same location, and

- there is reason to believe that the decision processes associated with the events might be related and the interpretation of the need for one action might bear on the crews decision regarding another action.

The more the above commonalities and similarities exist, the greater the potential for dependence among the HFEs (i.e., if the first act is not performed correctly, there is a higher likelihood the second, third... act(s) will also not be performed correctly; or vice versa if the act(s) are successful). For example, if nearly all or all of the above characteristics exist, very high or complete dependence should generally be assumed. If only one or two of the above characteristics exist, then analysts will need to evaluate the likely strength of their effect and the degree of dependence that should be assumed and addressed in quantification.

The resulting joint probability of the HEPs in an accident sequence should be such that it is in line with the above characteristics and the following guidance, unless justified otherwise:

- The total combined probability of all the HFEs in the same accident sequence/cut set should not be less than a justified value. It is suggested that the value not be below the ~0.0001 to 0.00001

range since it is typically hard to defend that other not specifically treated dependent failure modes (e.g., even heart attack) cannot occur. Depending on the independent HFE values, the combined probability may need to be higher.

- To the extent the joint HEPs are looked at separately, but a previous human action in the sequence has failed, then:

  - A factor of 3-10 higher than what would have been the independent HEP value for the subsequent act(s) exists for low to moderate dependence

  - 0.1 up to 0.5 is the resulting probability value used for the subsequent HEP(s) for high dependence

  - $\geq$0.5 exists for the subsequent HEP(s) for very high or 1.0 for complete dependence.

5.3.3.7 Good Practice #7:

Mean values for each HEP (excluding conservative HEPs) and an assessment of the uncertainty in the mean values should be performed at least for the dominant HEPs to the extent that these uncertainties need to be understood and addressed in order to make appropriate risk-related decisions. Assessments of uncertainty are typically performed by:

- assigning uncertainty distributions for the HEPs and propagating them thru the quantitative analysis of the entire PRA such as by a Monte Carlo technique, and/or

- performing sensitivity analyses that demonstrate the effects on the risk results for extreme estimates in the HEPs based on at least the expected uncertainty range about the mean value.

Note, in some cases, it may be sufficient to address the uncertainties by just qualitative arguments without the need to specifically quantify them (e.g., justifying why the HEP cannot be very uncertain or why a change in the HEP has little relevancy to the risk-related decision to be made).

In assessing the uncertainties and particularly when assigning specific uncertainty distributions, the uncertainties should include (a) those epistemic uncertainties existing because of lack of knowledge of the true expected performance of the human for a given context and associated set of performance-shaping factors (i.e., those factors for which we do not have sufficient knowledge or understanding as to the "correct" HEP, such as how time of day affects the bio-rhythm and hence, performance of operators), and (b) consideration of the combined effect of the relevant aleatory (i.e., random) factors to the extent they are not specifically modeled in the PRA and to the extent that they could significantly alter the context and performance-shaping factor evaluations for the HFE, and thereby the overall HEP estimate.

Concerning the latter, while it is best to specifically model the aleatory factors in the PRA (i.e., those factors that are random and could significantly affect operator performance, for example, the time of day the initiator occurs, whether or not other nuisance alarms or equipment failures may co-exist with the more important failures in the sequence, whether a critical equipment failure occurs early in the sequence or late in the sequence, etc.), this is often impractical and is typically not done

as it would make the PRA model excessively large and unwieldy. Thus in assigning the mean HEP and uncertainty distribution, analysts should reflect an additional contribution from random factors associated with the plant condition or overall action context. This can be done by considering the relevant aleatory (i.e., random) factors, their likelihoods of occurrence, and their effects on the HEP estimate.

For example, suppose for an accident sequence(s) it is judged that the human performance will be significantly affected by the number of "nuisance and extraneous failures," as opposed to when no or few nuisance/extraneous failures exist (and yet these two plant "states" are not explicitly defined by the PRA model). Further, based on the analyst considering how the HEP is affected, a value of $P_o$ would be estimated for when no or few nuisance/extraneous failures exist and a value of $P_1$ would be estimated for when many do exist, and the difference between $P_o$ and $P_1$ is significant (e.g., factor of 10). It is also judged that many nuisance/extraneous failures will occur about 50% of the time based on past experience. The resulting combined mean HEP value is $0.5P_o + 0.5P_1$ considering this random factor. The overall uncertainty about the combined mean HEP value should reflect the weighted epistemic uncertainties in $P_o$ and $P_1$ (such as by a convolution approach, via an approximation, or other techniques). While it is not expected that such a detailed evaluation be done for every random situation or for every HEP, the mean and uncertainty estimates for the most dominant HEPs should account for any such perceived important aleatory factors that have not otherwise been accounted for (i.e., the factors, considering their likelihoods and effects on the HEP, are anticipated to have a significant impact on the resulting overall HEP).

Whatever uncertainty distributions are used, the shape of the distributions (log-normal, normal, beta...) are typically unimportant to the overall risk results (i.e., the PRA results are usually not sensitive to specific distributions). Further, typical uncertainties include values for the HEP that represent a factor of 10 to 100 or even more between the lower bound value and the upper bound value that encompass the mean value.

5.3.3.8 Good Practice #8:

The post-initiator HEPs (excluding the conservative HEPs) should be reasonable from two standpoints:

- first and foremost, relative to each other (i.e., the probabilistic ranking of the failures when compared one to another), and

- in absolute terms (i.e., each HEP value) to the extent that the sensitivity of the risk-related decision is not important as to the absolute values for the HEPs.

This reasonableness should be checked based on consideration of actual plant experience and history, against other evaluations (such as for similar acts at other plants), and the qualitative understanding of the actions and the relevant contexts and performance-shaping factors under which the acts are performed.

It is suggested that a rank-ordered list of the post-initiator HFEs by probability be used as an aid for checking reasonableness. As part of such a list, it is particularly worthwhile to compare "like" HFEs for different sequences such as failure to manually depressurize in a BWR when all high

pressure injection is lost during a LOCA as compared to the same action but during a simple transient. For example, simple, procedure-guided actions with easily recognized cues and plenty of time to perform the actions, should have lower HEPs than complex, memorized, short time available type actions, all other factors being the same. Typical expectations of most post-initiator HEPs are in the 01 to 0.0001 range and depend particularly on the relevant contextual factors and proper consideration of dependencies as discussed under many of the Good Practices covered above. Helpful checks include:

- For a HFE, do any one or two dominant performance-shaping factors exist or is the cumulative effect of the relevant performance-shaping factors such that they are either so negative or so positive that a 'sanity check' would suggest a high HEP (e.g., 0.1) or a low HEP (e.g., 1E-4) respectively? Accordingly, this very high or low probability HFE should be one of the higher or lower probability HFEs relative to the other HFEs in the model. For example, while the manual scram action may need to be done in a short time, it is a proceduralized action, is often an early step in procedures, is performed often in training, and thus has become such an "automatic" action (the predominant positive factor) that a low HEP is justified.

- Are there seemingly balanced combinations of both positive and negative factors, or are there weak to neutral factor effects? If so, this is likely to lead to in-between values for the HEPs (e.g., ~0.01) placing these HFEs (relative to others) 'in the middle'.

- Do the individual HEPs and the relative ranking of the HFEs seem consistent with actual or simulated experience? For example, if it is known that operators 'have trouble with' a specific act(s) in simulations or practiced events, and yet the assigned HEP is very low (e.g., 1E-3 or lower), this may be a reason to question and revisit the assigned HEP.

- Do other similar plant and action analyses support the HEP evaluation? Recognize, however, that there may be valid reasons why differences may exist and thus this check is not likely to be as helpful as the others above.

### 5.3.4 POSSIBLE IMPACTS OF NOT PERFORMING GOOD PRACTICES AND ADDITIONAL REMARKS:

Besides the obvious concerns about inaccuracies in the HEP quantification and thus whether the HEPs "make sense", as well as the resulting potential misinformation about the dominant risk contributors if quantification is not done well, the following observations are noted.

- Use of conservative values is a useful and most often, necessary part of HRA so as to avoid the expenditure of resources on unimportant human events and accident sequences. The above guidance is aimed at allowing some conservative values without inadvertently and inappropriately allowing the analytical phenomenon of, for instance, multiplying four human events in the same sequence each at a conservative estimate of 1E-1 to yield a 1E-4 combined probability, without checking for dependencies among the human events. In such a case some human failure events and combinations of events, or even whole accident sequences, may inappropriately screen out of the PRA model entirely because the accident sequence frequency drops below a model threshold. Hence some of the dominant individual or combination contributors may be missed. This is why the conservative estimates both individually and for

combined events should not be too low. Further, if conservative values are left permanently assigned to some human failure events that should be assessed with more detail to obtain a more realistic assessment of risk (supposedly lowering the probability), the risk significance of these human failure events and related equipment failures are likely to be over-emphasized at the expense of improperly lessening the relative importance of other events and failures.

- It is important to be sure that dependencies among the various modeled HFEs including those with conservative values, have been investigated. Treating HFEs, whether with conservative values or based on more detailed analysis, as independent acts without checking for dependencies (thereby being able to multiply the individual HEPs) can inappropriately lessen the risk significance of those HFEs and related equipment failures in accident sequences. This can cause the inappropriate dropping out of accident sequences because the sequences quantitatively drop below a model threshold value as discussed above under screening. Proper consideration of the dependencies among the human actions in the model is necessary to reach the best possible evaluation of both the relative and absolute importance of the human events and related accident sequence equipment failures.

- The use of mean values and addressing uncertainties are a part of the Regulatory Guide 1.174[11] guidance and to the extent addressed therein, the HRA quantification needs to be consistent with that guidance when making risk-informed decisions.

- There can be a tendency for analysts to want to use an existing PRA model to address issues such as changes to the plant, without spending the appropriate time to revisit some of the underlying assumptions and modeling choices made to create the original PRA. A review should be done to see if these assumptions and choices still apply for the issue being addressed. In this case, some post-initiator human failure events may be quantified in the original model using conservative estimates and detailed failure probabilities that may not be appropriate for the new issue being addressed. As an example, where higher conservative values may have been acceptable for purposes of the original PRA, these may over-estimate the contribution of these human failure events for the issue being addressed. Further, the relative risk contribution of equipment and associated accident sequences with which the human failure events appear, may be artificially too high (and therefore other events too low) because of the conservative values. Hence it is good practice to revisit the use of conservative estimates and detailed human failure event probabilities in order to appropriately address the issue.

### 5.4 Adding recovery actions to the PRA

5.4.1 OBJECTIVE: To address what recovery actions can be credited in the post-initiator HRA and the requirements that should be met before crediting recovery actions. Adding recovery actions is common practice in PRA and accounts for other reasonable actions the operators might take to avoid severe core damage and/or a large early release that are not already specifically modeled. For example, in the PRA modeling of an accident sequence involving a loss of offsite power, subsequent station blackout, and loss of all injection, it would be logical and common to credit the operators attempting to recover offsite or onsite power (and thus ac-powered core cooling systems) as well as perhaps locally aligning an

independent firewater system (not affected by the station blackout) for injection. The failure to successfully perform such actions would subsequently be added to the accident sequence model thereby crediting the actions and further lowering the overall accident sequence frequency because it takes additional failure of these actions before the core is actually damaged. The following provides good practices for crediting post-initiator recovery actions while implementing the related Standard requirements.

5.4.2 CORRESPONDING ASME STANDARD REQUIREMENTS:

The Standard requires that recovery actions be modeled only if it has been demonstrated that the action is plausible and feasible for those sequences to which they are applied. There are multiple supporting requirements in the Standard that address what recovery actions can be credited as well as the need to consider dependencies among the HFEs and any recovery actions that are credited.

5.4.3 GOOD PRACTICES:

5.4.3.1 Good Practice #1:

Based on the failed functions, systems, or components, identify recovery actions to be credited that are not already included in the PRA (e.g., restoring offsite power loss, aligning another backup system not already accounted for...) and that are appropriate to be tried by the crew to restore the failure. The following should be considered in defining appropriate recovery actions:

- the failure to be recovered,
- whether the cues will be clear and provided in time to indicate the need for a recovery action, and the failure that needs to be recovered,
- the most logical recovery actions for the failure and based on the cues that will be provided,
- the recovery is not a repair action (e.g., the replacement of a motor on a valve so that it can be operated),
- whether sufficient time is available following the timing of the cues (for the sequence/cut set) for the recovery action to be diagnosed and implemented to avoid the undesired outcome,
- whether sufficient crew resources exist to perform the recovery(ies),
- whether there is procedure guidance to perform the recovery(ies),
- whether the crew has trained on the recovery action(s) including the quality and frequency of the training,

- whether the equipment needed to perform the recovery(ies) is accessible and in a non-threatening environment (e.g., extreme radiation), and

- whether the equipment needed to perform the recovery(ies) is available in the context of other failures and the initiator for the sequence/cut set.

