Chapter 4

Human-Reliability Analysis

4.1 INTRODUCTION

The purpose of this chapter is to provide a procedure for estimating the probabilities of human errors in the operation of nuclear power plants. This introductory section defines the scope, assumptions, limitations and uncertainties, and the product of a human-reliability analysis (HRA). The procedure for conducting a human-reliability analysis is then outlined, highlighting the major tasks involved. The recommended method is described in Section 4.3, followed by a listing of the information requirements in Section 4.4. A detailed procedure, each step of which is illustrated by example, is presented in Section 4.5. Also included in this chapter are recommendations for documentation and the display of final results (Sections 4.6 and 4.7, respectively), a discussion of uncertainty and variability (Section 4.8), and a sample of alternative methods, their strengths, and their limitations (Section 4.9). The chapter ends with recommendations on the assurance of technical quality.

For a greater understanding of the main method presented in this chapter, the reader is urged to study the practice exercises in a recent NRC publication (Bell and Swain, 1981). Additional examples, human-performance models, and estimates of generic human-error probabilities for tasks in nuclear power plants are available in the source document for most of this chapter, the Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications,* called simply the "Handbook" in the text that follows.

4.1.1 SCOPE

The HRA methods in this chapter are intended to support probabilistic risk assessments of light-water-reactor power plants. In such an assessment, the first effort at identifying the human-related events that affect system reliability is made by the system analysts. The human-reliability analysts then determine the associated human errors that are to be defined and analyzed. Drawing from the data in the Handbook, or on better sources of data if available, these analysts then estimate probabilities for these

---

*A. D. Swain and H. E. Guttmann, Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications, draft, NUREG/CR-1278, U.S. Nuclear Regulatory Commission, October 1980. This draft will have been substantially revised by the time of its final publication in late 1982, but the authors of this chapter have attempted to keep abreast of the current revisions. It should be noted that all chapter, table, and page numbers cited here for the Handbook refer to the October 1980 draft.

system-important errors and investigate their effects on the probability of system success. Criteria for system success and failure are established by the system analysts.

In a probabilistic risk assessment, it is necessary to consider the human tasks that are performed under normal operating conditions and those performed after accidents or abnormal occurrences. In the former situation, errors might be made during or after maintenance, calibration, or testing or in the normal operation of the plant. These errors may occur in or out of the control room. In the post-accident situation, most, but not all, of the system-safety-related errors occur in the control room.

In either situation, most of the errors identified and analyzed in this guide are those made in following plant procedures (written, oral, or standard shop practice). Only occasionally are extraneous acts considered. That is, in most cases, the analyst determines whether a given response procedure is followed correctly and does not attempt to determine which uncalled-for elements are manipulated.

The HRA method recommended and most fully described in this procedures guide employs the Technique for Human Error Rate Prediction (THERP) described in the Handbook. Unless specifically stated otherwise, all qualifying assumptions and limitations apply to the alternative methods discussed in Section 4.9.

4.1.2 ASSUMPTIONS

Only human errors are dealt with--mistakes made in the performance of assigned tasks. Malevolent behavior--deliberate acts of sabotage and the like--are not considered. It is assumed that all plant personnel act in a manner they believe to be in the best interests of the plant. Any intentional deviation from standard operating procedures is made because the employee believes his method of operation to be safer, more economical, or more efficient or because he believes performance as stated in the procedure to be unnecessary.

An important aspect of a human-reliability analysis is the qualitative assessment of the sources of human error. (This calls for identifying and understanding the underlying contributors to each error and for assessing the relative importance of each of these contributors to the system-failure events being analyzed.) However, since the PRA Procedures Guide is intended for probabilistic