In addressing the above issues and assessing which recovery action, or a few, to credit in the PRA, just as with any other HFE, all the good practices provided earlier in Sections 5.1, 5.2, and 5.3 apply to these recovery actions as well (i.e., the failure to recover is just another HFE like all the other post-initiator HFEs). In general, no recovery should be credited where any of the above considerations are not met (e.g., there is not sufficient time, there are no cues that there is a problem, there are not sufficient resources, there is no procedure or training, etc.). Exceptions may be able to be justified in unique situations, such as a procedure is not needed because the recovery is a skill-of-the-craft, non-complex, and easily performed; or the specific failure mode of the equipment is known for the sequence (this is usually not the case at the typical level of detail in a PRA) and so "repair" of the failure can be credited because it can be easily and quickly diagnosed and implemented. Any exceptions should be documented as to the appropriateness of the recovery action.

When considering multiple recoveries (i.e., how many recoveries to be credited in one accident sequence/cut set), the above considerations apply to all the recoveries. The analyst should also consider that one recovery may be tried (perhaps even multiple times) and then the second recovery may be tried but with even less time and resources available because of the attempts on the first recovery. Hence the failure probability of the second recovery should be based on more pessimistic characteristics (e.g., less time available, less resources, etc.) than if such a possibility is not considered.

5.4.3.2 Good Practice #2:

As stated above, all the good practices provided earlier in Sections 5.1, 5.2, and 5.3 apply. From these good practices, particular attention should be paid to accounting for dependencies among the HFEs including the credited recovery actions. More specifically, dependencies should be assessed:

- among multiple recoveries in the accident sequence/cut set being evaluated, and
- between each recovery and the other HFEs in the sequence/cut set being evaluated..

As part of this effort, the analyst should give proper consideration to the difficulties people often have in overcoming an initial mind-set despite new evidence (e.g., look how long the PORV remained open in the Three Mile Island accident despite new cues of the problem, different personnel, etc.). For this and similar reasons, the assessing of no dependence needs to be adequately justified to ensure the quantified credit for the recovery action(s) is not unduly optimistic.

5.4.3.3 Good Practice #3:

Quantify the probability of failing to perform the recovery(ies) by:

- using representative data that exists and deemed appropriate for the recovery event (i.e., a data-based approach such as using data that exists for typical times to recover offsite power)

- using the HRA method/tool(s) used for the other HFEs (i.e., using an analytical/modeling approach).

In performing the quantification, one should ensure that all the good practices under Section 5.3 are followed (for each individual recovery as well as for multiple/joint recovery credit). In addition, if using data, ensure the data is applicable for the plant/sequence context or that the data is modified accordingly. For example, a plant may use available experience data for the probability of failing to align a firewater system for injection but the experience data is based on designs for which all the actions can be taken from the main control room whereas for this plant, the actions have to be performed locally.

5.4.4 POSSIBLE IMPACTS OF NOT PERFORMING GOOD PRACTICES AND ADDITIONAL REMARKS:

The primary concern for not performing the above good practices is that recovery credit could be applied too optimistically; that is, the failure to recover is assigned too low a probability. Hence an under-estimate of the failure to recover is applied to the PRA accident sequence/cut set, making the affected sequence/cut set artificially too low in risk significance. This can subsequently affect the ranking of the important sequences, equipment failures, and human actions potentially leading to false conclusions of the dominant risk contributors.

## 6. HRA DOCUMENTATION

The ASME Standard[5] provides a set of requirements for documenting a human reliability analysis (HRA) in a manner that facilitates PRA applications, upgrades, and peer review. Specific requirements are provided. The following provides good practice for documenting a HRA building on those requirements.

Good Practice:

The level of detail that needs to be addressed in the documentation is dependent on the PRA application and the issue being addressed as well as the objectives, scope, and level of detail of the analysis. Whatever documentation is provided, the test for adequate documentation should be: "Can a knowledgeable reviewer understand the analysis to the point that it can be at least approximately reproduced and the resulting conclusion reached if the same methods, tools, data, key assumptions, and key judgments and justifications are used?" Hence, the documentation should include the following, but only to the extent it is applicable for the application:

- the overall approach and disciplines involved in performing the HRA including to what extent talk-throughs, walkdowns, field observations, and simulations were used,

- summary descriptions of the HRA methodologies, processes, and tools used to:

  - identify the pre-and post-initiator human actions,

  - screen pre-initiators from modeling,

  - model the specific HFEs including decisions about level of detail and the grouping of individual failures into higher order HFEs,

  - quantify the HEPs with particular attention to the extent to which plant and accident sequence-specific information was used, as well as how dependencies were identified and treated,

- assumptions and judgments made in the HRA, their bases, and their impact on the results and conclusions (generic or on a HFE-specific basis, as appropriate),

- for at least each of the HFEs important to the risk decision to be made, the PSFs considered, the bases for their inclusion, and how they were used to quantify the HEPs, along with how dependencies among the HFEs and joint probabilities were quantified,

- the sources of data and related bases or justifications for:

  - the screening and conservative values,
  - the best estimate values and their uncertainties with related bases,

- the results of the HRA including a list of the important HFEs and their HEPs, and

- conclusions of the HRA.

## 7. ERRORS OF COMMISSION (EOCs)

Explicit modeling of errors of commission (i.e., committing an incorrect act) are generally beyond current PRA practice and are not explicitly addressed in the ASME Standard HRA requirements. This is largely because of the seemingly unlimited set of acts that an operator might perform that are adverse to safe shutdown (i.e., fail or make unavailable equipment/functions relevant to mitigating the scenario, or otherwise exacerbate the scenario such as opening a PORV and causing an unwanted loss of coolant accident) even for what may appear to be justifiable reasons. Errors of omission (i.e., failure to perform the correct act) are typically modeled in PRAs because the set of correct acts is better known for each sequence thus limiting the number of human failures that need to be modeled.

Given the current state of the art, EOCs should not be expected to be explicitly treated in any HRA/PRA evaluation. Yet, the use of risk in any issue assessment should at least ensure that conditions that promote likely EOCs do not exist or have not been introduced by a plant change or modification. To the extent any EOCs are modeled, all the guidance in this document has been written with both types of errors in mind; that is, all the same good practices apply whether the error is one of omission or commission.

When considering the potential for situations that may make EOCs somewhat likely, the premise of any evaluation should be that:

- operators are performing in a rationale manner (e.g., no sabotage), and

- the procedural and training guidance is being used by the crew based on the plant status inputs they are receiving.

Using this premise, EOCs are considered to be largely the result of problems in the plant information/operating crew interface (wrong, inadequate information is present, or the information can be easily misinterpreted) or in the procedure-training/operating crew interface (procedures/training do not cover, very well, the actual plant situation because they provide ambiguous guidance, no guidance, or even unsafe guidance for the actual situation that may have evolved in a somewhat unexpected way).

With a present focus on reviewing potential applications of current PRAs and not on the whole-sale addition of EOCs to current PRAs, the following is offered as guidance in this area to aid in ensuring EOC-prone conditions do not exist or have not been introduced as part of a plant change. Hence, a review of a plant change should look for situations where one or more of the following characteristics are introduced as a result of the change and thus should be corrected if possible.

- To deal with the 'bad information' interface, an analysis/review should at least look for those acts that operators may take that (a) would fail or otherwise make unavailable a PRA function or system, or (b) would reduce the accident mitigating redundancy available, or (c) would exacerbate an accident challenge, because the change has caused such an action to be performed on the basis of just one primary input/indication for which there is no redundant means to verify the true plant status. Such a situation identifies a vulnerable case where EOCs may likely be performed based on just one erroneous (failed, spurious, etc.) input such as an alarm, indicator, or verbal cue of an observed condition.

.

  In identifying such cases, one should keep in mind that multiple indications may use the same faulty input (e.g., subcooling margin indication and primary system indication may use the same pressure transmitter(s); multiple reactor vessel level indications may rely on the same power supply) and hence a single fault may actually affect multiple inputs observable to the operator. Depending on the how the failure affects the indications (fail high, low, mid-scale, etc.), the failure may not be "obvious" and a EOC-prone situation may exist that may need to be rectified.

- To deal with the procedure-training interface, an analysis/review should at least look for those acts that operators may take that (a) would fail a PRA function or system, or (b) would reduce the accident mitigating redundancy available, or (c) would exacerbate an accident challenge, because the change has caused the procedure (including entry conditions) and/or training guidance:

  - to become ambiguous/unclear (e.g., vague criteria as to when to abandon the main control room),

  - to introduce a repetitive situation in the response steps where a way to proceed out of the procedure and/or the specific repetitive steps is not evident (e.g., at the end of a series of steps, the procedure calls for a return to a previous step with no clear indication as to how the operators ultimately get out of the repetitive loop of steps,

  - to place the operators in dilemma conditions without some guidance/criteria as to how to "solve" the dilemma (e.g., being vague as to whether or not to shutdown a diesel with a cooling malfunction when all other ac power is unavailable),

  - to require the operators to rely on memory especially for complex or multi-step tasks, or

  - to require the operators to perform calculations or make other manual adjustments especially in time-sensitive situations.

The above identify vulnerable cases where EOCs may likely be performed because the procedures and/or training do not adequately cover accident situations that may be faced by the operator or rely on techniques (require memory or adjustments) that may be difficult to perform properly especially when in a dynamic response situation. In these cases, mismatches between the actual event response that is required and the procedure/training guidance can become magnified making conditions potentially more prone to EOCs.

## 8. REFERENCES

[1] *Use of Probabilistic Risk Assessment Methods in Nuclear Activities: Final Policy Statement*, Federal Register, Vol 60, p. 42622 (60FR 42622), US Nuclear Regulatory Commission, August 16, 1995.

[2] *Code of Federal Regulations 10, Parts 1 to 50*, Office of the Federal Register National Archives and Records Admission, Revised as of January 1, 2001.

[3] *An Approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis*, Regulatory Guide 1.174, Rev. 1, US Nuclear Regulatory Commission, November 2002.

[4] *An Approach for Determining The Technical Adequacy of Probabilistic Risk Assessment Results For Risk- Informed Activities*, Draft Regulatory Guide 1.200, U.S. Nuclear Regulatory Commission, February 2004.

[5] *Standard for Probabilistic Risk Assessment for Nuclear Power Plant Applications,* ASME RA-S-2002, American Society of Mechanical Engineers, April 5, 2002.

[6] *Probabilistic Risk Assessment Peer Review Process Guidance*, NEI 00-02 Revision A3, Nuclear Energy Institute, March 2000

[7] *Use of Probabilistic Risk Assessment in Plant-Specific, Risk-Informed Decisionmaking: General Guidance,* NUREG-0800, Chapter 19, Rev. 1, US Nuclear Regulatory Commission, November 2002.

[8] A.D. Swain and H.E. Guttmann, *Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications - Final Report,* NUREG/CR-1278, SAND80-0200, Sandia National Laboratories, August 1983.

[9] Embrey, D. E., et al., *SLIM-MAUD: An Approach to Assessing Human Error Probabilities Using Structured Expert Judgment (Vols. I & II),* NUREG/CR-3518, Brookhaven National Laboratory, Upton, NY, 1984.

[10] *Technical Basis and Implementation Guidelines for A Technique for Human Event Analysis (ATHEANA),* NUREG-1624, Rev. 1, US Nuclear Regulatory Commission, May 2000.

[11] *Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants,* NUREG-1150, US Nuclear Regulatory Commission, December 1990

[12] Alan D. Swain, *Accident Sequence Evaluation Program Human Reliability Analysis Procedure,* NUREG/CR-4772, SAND86-1996, Sandia National Laboratories, February 1987.

[13] *SHARP1 - A Revised Systematic Human Action Reliability Procedure,* EPRI NP-7183-SL, Electric Power Research Institute, December 1990.

[14] *SPAR-H Method,* NUREG/CR- later, INEEL/EXT-02-10307, Idaho National Engineering and Environmental Laboratory, DRAFT, November 2002.

[15] J. Reason, *Human Error,* Cambridge, England: Cambridge University Press, 1990 and *Managing the Risks of Organizational Accidents,* Aldershot, UK: Ashgate, 1997.

[16] D.D. Woods et. al., *Behind Human Error: Cognitive Systems, Computers, and Hindsight,* Crew System Ergonomics Information Analysis Center (CSERIAC), The Ohio State University, Wright-Patterson Air Force Base, Columbus, OH, December 1994.

[17] M.R. Endsley, *Towards a theory of situation awareness in dynamic systems,* Human Factors, 37, pp. 65-84, 1995.

ACKNOWLEDGMENTS

We would like to thank Gareth Parry and Susan Cooper of the USNRC for several thorough reviews of this document at various stages of it development and for many important comments and suggestions. Similarly, we would like to thank Dennis Bley of Buttonwood Consulting and Bruce Hallbert of INEEL for their reviews and comments.

## APPENDIX A

### Guidance on Consideration of Performance-Shaping Factors for Post-Initiator HFEs

The following provides more detail on the performance-shaping factors presented in Section 5.3.3.5, including some key characteristics to consider when assessing the influence of these performance-shaping factors on the failure probability for a human failure event (HFE). Included are important interactions among the factors that should also be examined when assessing the holistic impact of the performance-shaping factors on operator performance. These factors need to be assessed on a plant-specific and accident sequence-specific basis considering the relevant context and the act to be performed.

It is important to re-iterate that this Appendix is written for the specific purpose of addressing post-initiator HFEs in a risk assessment for commercial nuclear power plant operations occurring nominally at full power, and for internal initiating events. However, much of it is considered useful to other modes of operation and for other industry applications such as safety assessments of chemical plants, space mission risk assessments, and others. Similarly, much of it is considered applicable for external initiating events but it should be used with the additional context of such events in mind (e.g., shaking during a seismic event). Additionally, portions of this Appendix may be of benefit in examining human actions related to nuclear materials and safeguard types of applications.

Specific HRA methods and tools used by the industry may define and "measure" these performance-shaping factors somewhat differently than described here. That is, they may use a different explicit set of performance-shaping factors that 'roll-up' many of the factors listed below into the definitions of their specific factors (e.g., stress, workload). Nevertheless, these summaries are provided as one means with which to assess that the specific HRA method/tool has been used such that the characteristics described here have indeed been accounted for in the evaluation of post-initiator human error probabilities (HEPs).

While quantitative guidance is not provided (specific quantification depends on the method/tool that is used), the following should be useful in arriving at whether a performance-shaping factor, regardless of the method/tool, qualitatively is a weak/strong positive, neutral (or not applicable), or negative influence. The method/tool that is used, should have established scales and corresponding definitions for assessing each PSF qualitatively (e.g., "good', "adequate", "poor") and a way to interpret the result into a quantified HEP.

The performance-shaping factors are addressed below.

Applicability and suitability of training/experience. For both in-control room and local actions, this is an important factor in assessing operator performance. For the most part, in nuclear plants, operators can be considered "trained at some minimum level" to perform their desired tasks.

However, from a HRA perspective, the degree of familiarity with the type of sequences modeled in the PRA and the actions to be performed, can provide a negative or positive influence that should

be assessed on the likelihood of operator success. In cases where the type of PRA sequence being examined or the actions to be taken are not periodically addressed in training (such as covered in classroom sessions or simulated every one to two years or even more often) or the actions are not performed as part of their normal experience or on-job duties, this factor should be treated as a negative influence. The converse would result in a positive influence on overall operator performance.

One should also attempt to identify systematic training biases that may affect operator performance either positively or negatively. For example, training guidance in a pressurized water reactor (PWR) may provide a reluctance to use "feed and bleed" in a situation where steam generator feed is expected to be recovered. Other biases may suggest operators are allowed to take certain actions before the procedural steps calling for those actions are reached, if the operators are sure the actions are needed. Such training "biases" could cause hesitation and hence higher HEPs for the desired actions, as in the first case above, or as in the case of not waiting to take obvious actions, be a positive influence.

It is incumbent on the analyst to ensure that training and/or experience is relevant to the PRA sequence situation and desired actions. The more it can be argued that the training is current, "is like the real event," is varied enough to represent differences in the way the event can evolve, and proficiency is demonstrated on a periodic basis, the more positive this factor. If there is little or no training/experience or there are potentially negative training biases for the PRA sequence being examined, this factor should be considered to have a negative influence.

Suitability of relevant procedures and administrative controls. For both in-control room and local actions, this is an important factor in assessing operator performance. Similar to training, for the most part, procedures exist to cover many types of sequences and operator actions.

However, from a HRA perspective, the degree the procedures clearly and unambiguously address the types of sequences modeled in the PRA and the actions to be performed, dictates whether they are a negative or positive influence on operator performance. Where procedures have characteristics like those below related to the desired actions for the sequences of interest, this factor should be considered a negative influence:

- ambiguous/unclear/non-detailed steps for the desired actions in the context of the sequence of interest,
- situations can exist where the operators are likely to have trouble identifying a way to proceed forward,
- there is a requirement to rely on considerable memory,
- operators must perform calculations or make other manual adjustments especially in time-sensitive situations,

- there is no procedure or the procedure is likely to not be available especially when taking local actions "in the heat of the scenario" and it cannot be argued that the desired task is simple and a "skill of the craft" or automatic/memorized activity that is trained on or there is routine experience.

Otherwise, this factor should be considered as adequate or even a positive influence.

Talk-throughs with operations and training staff can be helpful in uncovering 'difficulties' or 'ease' in using the relevant procedures considering the associated training that the operators receive and the way the operators interpret the use of the procedures.

Availability and clarity of instrumentation (cues to take actions as well as confirm expected plant response). For both in-control room and local actions, this is an important factor since operators, other than for immediate and memorized response actions, take actions based on diagnostic indications and look for expected plant responses to dictate follow-on actions. For in-control room situations, typical nuclear plant control rooms have sufficient redundancy and diversity for most important plant parameters. For this reason, most HRA methods inherently assume that adequate instrumentation typically exists. Nevertheless, this should be verified looking for the following characteristics that could make this a negative performance-shaping factor, particularly in situations where there is little redundancy in the instrumentation associated with the act(s) of interest:

- the key instrumentation associated with an act is adversely affected by the initiating event or subsequent equipment failure (e.g., loss of DC power causing loss of some indications, spurious or failed as a result of a hot short from a fire)

- the key instrumentation is not readily available and may not be typically scanned such as on an obscure back panel

- the instrumentation could be misunderstood or may be ambiguous because it is not a direct indication of the equipment status (e.g., PORV position is really the position of the solenoid valve and not the PORV itself)

- the instrumentation is operating under conditions for which it is not appropriate (e.g., calibrated for normal power conditions as opposed to shutdown conditions)

- there are so many simultaneous changing indications and alarms or the indication is so subtle, particularly when the time to act is short, it may be difficult to "see and pick out" the important cue in time (e.g., a changing open-close light for a valve without a concurrent alarm or other indication, finding one alarm light among hundreds).

The above also applies to local actions outside the control room, recognizing that in some situations, less instrumentation may exist (e.g., only one division of instrumentation and limited device actuators on the remote shutdown panel). However, on the positive side, local action indications often can include actual/physical observation of the equipment (e.g., pump is running, valve stem shows it is closed) that compensates for any lack of other indicators or alarms.

It is incumbent on the analyst to ensure that adequate instrumentation is available and clear so that the operators will know the status of the plant and when certain actions need to be taken.

Time available and time required to complete the act, including the impact of concurrent and competing activities. This can be an important influence for both in-control room and local actions since clearly, if there is not enough or barely enough time to act, the estimated HEP is expected to be quite high. Conversely, if the time available far exceeds the time required and there are not multiple competing tasks, the estimated HEP is not expected to be strongly influenced by this factor.

It is important that the time available and the time needed to perform the act be considered *in concert with* many of the other performance-shaping factors and the demands of the sequence. This is because the thermal-hydraulic inputs (e.g., time to steam generator dryout, time to start uncovering the core), while important, are not the only influences. (Note, it is best if the thermal-hydraulic influences are derived from plant-specific or similar analyses rather than simple judgments).

The time to perform the act, in particular, is a function of the number of available staff, the clarity and repetitiveness of the cues that the act needs to be performed, the HMI, the complexity involved (discussed later), the need to get special tools or clothing (discussed later), consideration of diversions and other concurrent requirements (discussed later), where in the procedures the steps for the act of interest are called out, crew characteristics such as whether the crews are generally aggressive or slow and methodical in getting through the procedural steps, and other potential 'time sinks'.

Clearly there is judgment involved, but as described here, it is not as simple as watching an operator perform an act in ideal conditions with a stop watch to determine the time required to perform the act. Only when the sequence context is considered holistically with the interfacing performance-shaping factors that have been mentioned here, can more meaningful "times" be estimated. Hence, especially for complex acts and/or situations, walkdowns and simulations can be helpful in ensuring overly optimistic "times" have not been estimated. Whatever HRA method/tool is used, determination of these times should include the considerations provided here.

Complexity of required response along with workload, time pressure, the need for special sequencing, and the familiarity of the situation. This is one of those catch-all type factors that attempts to measure the overall complexity involved for the situation at hand and for the act itself (e.g., many steps have to be performed by the same operator in rapid succession vs. one simple skill-of-the craft action). Many of the other performance-shaping factors address elements of the overall complexity such as the need to decipher numerous indications and alarms, many and complicated steps in a procedure, poor HMI, etc. Nevertheless, this factor should also capture 'measures' such as the ambiguity of the task, the degree of mental effort or knowledge involved, whether this is a multi-variable or single variable associated task, the overall task load and time pressure on the operators, whether special sequencing is required in order for the act to be successful especially if it involves multiple persons in different locations, whether the activity may require very sensitive and careful manipulations by the operator, etc. The more these "measures"

describe an overall complex situation, this performance-shaping factor should be found to be a negative influence. To the extent these "measures" suggest a simple, straightforward, unambiguous process, this factor should be found to be nominal or even ideal (i.e., positive influence).

Team/crew dynamics and crew characteristics (degree of independence among individuals, operator attitudes - biases - rules, use of status checks, approach for implementing procedures, e.g., aggressive vs. slow and methodical crew). This is another catch-all type of factor which can be important particularly to in-control room actions where the early responses to an event occur and the overall strategy for dealing with the event develops. In particular, the way the procedures are written and what is (or is not) emphasized in training (may be related to an organization influence), can cause systematic and nearly homogeneous biases and attitudes in most or all the crews that can affect overall crew performance. A review of this factor should include looking for such characteristics as:

- are independent actions encouraged or discouraged among crew members (allowing independent actions may shorten response time but could cause inappropriate actions going unnoticed until much later in the scenario)

- are there common biases or 'informal rules' such as a reluctance to do certain acts, whether the overall philosophy is to protect equipment or run it to destruction if necessary, or the way procedural steps are interpreted

- are periodic status checks performed (or not) by most crews so that everyone has a chance to 'get on the same page' and allow for checking what has been performed to ensure the desired activities have taken place

- is the overall approach of most crews to aggressively respond to the event, including taking allowed shortcuts through the procedural steps (which will shorten response times), or are typical responses slow and methodical (we trust the procedures type of attitude) thereby tending to slow down response times but making it less likely to make mistakes.

Observing simulations and using talk-throughs and walkdowns can provide valuable insights into the overall crew response dynamics, attitudes, and the typical times it takes them to get through various procedure steps and deal with unexpected failures or distractions. This knowledge can be a key input into the HEP evaluation including determining the typical time to respond (see that factor above).

Consideration of 'realistic' accident sequence diversions and deviations (e.g., extraneous alarms, outside discussions, the sequence evolution is not exactly like that trained on...). Particularly for in-control room actions where the early responses to an event occur and the overall strategy for dealing with the event develops, this can be an important factor to be considered. Through simulations, training, and the way the procedures are written, operators 'build up' some sense of expectations as to how various types of sequences are likely to proceed; even to the extent of recognizing alarm and indication patterns and what actions will likely be appropriate. To the extent the actual sequence may not be 'just like in the simulator,' such as involving other unimportant or

spurious alarms, the need for outside discussions with other staff or even offsite entities such as a fire department, differences in the timing of the failed events, and behavior of critical parameters, etc., all can add to the potential diversions and distractions that may delay response timing or in the extreme, even confuse the operators as to the appropriate actions to take.

Hence the 'signature' of the PRA accident sequence and the potential acts of interest should be examined against the expectations of the operators to determine if there is a considerable potential for such distractions and deviations. Observing simulations and talking with the operators can help in discovering such possibilities. This could impact the HEP mean value estimate as well as the uncertainty in the HEP, which may be important to assessing the potential risk or in establishing the limits for doing sensitivity studies with the HEP.

Available staffing/resources. For in-control room actions, this is generally not an important consideration (i.e., not a particularly positive or negative factor) since plants are supposed to maintain an assigned minimum crew with the appropriate qualified staff available in or very near the control room.

However, for ex-control room local actions, this can be an important consideration particularly dependent on (a) the number and locations of the necessary actions, (b) the overall complexity of the actions that are required to be taken, and (c) the time available to take the actions and the time required to perform the actions (see above for more on these related factors). For instance, where the number of actions are few and complexity is low and time available is high, one or two personnel available to perform the local actions may be more than enough and thus the available staffing can be considered to be adequate or even a positive factor. On the other hand, where the number of actions and their complexity is high, and with little time, perhaps three or more staff may be necessary. Additionally, the time of the day the initiating event occurs may be a factor since typically, night and graveyard shifts have fewer people available than the day shift (see more on this particular factor, below).

It is incumbent on the analyst to demonstrate that the available staffing is sufficient to perform the desired actions and/or assess the HEP(s) accordingly.

Human-machine interface (HMI). This is generally not an important factor relative to in-control actions since, given the many control room design reviews and improvements and the daily familiarity of the control room boards and layout, problematic human-machine interfaces have been taken care of or are easily worked around by the operating crew. Of course, any known very poor human-machine interface should be considered as a negative influence for an applicable action even in the control room. For example, if common workarounds are known to exist that may negatively influence a desired act, this should be accounted for in the HEP evaluation. Furthermore, it is possible that some unique situations may render certain human-machine interfaces less appropriate and for such sequences, the relevant interfaces should be examined.

However, since local actions may involve more varied (and not particularly human-factored) layouts and require operators to take actions in much less familiar surroundings and situations, any problematic human-machine interfaces can be an important negative factor on operator success.

For instance, if to reach a valve to open it manually requires the operator to climb over pipes and turn the valve with a tool while in a laid out position, or in-field labeling of equipment is generally in poor condition and could lengthen the time to find the equipment, etc., such 'less ideal' human-machine interfaces could mean this is a negative performance-shaping factor. Otherwise, if a review reveals no such problematic interfaces for the act(s) of interest, this influence can be considered adequate or even positive.

Walkdowns and field or simulator observations can be useful tools in discovering problems (if any exist) in the human-machine interface for the actions of interest. Sometimes, discussions with the operators will reveal their own concerns about issues in this area.

Environment in which the act needs to be performed. Except for relatively rare situations, this factor is not particularly relevant to in-control room actions given the habitability control of such rooms and the rare challenges to that habitability (e.g., control room fire, loss of control room ventilation, less lighting as a result of station blackout). However, for local actions, this could be an important influence on the operator performance. Radiation, lighting, temperature, humidity, noise level, smoke, toxic gas, even weather for outside activities (e.g., having to go on a potential snow-covered roof to reach the atmospheric dump valve isolation valve), etc., can be varied and far less than ideal. Hence any HEP assessment should ensure that the influence of the environment where the act(s) needs to take place is accounted for as a performance-shaping factor. This factor may be non-problematic (adequate) or a negative influence (even to the point of not being able to perform the act).

Accessability and operability of the equipment to be manipulated. As with the environment factor, this factor is not particularly relevant most of the time to in-control room actions except for special circumstances such as loss of operability of indications or controls as a result of the initiator or equipment failures (e.g., loss of DC). However, for local actions, accessability and the operability of the equipment to be manipulated may not always be ensured, and needs to be assessed in the context with such influences as the environment, the need to use special equipment (discussed later), and HMI. Hence any HEP assessment should ensure that this factor, for where the act(s) needs to take place, is accounted for as a performance-shaping factor. This factor may be non-problematic (adequate) or a negative influence (even to the point of not being able to perform the act).

The need for special tools (keys, ladders, hoses, clothing such as to enter a radiation area...). As for the environment and accessability factors, this factor is not particularly relevant to in-control room actions with the common exception of needing keys to manipulate certain control board switches or similar controls (e.g., key for explosive valves for standby liquid control injection in a BWR). However, for local actions, such needs may be more commonplace and necessary in order to successfully perform the desired act. If such equipment is needed, it should be ensured that the equipment is readily available, its location is readily known, and it is either easy to use or periodic training is provided, in order for this factor to be considered to be positive or adequate. Otherwise, this factor should be considered to have a negative influence on the operator performance, perhaps even to the point of making the failure of the desired action very high.

Communications (strategy and coordination) as well as whether one can be easily heard). For in-control room actions, this factor is not particularly relevant although there should be verification that the strategy for communicating in the control room is one that tends to ensure that directives are not easily misunderstood (e.g., it is required that the board operator repeat the act to be performed and then wait for confirmation before taking the act). Generally, it is expected that this will not be problematic; but any potential problems in this area (such as having to talk with special air packs and masks on in the control room in a minor fire) should be accounted for if they exist.

For local actions, this factor may be much more important because of the possible less than ideal environment or situation. It should be assured that the initiating event (e.g., loss of power, fire, seismic) or subsequent equipment faults are not likely to negatively affect the ability for operators to communicate as necessary to perform the desired act(s). For instance, having to set-up the equipment and talk over significant background noise and possibly having to repeat oneself many times should be a consideration - even if just as a possible 'time sink' for the time to perform the act. Additionally, there should be training on the use of the communication equipment, its location is readily known, and its operability periodically demonstrated and shown to be in good working condition. Depending on the status of these characteristics, this factor may be non-problematic (adequate) or a negative influence (even to the point of not being able to perform the act).

Special fitness needs: While typically not an issue for in-control room actions, this could be an important factor for a few local actions depending on the specific activity involved. Having to climb up or over equipment to reach a device, needing to move and connect hoses, using an especially heavy or awkward tool, are examples of where this factor could have some influence on the operator performance. In particular, the response time for an action may be increased for successful performance of the act. Physically demanding (or not) activities should be considered in the evaluation of any HEP where it is appropriate to do so. Talk-throughs or field observations of the activities involved can help determine whether such issues are relevant to a particular HFE.

Time of day: While it is recognized that time of day and similar influences such as day of shift can affect the bio-rhythm of personnel and potentially their performance, not much is understood on how to quantify such effects. Moreover, it is typically the PRA's intent to measure an average risk for the whole year, as opposed to at a specific point in time. For these reasons, time of day is not typically specifically treated in a HEP evaluation.

However, at least one easily measurable effect of the time of day is on the available level of staffing during the early stages of a transient response (see available staffing factor above). Especially if there are significant differences in the staffing levels depending on the time of day, it is advisable to either treat the staffing level in a HEP evaluation as the minimum available depending on the shift, or probabilistically account for these aleatory differences more explicitly in the PRA model.

# Expert elicitation approach for performing ATHEANA quantification

John Forester[a,*], Dennis Bley[b], Susan Cooper[c], Erasmia Lois[c], Nathan Siu[c],

Alan Kolaczkowski[d], John Wreathall[e]
[a]Sandia National Laboratories, P.O. Box 5800, MS0748, Albuquerque, NM 87185-0748, USA
[b]Buttonwood Consulting, Inc., Oakton, VA, USA
[c]US Nuclear Regulatory Commission, Washington, DC, USA
[d]Science Applications International Corporation, San Diego, CA, USA
[e]John Wreathall and Co., Dublin, OH, USA
Abstract
An expert elicitation approach has been developed to estimate probabilities for unsafe human actions (UAs) based on error-forcing contexts (EFCs) identified through the ATHEANA (A Technique for Human Event Analysis) search process. The expert elicitation approach integrates the knowledge of informed analysts to quantify UAs and treat uncertainty ('quantification-including-uncertainty'). The analysis focuses on (a) the probabilistic risk assessment (PRA) sequence EFCs for which the UAs are being assessed, (b) the knowledge and experience of analysts (who should include trainers, operations staff, and PRA/human reliability analysis experts), and (c) facilitated translation of information into probabilities useful for PRA purposes. Rather than simply asking the analysts their opinion about failure probabilities, the approach emphasizes asking the analysts what experience and information they have that is relevant to the probability of failure. The facilitator then leads the group in combining the different kinds of information into a consensus probability distribution. This paper describes the expert elicitation process, presents its technical basis, and discusses the controls that are exercised to use it appropriately.
The paper also points out the strengths and weaknesses of the approach and how it can be improved. Specifically, it describes how generalized contextually anchored probabilities (GCAPs) can be developed to serve as reference points for estimates of the likelihood of UAs and their distributions.

Keywords: Human reliability analysis; HRA; ATHEANA; Probabilistic risk assessment; PRA; Uncertainty; Human performance; Expert elicitation

## 1. Introduction

A Technique for Human Event Analysis (ATHEANA) [1] is a human reliability analysis (HRA) method that was developed by the US Nuclear Regulatory Commission (USNRC) to increase the degree to which an HRA can represent the kinds of human behaviors seen in accidents and near-miss events at nuclear power plants and at facilities in other industries that involve broadly similar kinds of human/system interactions. The method provides a detailed search process for identifying important human actions and the contexts that can lead to their success or failure. However, an accepted model of human behavior suitable for formally supporting the quantification of human actions does not exist. While ATHEANA also provides guidance for quantifying human actions for PRA purposes, the final steps of the quantification process (as presented in NUREG-1624 [1]) suggest that analysts translate the important contextual information identified with the search process into human error probabilities (HEPs) using existing HRA methods such as THERP [2].

The problem with this approach (from our perspective) is that existing quantification methods are not adequately structured to guide appropriate incorporation of the broad contextual information identified using ATHEANA, and therefore significant judgment must be exercised by the analysts performing the quantification. In fact, a significant amount of creativity and insight on the part of the analysts would be necessary to use existing HRA quantification methods to address the error-forcing contexts (EFCs) identified using ATHEANA. As a result, the originators of ATHEANA have recently adopted a facilitator-led group consensus expert elicitation approach

0951-8320/$ - see front matter q 2003 Published by Elsevier Ltd.
doi:10.1016/j.ress.2003.09.011

Reliability Engineering and System Safety 83 (2004) 207–220
www.elsevier.com/locate/ress
* Corresponding author. Tel.: þ505-844-0578; fax: þ505-844-2829.
E-mail address: jafores@sandia.gov (J. Forester).

for quantifying human actions and treating uncertainty (quantification-including-uncertainty[1]).

This paper describes the expert elicitation process, presents its technical basis, and discusses the controls that should be exercised to use it appropriately. The paper also points out strengths and weaknesses of the approach and how it can be improved. Specifically, it describes how generalized contextually anchored probabilities (GCAPs) can be developed to serve as reference points for estimates of the likelihood of unsafe actions (UAs) and their distributions.

## 2. Description of the quantification process

The basic formulation of the quantification process is described first, followed by a more detailed discussion of the critical elements of the process and how it is currently implemented.

### 2.1. Basic formulation

Quantification as part of HRA involves the derivation of a probability distribution for basic events modeled in a probabilistic risk assessment (PRA), referred to as human failure events (HFEs). Each HFE comprises one or more UAs, and each UA could occur under one or more EFCs[2]. In the simplest case of one UA for the HFE modeled in the PRA accident scenario (S), the UA can still operate under several EFCs, so quantification of the HFE is calculated as:

$$P(\mathrm{HFE}\,|\,S) = \sum_i P(\mathrm{EFC}_i\,|\,S)\, P(\mathrm{UA}\,|\,\mathrm{EFC}_i; S)$$

This equation does not imply a mechanistic calculation of P(HFE|S). Rather, it alerts us the need to examine a wide range of EFCs, given a particular UA associated with an accident scenario (i.e. S). That is, a range of possible plant conditions and 'levels' of performance shaping factors (PSFs) could be consistent with a particular accident scenario represented in the PRA.

As can be seen from this equation, quantification in ATHEANA is a two-step process: (1) quantification of the EFC (plant conditions and performance shaping factors (PSFs or 'human-related conditions')); and (2) quantification of the UA, given the EFC. Iterations may be required to ensure that EFCs potentially significant to risk are identified. The probability of a plant condition can generally be quantified using a standard PRA processes. Typically, the plant conditions are given by the PRA model, and the assessment of the probability of those conditions is not directly an HRA issue. However, when the HRA search process identifies variations on the nominal (or most expected) PRA conditions that should be considered, the probability of those conditions must be determined. The associated human-related conditions, i.e. the PSFs, are specified for the UA in question by the search process for plant conditions and plant-specific factors that can affect human performance.

The strength of the effect of a PSF is a matter of the overall EFC. In other words, a PSF is identified as being relevant to the occurrence of a UA or not, depending on the overall EFC in which the potential for a UA is being examined. If it is judged that potential variations in a PSF will not affect the probability distribution of a UA (or the uncertainty associated with its occurrence) given the particular set of plant conditions and other PSFs thought

to be relevant, then it is simply dropped from consideration. When it is thought to be relevant, the degree of its contribution to the likelihood of a UA will be assessed in the overall plant-specific context in which the operating crew functions.

Some of the major differences of ATHEANA from earlier methods are that ATHEANA probes more deeply into plant conditions and plant-specific human influences; it evaluates each UA over a range of possible EFCs, rather than assuming a single, nominal EFC; and it does not require the analysts to deal with only a limited set of PSFs, to evaluate them by 'rating' their influence, and to assume that they are independent. In ATHEANA, PSFs are evaluated in an integrated manner: The potential impact of the PSFs is evaluated in the context of the other PSFs and plant conditions identified as important. Analysts strive to consider potential interactions among the set of factors that might strengthen, weaken, or eliminate the usual influence of individual factors. Such effects can occur in several different ways. In some cases the strong effect of one or several factors essentially make the effects of other factors logically irrelevant. For example, workload may be high in terms of the number of tasks (task load) that needs to be performed, but if the crew trusts their procedures and insists on a steady, methodical application of the procedures, then workload (at least in terms of task load) as a PSF is irrelevant.

Other types of interactions that may be important are those commonly seen in experimental research where the effects of one factor are altered by the presence of another factor. An obvious example is that stress-related anxiety may facilitate operator performance on an easy, mundane task, but hinder performance on a complex task. However, in some contexts, stress-related anxiety may also facilitate performance on complex tasks by inducing individuals to use various problem solving 'heuristics' that permit more

[1] This paper uses the phrase 'quantification-including-uncertainty' as interim language, because the two words have acquired separate meanings to many members of the risk and safety community. In fact, quantification includes uncertainty, because anything else would be incomplete. The only generally defensible point estimate is the one derived from the uncertainty distribution of the final result.

[2] Note that in the ATHEANA methodology [1], the term error-forcing context (EFC) is defined as the situation that arises when particular combinations of performance shaping factors and plant conditions create an environment in which unsafe actions are more likely to occur.



efficient information processing [3]. While empirical data supporting potential interaction effects that may be important in power plant control rooms are still limited, in the expert elicitation approach being described, the analysts attempt to examine such relationships among important factors and take them into account during quantification.

In performing quantification, uncertainty is addressed by identifying the key factors (plant conditions and PSFs) that can affect the results in the case being examined. Important factors identified from the search process are considered first. Next, a relatively long list of factors developed to support the quantification process (currently about 50 items) is examined to ensure that all potentially important influencing factors have been considered, especially to determine those with a potentially strong effect on uncertainty. The general approach to uncertainty is Bayesian: It treats probability as a representation of the analysis team's state of knowledge (subjective probability), it incorporates aleatory and epistemic uncertainty, and it

brings all available evidence to bear on the issues at stake.

2.2. The expert elicitation process

2.2.1. Identification of contextual information

Application of the process makes use of a facilitator to present an initial description of the event being analyzed and the expected or proposed EFC for the event. This initial description includes as many as possible of the following steps of the ATHEANA search process [1]:

1. From Step 1 (Define and Interpret the Issue) and Step 2 (Define the Scope of the Analysis):
- the type of initiating event
- necessary combinations of systems or component failures and any previous human actions or failures that have occurred in the scenario, i.e., S
- expected relevant indications, alarms, and compelling signals available to the crew.

2. From Step 3 (Describe the Base Case Scenario):
- expected behavior of critical plant parameters
- procedures and procedure steps indicated by the behavior of the plant parameters, including continuous action statements that allow operators to respond at any time to particular situations
- event priority in terms of how important the crew would view the needed action
- crew training on relevant scenarios.

3. From Step 4 (Define HFEs and/or UAs), Step 5 (Identify Potential Vulnerabilities in Operators' Knowledge Base), and Step 6 (Search for Deviations from the Base Case Scenario):
- informal rules that could contribute to UAs given the conditions (that is, rules or tendencies exist that are not explicitly part of the relevant written procedures operators have come to follow and that are usually correct, but which in some contexts could lead to inappropriate actions, e.g., a focus on core damage under conditions where pressurized thermal shock is actually a greater concern or more simple biases such as "protect the pumps")
- timing of the scenario including the time at which critical indicators should occur and time when critical actions should be completed
- expected workload during the relevant time interval, i.e. how many others tasks and evaluations (including unrelated tasks) must be completed in the time interval and how many staff members would be available to perform the tasks
- potential distractions
- important PSFs associated with the EFC identified during the search process
- other findings from the identification of vulnerabilities and search for potential variations or deviations in the accident scenarios that might cause the crew problems.

4. From ATHEANA Section 7 (Preparation for Applying ATHEANA):
- results of observations of relevant simulator exercises, including the pacing of procedure implementation (will the crews general approach to procedure implementation, given constraints such as workload, allow them to reach important procedural steps in time), crew communication and coordination, and other crew dynamics that could influence the outcome of the scenario
- management and organizational factors that could

influence the outcome of the scenario (as identified from discussions with plant personnel regarding management and organizational factors that might influence their response to the scenario and inferred from their behavior in the simulator)
- any historical experiences of similar events at the plant (successes or failures) or in the simulator.

Ideally, the facilitator presents all of the above information to the analysts participating in the quantification process. In most cases, at least some of the analysts will have participated in the ATHEANA search process and will already be somewhat familiar with the relevant information. However, in practice it may be difficult for the facilitator to compile all this information for each event. This is one reason why it is important that the analysis team includes experienced personnel from training, operations, and the PRA/HRA group. With people knowledgeable in these areas, much of the above information can be obtained during initial discussions of the event being analyzed (i.e. all team members can contribute to the identification of the event EFC). As noted earlier, the process is one of investigating evidence identified by knowledgeable analysts and therefore the identification of EFC should be seen as an iterative process where the analysts develop a clear understanding through presentations and discussions of relevant information.



Once the above information has been established for the UAs and HFE being analyzed, the separate list of about 50 factors, which includes plant conditions and PSFs, is reviewed to help ensure that all potential influences on quantification-including-uncertainty are considered. The list was structured around PRA and plant operations topics. Although it appears to be reasonably complete, additions and refinements may take place throughout the evaluation process. Note that, in spite of the fact that it appears to be a long list, its use in recent applications has been relatively easy. Experience shows that most of the important factors are identified during the search phase of the work and that a quick review of the remaining factors in the list can provide a useful check for completeness.

#### 2.2.2. Translating the contextual information into probability distributions for HFEs

Once the scenario-relevant information and apparent important influences on human performance have been established, the analysis team begins to investigate the relationships among the factors and to determine how the factors (when considered together) would be expected to influence the likelihood of a UA. The discussions have several initial objectives. First, the analysts, with the help of the facilitator, discuss the expected ‘strength of the signature’ of the event for operating crews at the plant. In other words, given the indications, alarms, signals, and identifiable plant behavior that will occur because of the scenario, given the crews’ training and general familiarity with the event, and given the existing procedures relevant to the event, to what extent will the pattern of plant behavior clearly identify the event and needed action to the crews? The judgment is whether the pattern of cues will have significant meaning to the crews and thereby facilitate their diagnosis of the situation and needed response. A strong signature may exist even when some instrumentation is modeled as failed, if there are strong alternative indications supporting a correct

diagnosis.
Next, the analysts discuss the ability of the procedures to lead the crew to the correct response even if the signature of the event is not necessarily strong. Given the symptombased procedures used in plants today, in most cases the behavior of individual parameters can lead operators to take appropriate actions (or not take inappropriate actions) even when the 'meaning' of the combined available information may not be obvious. Even when procedures appear strong, however, analysts must strive to ensure no reasonable variations in the range of scenario conditions for which the symptom-based procedures might not be perfectly matched. For cases where the procedure may be vulnerable, the potential for correct diagnosis may be the most important factor. Alternatively, if the relevant procedures are robust with respect to the needed action given significant variation in the behavior of parameters, then familiarity with the event may be less important.
Other important aspects considered in these initial discussions include several of the factors mentioned above, such as the timing of the scenario, the tasks, and evaluations (including unrelated tasks) that must be completed in the available time; the staff available to perform the tasks; potential distractions; the pacing of procedure implementation (will the crews' general approach to procedure implementation allow them to reach important procedural steps in time); and aspects of crew communication and coordination that could influence the outcome of the scenario. It is believed that all of these factors have the potential to strongly influence the ability of the crews to respond appropriately. Thus, in this step of the process, judgments with respect to their contribution to the likelihood of failure or success need to be made and discussed by the analysts. This is not to say that other factors will not have strong influences. However, by addressing the initial objectives and factors listed (in the context of the results of the ATHEANA search process), factors likely to be critical are discussed in an integrated fashion. For example, workload is neither independent of how strong the meaning of the pattern of cues is, nor is it independent of the usual pacing of procedure implementation. Considering the relevant factors in an integrated manner will result in a reasonable understanding of the functional EFC. Furthermore, other potentially relevant factors can be considered in the context of these initial evaluations.
Once the initial discussions regarding the factors likely to influence success have been held, the analysis team identifies the key or driving factors of the plant-specific EFC expected to influence the crew with respect to the UA being quantified and any other factors that could significantly contribute to uncertainty about the probability of the UA. Each analyst is then asked to independently develop an uncertainty distribution for the UA probability. To help control for anchoring bias (see Section 2.3), they are asked to identify the 1st, 10th, 25th, 50th, 75th, 90th, and 99th percentiles of the distribution while accounting for all identified sources of uncertainty. They are encouraged to begin by asking what the worst case for the probability of failure would be (this is interpreted as the 99th percentile). In determining the worst case, they are to:
† consider all of the 'bad' factors that could occur randomly (i.e. not explicitly part of the identified EFC, but not excluded by its definition either) and that have been identified as being reasonably credible (e.g. the

event occurs in the middle of the night, the weakest crew is on shift, etc.)
† assume that the negative effects of all important factors are at their strongest (e.g. worst case for possible level of workload).

They are next asked to determine what the 'best' case for the probability of failure would be (this is interpreted as the 1st percentile). For example, they might assume that the best



crew is on shift, it is daytime, there are few distractions, 'ideal' control room conditions exist (e.g. all instruments are working), the behavior of critical parameters matches exactly those experienced during training, etc. Which factors are actually considered will be those that have already been identified as likely contributors to the probability of the action and its associated uncertainty.

With the 1st and 99th percentiles identified, the analysts are asked to identify the remaining values:

† the median (50th percentile) should be selected such that it is equally likely that the true value is higher or lower than this point
† the upper quartile (75th percentile) should be selected such that, if the true value is above the median, it is equally likely to be above or below this value
† the lower quartile (25th percentile) should be selected such that, if the true value is below the median, it is equally likely to be above or below this value
† the 90th percentile, another high-end value, should be selected such that there is a 90% chance that the true value is no higher than this value (i.e., the evaluator should be willing to offer 9 to 1 odds that the true value is less)
† the 10th percentile, another low-end value, should be selected such that there is only a 10% chance that the true value is lower

As noted above, the development of the distribution is done independently and each analyst is encouraged to draw the distribution they are developing. With seven percentiles selected, each analyst can simply 'fair in' a curve through these seven points. As described later, the facilitator should pose questions back to the analyst, based on this probability distribution, to test whether it truly describes the analyst's intent.

To assist the analysts (who may not have strong backgrounds in probability) in making their judgments regarding the probability of events, some basic guidance is provided. In thinking about what a particular probability for a UA will be, they are encouraged to try to imagine how many times out of 10, 100, 1000, etc. would they expect crews to commit the UA, given the identified EFC. The following examples of what different probabilities mean are provided to the analysts:

† 'Likely' to fail ,0.5 (5 out of 10 would fail)
† 'Infrequently fails' ,0.1 (1 out of 10 would fail)
† 'Unlikely' to fail ,0.01 (1 out of 100 would fail)
† 'Extremely unlikely' to fail ,0.001 (1 out of 1000 would fail)

The analysts are allowed to select any values to represent the probability of the UA. That is, other values (e.g. 3E-2, 5E-3) can be used. However, the analyst must provide numeric probabilities. The qualitative descriptions above are provided initially to give analysts a simple notion of what a particular probability means.

After all analysts have developed their distributions, they

each provide the distributions along with an explanation of the basis for their estimates to the group. They describe what factors are thought to be key contributors, the evidence used, and why the resulting distribution would be appropriate. The facilitator attempts to draw all distributions being reported so that differences in the shapes of the distributions can be discussed. He also tests for bias and reasonableness by probing the team members on the basis for their distributions and by ensuring that no one's ideas are ignored or easily dismissed. Other approaches for controlling bias and reasonableness are discussed in Section 2.4. Finally, all analysts work toward a consensus (final) distribution.

Preliminary applications of this approach in three plantspecific PRAs examining the risk of pressurized thermal shock have appeared to work well and generate reasonable results. They have worked well in the sense that, given a careful explanation of each UA and associated EFC and given a thorough sharing of information relevant to the issue, the analysts have developed reasonably consistent estimates and found it easy to agree on consensus distributions. The time involved in the assessment seemed reasonable to both government contractors and utility PRA managers. The results are reasonable in the sense that, when the analysts were asked to review propositions implied by their probability distributions, they generally confirmed that they agreed with those propositions. Also, in cases of disagreement, the analysts were able to identify and explain them in terms of their underlying assumptions/understanding and reach consensus.

2.3. Basis for using expert elicitation process

We have elected to use an approach for quantification based on the elicitation of consensus expert judgment for two reasons: (1) there are neither identified quantitative data directly appropriate to the specific cases being addressed nor accepted explicit phenomenological models of human performance, and (2) the group consensus approach provides a reasonable means of quantifying situations where a broad range of indirect evidence exists and formal models for treating this evidence are lacking.

The data problem is associated with the nature of human performance, particularly in the presence of strong EFC, and with the traditional assumptions made in quantifying human actions in PRA accident scenarios. It can be argued that the traditional model for quantifying human actions in PRAs has focused on the 'best case' conditions for the accident scenario being examined. For example, it has traditionally been assumed that the only equipment failures are the one or two explicitly involved in the PRA accident sequence (i.e. S) and that the behavior of the critical plant parameters will explicitly match the crews' expectations (e.g.



the parameters behave as in the traditional accident scenarios on which they have been trained). From the perspective of Hollnagel's argument [4] that there are multiple modes of human performance driven by context and control, the 'traditional' model can be seen as examining regimes with usually mild EFC and good crew control. Here one might claim (as does THERP [2]) that nominal failure rates for well-defined tasks have been clearly described by traditional task analysis and these nominal failure rates can be established through data analysis (or in principle by experiment). In this mode, one might further claim that the effect of mild EFC can be

viewed as a multiplier, a modifier of the nominal failure rate, and that this multiplier is a function of nominally independent PSFs. In this case, the work focuses on establishing these nominal failure rates and the 'delta' due to the PSFs. However, it can also be argued that the nominal or 'best case' scenarios may not always be the most frequent and that there may be more important and troublesome cases for which such data would not be appropriate and where other applicable data have not been identified or developed. For instance, the assumption that the only equipment failures in an accident sequence are those explicitly modeled could be wrong. This assumption is fostered by the idea that additional failures reduce the probability of the event. But if we consider that a plant has thousands of components, the most likely state is one with several failures, even though they may not have a direct effect on the crew (see Ref. [1], pp E1-E4 for a simple illustration calculation of this concept). Furthermore, as is argued and documented in ATHEANA [1], multiple effects and variations can be associated with a particular accident scenario that can confuse and hinder the crew. Thus, in reality a variety of 'near-nominal' conditions are quite likely. In such cases, EFC is strong and unexpected, and crew control can begin to slip away. Here the traditional approach of assuming a nominal failure rate and making small adjustments for a well-defined task may break down. The important issue becomes loss of crew control.Nowthe failure rate is no longer related to the nominal or best case, and the crew is not embarked on a well-defined task. It no longer matters what the task is; the mental condition makes any task difficult. In fact, just deciding what task to carry out can be daunting. Here the PSFs are strong; they act in concert, potentially producing interactive effects that are far stronger and different than their independent assessment would indicate. The consensus expert judgment approach helps to deal with these types of situations and other non-nominal cases such aswhenthere aremismatches between the scenario 'expected' by the operating crews and the actual evolution of the accident (e.g. Reason's latent pathway effects [5] or situations where the usual pacing of procedure implementation does not match that demanded by the scenario).

The second issue of a group consensus expert elicitation process has been clearly addressed by the Senior Seismic Hazard Analysis Committee, authors of the 'SSHAC' report [6]. In Appendix J of that report, the authors provide a nice comparison of mathematical and behavioral schemes for aggregation of information from multiple experts. Both approaches have been used in HRA. SSHAC cites the following advantages for mathematical aggregation:

† logic is clear and can be reviewed and tested
† mathematical formulae can separate 'assessments of dependence, expertise, and overlap, so that sensitivity studies are straightforward.'

SSHAC then notes that the current state of the art presents serious disadvantages:

† no credible mathematical models for aggregation include all the important factors
† current alternative models are not fully applicable for all cases of interest and ignore important dependence effects
† no single objective model can fit all cases.

The alternative, 'behavioral' approaches seek some type of consensus. There are a number of behavioral approaches in these literature, such as Delphi methods [7]. We prefer the expert 'information' focused group interaction [8] as

described in Section 2.2 above. SSHAC points out that the primary advantage of the group consensus approach is that, 'if the information exchange is full and unbiased, and if the result truly reflects each expert's state of information, then the consensus result is credible and non-controversial' ([6], p. 33). The primary concerns include:

† the result may not be a true consensus reflecting the combined expertise and experience of the group, but some negotiated position
† strong personalities may influence the result
† the group has limited the discussion, which results in uncertainty being understated.

The formalism of structured expert elicitation (e.g. see the SSHAC report [6] and a generalized paper [9]) permits us to address the concerns raised above. Such formalism can also address the highly complex and interdependent conditions involved in the evaluation of human reliability under strong EFC. It can also provide good results for cases with less severe, but well-defined EFC by allowing analysts to consider a realistic set of evidence in a holistic manner. The SSHAC report offers an effective structure to make the elicitation process consistent, and we have adopted key aspects of that structure and specialized it to fit HRA.

The preceding discussion provides good reasons for using a consensus expert elicitation process or at least for not abandoning such an approach until more structured mathematical methods are developed that provide the same advantages. However, to gain the advantages of the expert evidence/consensus approach, a strong facilitator, or set of strong analysts, who understands the process and enforces



a formal, structured interaction, is required. As discussed earlier in this paper, each analyst is required to develop their distribution independently and to defend their position with all the evidence of which they are aware. No one is allowed 'off-the-hook' (i.e. to capitulate to another analyst's unsupported opinion). In our experience, the process levels the playing field, as everyone shares evidence and the basis for their opinions, often simplifies the issues, sometimes splits the current question into two or more related conditions, and makes it easy to reach consensus.

### 2.4. Controls for unintentional bias

One of the most important concerns associated with the use of a consensus expert judgment process is that of unintentional bias. In the subjective process of developing probability distributions, strong controls are needed to prevent bias from distorting the results (i.e. to prevent results that don't reflect the team's state of knowledge). Perhaps the best approach is to thoroughly understand how unintended bias can occur. With that knowledge, the facilitator and team can guard against its influence in their deliberations. A number of issues need to be considered, as discussed briefly below.

A number of studies present substantial evidence that people (both naive analysts and subject matter (domain) experts) are not naturally good at estimating probability (including uncertainty in the form of probability distributions or variance) [10–12]. For example, Hogarth [10] notes that psychologists conclude that man has only limited information processing capacity. This in turn implies that his perception of information is selective, that he must apply heuristics and cognitive simplification mechanisms, and that he processes information in a sequential fashion. These characteristics, in turn, often lead to a number of

problems in assessing subjective probability. Evaluators often:
† ignore uncertainty (this is a simplification mechanism); uncertainty is uncomfortable and complicating, and beyond most people's training.
† lack an understanding of the impact of sample size on uncertainty. Domain experts often give more credit to their experience than it deserves (e.g. if they have not seen it happen in 20 years, they may assume it cannot happen or that it is much more unlikely than once in 20 years).
† lack an understanding or fail to think hard enough about independence and dependence.
† have a need to structure the situation, which leads people to imagine patterns, even when there are none.
† are fairly accurate at judging central tendency, especially the mode, but tend to significantly underestimate the range of uncertainty (e.g. in half the cases, people's estimates of the 98% intervals fail to include the true values).
† are influenced by beliefs of colleagues and by preconceptions and emotions.
† rely on a number of heuristics to simplify the process of assessing probability distributions; some of these introduce bias into the assessment process.
Examples of this last area include:
† Representativeness. People assess probabilities by the degree to which they view a known proposition as representative of a new one. Thus stereotypes and snap judgments can influence their assessment. In addition, representativeness also ignores the prior probability [13] (i.e. what their initial judgment of the probability of the new proposition would be, before considering the new evidence—in this case their assumption of the representativeness of the known proposition). Clearly the prior should have an impact on the posterior probability, but basing our judgment on similarity alone ignores that point. This also implies that representativeness is insensitive to sample size (since they jump to a final conclusion, based on an assumption of similarity alone).
† Availability. People assess the probability of an event by the ease with which instances can be recalled. This availability of the information is confused with its occurrence rate. Several associated biases have been observed:
– biases from the retrievability of instancesrecency, familiarity, and salience
– biases from the effectiveness of a search set-the mode of search may affect the ability to recall
– biases of imaginability-the ease of constructing inferences is not always connected with the probability
† Anchoring and Adjustment. People start with an initial value and adjust it to account for other factors affecting the analysis. The problem is that it appears to be difficult to make appropriate adjustments. It is easy to imagine being locked to one's initial estimate, but anchoring is much more sinister than that alone. A number of experiments have shown that even when the initial estimates are totally arbitrary, and represented as such to the participants, the effect is strong. Two groups are each told that a starting point is picked randomly just to have a place to work from. The one given the higher arbitrary starting point generates higher probability. One technique found to be helpful is to develop estimates for the upper and

lower bounds before addressing most likely values.
Lest we agree prematurely that people are irretrievably poor at generating subjective estimates of probability, it is significant to realize that many applications have been successful. Hogarth [10] points out that studies of experienced meteorologists have shown excellent agreement with actual facts. Thus, an understanding is needed of what techniques can help make good assessments. In addition, in his comments published with the Hogarth

paper [10], Edwards observes that humans use tools in all tasks, and tools can help us do a very good job in the elicitation process.
Winkler and Murphy [14] make a useful distinction between two kinds of expertise or 'goodness.' 'Substantive' expertise refers toknowledge of the subject matterof concern. 'Normative' expertise is the ability to express opinions in probabilistic form. Hogarth [10] points out that the subjects in most of the studies were neither substantive nor normative experts. A number of studies have shown that normative experts (whose domain knowledge is critical) can generate appropriate probability distributions, but that substantive experts require significant training and experience, or assistance (such as provided with a facilitator), to do well.
By understanding how these inadequacies and biases occur, the information can be used to combat their influences. The inadequacies of individuals can be dealt with by selecting analysts with a variety of expertise and by facilitating the process, challenging participants to explain the basis for their judgments. Biases can be directly addressed by a facilitator. For example, representativeness bias involves ignoring available information and replacing a careful evaluation of that information with quick conclusions based on an over-focus on part of the information or allowing irrelevant information to affect conclusions. The facilitator must challenge analysts, asking them to explain their opinions. The facilitator must use his own judgment to sense when an individual is not using the full information.
Moreover, by understanding the heuristics that people generally use to develop subjective probability distributions and the biases that attend those techniques, that awareness can help analysts avoid the same traps. Through understanding which framings for eliciting distributions cause problems, we can use those that work better. Because the facilitator is familiar with the potential biases, she can test the group's ideas and push them in the right direction. The strategies presented below should be used either explicitly or implicitly through the questioning of the facilitator, as described in the SSHAC report [6]. In addition, Tversky and Kahneman [12] give many detailed examples useful for helping facilitators develop awareness of such useful aids.
Some of the simplest and best aids include:
† constructing simple models of the maximum and minimum points of the distribution, avoiding focus on the central tendency until the end points are studied to avoid anchoring; test these models to examine the evidence supporting them rather than relying on opinion alone.
† seeking consensus on the evidence considered by the analysis team.
† testing distributions by asking if the assessor agrees it is equally likely for the real answer to lie between the 25th to 75th percentiles or outside them; or between the 40th to 60th percentiles and outside the 10th and 90th percentiles. Sometimes these questions must be phrased in ways to

avoid suggesting the answer.
† establishing a strong facilitator who ensures each participant must individually put his evidence on the table and justify it [6]. The facilitator must use his judgment on when to push the participants, rather than going through a long and tedious checklist.
† being careful when assessing parameters that are not directly observable. The distribution is supposed to reflect the analyst's evidence concerning a particular parameter. If the analyst has little direct experience with the parameter, it can be difficult to justify an informative prior distribution.

In general, the process described in Section 2.2 employs techniques designed to minimize bias and enhance the development of all available information. Our experience indicates that once the full story is developed and shared, qualified analysts (e.g. a team consisting of operators, trainers, emergency procedure developers, and PRA/HRA personnel) provide reasonably consistent judgments. They almost always see the merit of arguments offered by their colleagues and can reach agreement on a consensus distribution that gives weight to each argument; this means that they also can reach agreement on the relative strength of those arguments and its supporting evidence.

2.5. Issues associated with adequate treatment of uncertainty

A formal approach for the treatment of uncertainty has been developed for use with the expert elicitation process and is consistent with recent USNRC emphasis on improving the treatment of uncertainty in PRA [15]. In particular, the approach emphasizes that all factors that can affect quantification-including-uncertainty should be included (at least conceptually). Of these many factors, those that have the potential to affect quantificationincluding-uncertainty must be examined in detail (as noted in the description of the process). This examination involves a determination of which of the three cases vis-a`-vis uncertainty apply:

† Deterministic Case—When there is no variability or there is no imperfect state-of-knowledge that leads to variability in the results.
† Aleatory uncertainty—When there is random variability in any of the factors that lead to variability in the results.
† Epistemic uncertainty—When the state of knowledge about the effects of specific factors is less than perfect.

A more operational point of view is that uncertainty is aleatory if:

† it is (or is modeled as) irreducible or
† the uncertainty is observable (i.e. repeated trials yield different results) or



† repeated trials of an idealized thought experiment will lead to a distribution of outcomes for the variable and thus this distribution is a measure of the aleatory uncertainties in the variable.

The uncertainty is epistemic if:

† we are dealing with uncertainties in a deterministic variable whose true value is unknown or
† repeated trials of a thought experiment involving the variable will result in a single outcome, the true value of the variable, or
† it is reducible (at least in principle).

The approach for the treatment of uncertainty

implements the subjective framework for treating probabilities in PRA described by Apostolakis [16]. By implementing the subjective framework, the approach helps to ensure consistency between the HRA (in the interpretation of the HFE probabilities, including uncertainties) and the rest of the PRA. An additional benefit is a clearer (and potentially simplified) elicitation-based quantification process. This benefit arises from the subjective framework's distinction between aleatory and epistemic uncertainties, which requires a careful examination of the factors contributing to P(HFE) and its uncertainty, resulting in a clearer definition of the issues being addressed during the elicitation.

In the subjective framework, P(EFC$_i$|S), P(UA|EFC$_i$,S), and P(HFE|S) are measures of aleatory uncertainty—they quantify the uncertainties in the occurrence of observable events (the occurrence of EFC$_i$; the occurrence of the UA, given EFC$_i$; the occurrence of the HFE) arising from random variability. These terms are conditioned on the occurrence of a particular scenario, S (typically a set of successes and failures of PRA basic events, which represents a bundle of physical sub-scenarios), leading up to the operator action modeled by the HFE. The degree of epistemic (or 'state of knowledge') uncertainty in the values of these terms varies according to how the EFCs and UAs are specifically defined.

In our current implementation of the subjective framework, the EFCs and UAs are defined in such a manner that the epistemic uncertainties in P(HFE|S) arise primarily from the P(UA|EFC$_i$,S) terms. Thus, the EFC definition includes deterministic factors (e.g. what specific actions are dictated by a particular procedure when a particular set of cues is recognized) and aleatory factors whose epistemic uncertainties either appear to be small (e.g. the time of day when the scenario occurs) or can be assessed in a relatively straightforward manner (e.g. the state of hardware not explicitly modeled by the PRA). It can be seen that this definition relegates the substantial epistemic uncertainties in operator performance (e.g. the effect of faulty or missing plant indications, the effect of time of day) to the P(UA|EFC$_i$,S) term. These uncertainties are addressed through the quantification-including-uncertainty elicitation process described earlier.

The formal separation of the P(EFC$_i$s|S) and P(UA|EFC$_i$,S) terms aids the assessment of epistemic uncertainties in P(UA|EFC$_i$,S) in two ways. First, it narrows the range of situations to be considered in a particular elicitation. For example, the elicitation can be conditioned on a particular time of day, so variability in this factor is eliminated from the analysts' consideration. Second, because the elicitation process requires that each analyst provide an explicit basis for his/her judgment, the separation reduces the possibility of different analysts assuming different underlying conditions when developing their judgments, which can lead to different assessments of 'best' and 'worst' cases.

Of course, a high degree of detail in the definition of the EFCs will require a large number of evaluations for P(UA|EFC$_i$,S). When time or budget is not available for these multiple evaluations, it is necessary to evaluate all the epistemic uncertainty in P(UA|EFC$_i$,S) at one time. However, in addition to losing the benefits discussed above, such an approach has a potentially significant additional conceptual penalty: analysts need to mix

inherently (from the standpoint of the PRA model) aleatory conditions (which only occur some of the time) with epistemic issues (which can be time independent). This has practical implications: participating analysts have been able to identify the end points of the P(UAlEFCi,S) distributions with a reasonable degree of comfort, but some confusion is added to the process of filling in the rest of the cumulative distribution.

3. Additional formalisms: can they help?

The expert elicitation approach to quantification provides a reasonable means to appropriately consider an EFC and translate the impact of the EFC into an uncertainty distribution of the likelihood of the UA. The approach provides a structure and guides analysts in identifying and considering the EFCs for UAs in a way that strives to facilitate the derivation of realistic quantitative estimates (i.e. estimates based on an understanding of the actual EFC driving the crew's response), including uncertainty, and permits incorporation of all sources of information, including partially-relevant data. It encourages the analyst to consider all sources of information (including data that are only relevant to some aspects of the problem) in developing estimates. However, as discussed above, in spite of the way in which the investigation of expert evidence is conducted and the striving for consensus in the uncertainty distributions derived by the analysts, potential limitations will always exist in any approach that relies on individuals to directly transform qualitative or even semi-quantitative information into probability distributions. Thus, while no



solution to the problem is completely satisfying, steps can be taken to strengthen aspects of the existing process.

It can be argued that the current expert elicitation process addresses several desirable criteria or 'desiderata':

† Considers all elements of context
† Considers all aspects of uncertainty
† Incorporates all relevant information (including data)
† Ensures intra- and inter-rater reliability
† Provides traceability
† Guards against bias.

However, further improvement of the process is needed. For example, questions could arise about the reliability of the current elicitation approach when analysts have less experience with operations, PRA, HRA and facilitation of expert elicitation sessions. The authors hope to investigate more formal processes (possibly a more explicitly mathematical process, in spite of some of the disadvantages discussed in Section 2.3). Currently, several general approaches have been suggested to improve the existing process with respect to the desiderata.

3.1. Contextually anchored probabilities

First, an observation on Bayesian analysis. The process described earlier in Section 2, where expert elicitation is used to build a prior distribution, is essentially the first step in a complete Bayesian analysis. That is, the distributions developed in the process of Section 2 are really a direct estimate of a subjective prior for the human failure rates. If data can be developed that are relevant to the specific HFE and EFC in question, then it can be used in a Bayesian updating process to calculate the most relevant posterior distribution. One challenge is the definition of the likelihood function if the data are not of the 'r-out-of-n' type.

The expert elicitation process itself can be generally described as an extension of a Delphi technique (e.g.

Linstone and Turoff [7]), which is characterized as a structured group communication process, with added controls and a focus on expert evidence rather than expert opinion. In addition, the SSHAC report [6] highlights a distinction between the use of mathematical aggregation of expert opinion as opposed to 'behavioral schemes' for integrating expert knowledge. SSHAC argues for seeking a consensus probability distribution. The existing expert elicitation approach follows a similar practice. The quantification process seeks formalisms to strengthen our ability to meet particular desiderata, such as ensuring intraand inter-rater reliability, providing traceability, ensuring consistency over time, and providing additional guards against bias. The process has already added controls including many of the items discussed in Section 2.4 to control for bias, particularly the emphasis on the use of a 'technical facilitator/integrator' as described in the SSHAC report [6]. Thought is being given to the possibility of additional formalisms (including mathematical approaches) to strengthen the controls.

As we think of new or modified methods, we note that there is a difference between:

† a method for uncovering the mental mechanisms associated with 'human error' for the purpose of determining P(UAlEFC) (where we are now, i.e., ATHEANA) and

† an HRA calculational method that could incorporate the results of many such ATHEANA analyses into a simple-to-apply high level tool.

The former method is the expert elicitation approach as it is currently being used and, as discussed in Section 2.3, it is needed to evaluate the range of realistic and cognitively demanding situations likely to be found in actual accident scenarios. The result for each case that is generated—a probability distribution and associated description—we call a contextually anchored probability (CAP), because the ATHEANA approach is context-driven. We also note that, in the long-term, it may be possible to generate additional CAPs from other sources, for example:

† ATHEANA quantification of real operating event descriptions

† experimental results for certain actions under a limited range of EFCs

† extraction of CAPs from existing experimental data.

Once ATHEANA and the expert elicitation approach have been used to analyze and derive a great many CAPs—plant and EFC-specific situations (or suitable experiments have been performed)—we could build a library of these results. With that database in hand, it might be possible to synthesize a more direct quantification approach. Each entry in the library would include qualitative and quantitative descriptions of each situation:

† detailed plant-specific, UA-specific, and EFC-specific description of the event

† its quantification-including-uncertainty.

Both results could be used in the new alternative approaches. Fig. 1 provides an example picture of the form of these quantification-including-uncertainty results for five example cases, 'a' through 'e.'

These results are sorted and rank ordered ($u_1$ to $u_5$) in Section 3.2 to develop alternative approaches for new HRA calculational methods based on the catalog of ATHEANA results. Let us call each individual ATHEANA result (e.g. each of the five example cases in Fig. 1) a CAP. The result

of the ATHEANA process is a probability distribution anchored to the EFC of the event.



### 3.2. One path forward: development and use of generalized contextually anchored probabilities

One promising approach involves the use of the library of individual ATHEANA quantification results, the CAPs of the last section, to provide reference cases for the quantification of new events (situations). To that end, we propose development of generalized CAPs (GCAPs). These GCAPS would be selected to emphasize certain high-level characteristics of each set. Here again our thoughts are preliminary but appear to offer a path toward a simplified HRA approach.

GCAPs, then, are descriptions of classes of events that have been characterized according to the factors driving their occurrence and that have an associated probability distribution that is also strongly affected by those factors. The characterizations are consistent with the emphasis on the role of context in producing UAs. The goal is to develop generalized descriptions relevant to quantitative results of the critical characteristics of events and what caused them. GCAPs would provide a basis to:

† standardize or normalize the judgments applied in assessing the probabilities of different UAs

† use operational experience from other events as reference points in assessing the probabilities of the different UAs

† explain the assessments to peers and outside reviewers.

The GCAPs will not always provide direct ties to specific UA probabilities for all events and scenarios because of the complex interactions among UAs, plant conditions, PSFs and error mechanisms. The quantitative significance of each GCAP element would depend on the details of the specific case at hand. Thus, evaluators will have to temper the use of these quantitative anchors with the knowledge of the range of conditions or EFC over which they apply. Substantial development of GCAPs is needed to provide a broad basis for quantifying a range of EFCs. The approach would begin with a search of events and databases that might be used to develop catalogs of specific CAPs. These descriptions of actual events would later be generalized into GCAPs because it is thought that the generalized descriptions of classes of events will be more straightforward to use.

To see what the GCAP curves may look like, consider the set of CAPs shown in Fig. 1. Suppose that these all apply to a group of UAs similar enough in event signatures that we wish to consider them a GCAP group. Therefore, they will serve as a basis for an event signature-driven GCAP. Fig. 2 shows the same CAPs as Fig. 1. Now they have been reordered by probability (mean value of the distributions). Parameter u orders the individual EFCs, $u_i$. If one fairs in a continuum of EFCs between the individual $u_i$ probability distributions, one gets the dashed lines (probability surface) of Fig. 2. This surface is the event signature-driven GCAP for this case and would be used for quantification of events fitting the applicable event-signature group in Section 3.3 below.

The overall distribution is for the GCAP, under all applicable EFCs ($u_i$), and can be envisioned as a threedimensional surface rising out of the page. Each of the sketched distributions are cuts of that surface at that particular value of u.

The next question is how to select groups of CAPs by similar event signature for use as GCAPs. At the current

time, we envision partitioning the CAPs into three types of GCAPs, based on high-level characteristics of the events (or situations):

† if its quantification result was primarily from the kind of deviation or mismatch that occurred (i.e. if the EFC was very strong) such that the error (or, rather, the UA) would be very likely regardless of the specific action involved, add it to the EFC-Driven GCAP type

† if its quantification result was primarily from a lack of deviation or mismatch (i.e., if the EFC was very strongly positive in that it was familiar, it matched training and

Fig. 1. Results from Five ATHEANA Calculations.

Fig. 2. Event Signature-Driven GCAP.



emergency operating procedures, and the timing was favorable) such that the error (or, rather, the UA) would be very unlikely regardless of the specific action involved, add it to the best case or Success-Driven GCAPtype

† if its quantification result was primarily from the specific event signature (the PRA scenario, S, and the EFC) add it to the Event-Signature-Driven GCAP type; the difference here is that the GCAP is strongly affected by the specific PRA scenario (the specific initiating event and the specific sequence of plant failures and operator actions, rather than by more general pressures as in the first case above).

Note that the third Event-Signature-Driven GCAP bin will be a large one at this time. We have several ideas on how to partition it into similar signature groups, such as:

† by initiating event

† by initiating event and UA

† by severity of EFC from near-nominal to near-extreme

† by contextual characteristics (cognitive characteristics), which could include plant conditions such as instrument or support system problems, as well as human characteristics such as crew interaction and strategies.

The EFC descriptions of the individual CAPs probably need to be abstracted to flag the salient cognitive characteristics of each CAP (often these will be a direct consequence of specific unusual plant conditions). This flagging would be a first step in grouping them. We believe that the most effective grouping schemes will evolve from a combination of theoretical concepts (which remain to be developed) and examination of the EFC descriptions in the catalog of CAPs (which also remain to be developed to a meaningful extent).

When smaller sets of CAPs are selected for a GCAP, the shape of the GCAP curve may require a less complex representation. For example, curves for the EFC-Driven GCAP are expected to be presented as a simpler family of curves as illustrated in Fig. 3, which shows a set of curves for the strong EFC case.

The most severe may be a distribution that is narrow and centered near 0.5. This example was prepared based on a common EFC observed in several serious accidents as described in NUREG-1624 [1]. These accidents began with the plant operating outside its expected range of conditions (a severe deviation from operator and designer expectations (e.g. Three Mile Island-2, Chernobyl)) that became situations in which the plant’s behavior was not understood (an event that did not match the operator’s mental model), indications of the actual plant state and behavior were not recognized as cues by the operators (mindset or other biases allowed the operators to miss the significance of available

information), and prepared plans or procedures were not applicable or helpful. The example in Fig. 3 includes uncertainty where the remainder of the family are expected to be less severe cases, with the curves broadening and moving off to the left (lower values).

### 3.3. Alternative possibilities for using GCAPs in quantification

Next, a process would be needed to map GCAPs developed for classes of EFCs and UAs to new events with specific EFCs and UAs. Such a process would require us to define measures of the 'distance' between the UA-EFC sets being analyzed and specific GCAPs. In other words, because the GCAPs are being developed to assist analysts in transforming identified EFCs into quantitative estimates for new events, a process for mapping the event EFCs to GCAPs would help reduce the subjectivity of the transformation process. Some preliminary ideas on the mapping process are discussed briefly below.

Analysis begins following the usual ATHEANA search process, including the identification of the full EFC of interest. The analyst will use the new approach to simplify quantification, drawing on the library of existing results. The envisioned approach for the types of GCAPs discussed in the last section proceeds as follows. Consider the EFC of the new event to be quantified in the following manner:

† if the EFC is such that the deviation/mismatch signature is strong, then use the EFC-Driven GCAP family of probability curves (Fig. 3). The analysts must use judgment to define the strength of the deviation/-mismatch. (Exactly how this is done remains to be developed.) Then they select the associated curve as their quantification result.

† if the EFC has no significant deviation/mismatch component, then use the Success-Driven GCAP family of probability curves. Again the analysts must use judgment to define the strength of the positive EFC.

Fig. 3. Deviation/Mismatch Signature-Driven GCAP.



Then they select the associated curve as their quantification result.

† For other cases, select the GCAP with cognitive characteristics that best match the new event's EFC. The analysts must use judgment to compare their EFC with those associated with the tabulated specific-EFCs, ($u_i$) in the GCAP (i.e. they select the region between two existing $u_i$'s and read off the associated probability distribution (a cut through the probability surface)). If more than one GCAP approximately matches the new event, then analysts can select from weighted combinations of these GCAPs.

All this sounds simple in concept, but it may prove much more difficult in practice. It has not yet been implemented.

### 3.4. A real shortcut

One desired product is a method that can be used by PRA analysts who may not have all the knowledge, experience, and expertise to confidently use the expert elicitation processes currently being used with ATHEANA. For convenience, let us call this method HRA-QUANT. The previous discussion of the use of GCAPs hints at one possible approach, after the collection of CAPs into a library is extensive. In this case, one could use all the CAPs for a UA to develop probability distributions for that UA (perhaps sorted by initiating event). Such distributions

would be very broad. They might be generated as the weighted sum of each of the CAPSs weighted by their chance of occurrence. One caveat is that the library of CAPs must be large enough to ensure that the rare but highprobability-of-failure cases are represented. Presumably the less likely but severe EFCs will dominate the lowprobability high-likelihood-of-failure tails of these distributions and the more likely nominal (or best EFC) case will set the low end tail of the distribution.

If a method along the lines of HRA-QUANT can be developed, it could be used by personnel with much more limited expertise than that required for the current ATHEANA expert elicitation process. However, the uncertainty in the results would be expected to be very broad, because it would cover all possible EFCs. When the PRA/HRA results indicate the UA is a significant risk contributor using a method like HRA-QUANT, more experienced analysts should return to the database to determine what kinds of EFC are contributing to each range of the uncertainty distribution to understand the reason for the risk contribution and what might be done to improve the situation.

Note that the wide range of EFCs would include the optimistic one assumed in much previous work (i.e. the case where everything is fine except for the specific failures identified in the set of PRA sequences, $S_i$).

In closing, these ideas are just evolving, but it already appears that the approach could improve quantification for the 'near-nominal' conditions discussed above in Section 2.3. As was noted, it appears that the 'nominal' case in many 'data' sources is more representative of a best case value and this case may not be realistic. The most significant events (e.g. Three Mile Island-2 and the Browns Ferry fire) are often removed from the data because the problem in each of those events has been 'fixed.' Furthermore, we have argued that the assumption that the most likely condition of the plant is one with no failures other than the one or two involved in the particular PRA scenario, $S_i$, is wrong, because of the large number of components in a plant. Recent efforts in pressurized thermal shock HRA quantification indicate that some of the near-nominal (reasonably high-probability cases) can have significantly higher P(UAlEFC). Thus, the near-nominal cases (some very likely, some less) can fill in the remainder of the distribution between the extremes discussed above.

### 3.5. Next steps

At this time, as part of ATHEANA applications, we are beginning to collect a catalog of CAPs and giving initial thought to possibilities for generalization into GCAPs along the lines described in Section 3.2. Furthermore, the USNRC is sponsoring work (performed by the Idaho National Engineering and Environmental Laboratory) to develop an HRA data repository using information from various sources, including nuclear power plant operational experience, simulator studies, and the open psychological literature. It is anticipated that this effort will support the development of more formalized quantification methods for HRA, including the development of GCAPs.

### Acknowledgements

This work was funded by the U.S. Nuclear Regulatory Commission and performed at Sandia National Laboratories. Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DEAC04-

94AL85000. The opinions expressed in this paper are those of the authors and not of the U.S. Nuclear Regulatory Commission.

## References


[1] US Nuclear Regulatory Commission (USNRC). Technical basis and implementation guidelines for a technique for human event analysis (ATHEANA). Washington, D.C.: NUREG-1624, Rev. 1; 2000

[2] Swain AD, Guttmann HE. Handbook of human reliability analysis with emphasis on nuclear power plant applications. Washington, D.C.: NUREG/CR-1278, U.S. Nuclear Regulatory Commission, 1983

[3] Klein G. The effect of acute stressors on decision making. In: Driskell JE, Salas E, editors. Stress and human performance. Mahwah, New Jersey: Lawrence Erlbaum Associates; 1996.

J. Forester et al. / Reliability Engineering and System Safety 83 (2004) 207–220 219

[4] Hollnagel E. Cognitive reliability and error analysis method: CREAM. New York: Elsevier Science Inc.; 1998.

[5] Reason J. Managing the risks of organizational accidents. Brookfield, Vermont: Ashgate; 1997.

[6] Budnitz RJ, Apostolakis G, Boore DM, Cluff LS, Coppersmith KJ, Cornell CA, Morris PA. Recommendations for probabilistic seismic hazard analysis: guidance on uncertainty and use of experts. Washington, DC: NUREG/CR-6372, US Nuclear Regulatory Commission; 1997

[7] Linstone HA, Turoff M. The Delphi Method: techniques and applications. Reading, Massachusetts: Addison-Wesley; 1975.

[8] Bley DC, Kaplan S, Johnson DH. The strengths and limitations of PSA: where we stand. Reliability Engineering and Systems Safety 1992;38(1/2):326.

[9] Budnitz RJ, Apostolakis G, Boore DM, Cluff LS, Coppersmith KJ, Cornell CA, Morris PA. Use of technical expert panels: applications to probabilistic seismichazard analysis. Risk Analysis1998;18(4):463–9.

[10] Hogarth RM. Cognitive process and the assessment of subjective probability distributions. Journal of the American Statistical Association 1975;70(350):271–94.

[11] Kahneman D, Slovic P, Tversky A. under uncertainty: heuristics and biases. Great Britain: Cambridge University Press; 1982.

[12] Tversky A, Kahneman D. Judgment under uncertainty: heuristics and biases: biases in judgments reveal some heuristics of thinking under uncertainty. Science 1974;185:1124–31.

[13] Siu NO, Kelly DL. Bayesian parameter estimation in probabilistic risk assessment. Reliability Engineering and System Safety 1998;62: 89–116.

[14] Winkler RL, Murphy AH. ‘Good’ probability assessors. Journal of Applied Meteorology 1968;7:751–8.

[15] Siu N, Malik S, Bessette D, Woods H. Treating aleatory and epistemic uncertainties in analyses of pressurized thermal shock. In: Kondo S, Furuta K, editors. PSAM 5-Probabilistic Safety Assessment and Management: Proceedings of the 5th International Conference on Probabilistic Safety Assessment and Management held on November 27-December 1, 2000, Osaka, Japan, vol. 1/4. Tokyo: Universal Academy Press, Inc.; 2000. p. 377–82.

[16] Apostolakis G. The concept of probability in safety assessments of technological systems. Science 1990;250:1359–64.


J. Forester et al. / Reliability Engineering and System Safety 83 (2004) 207–220 220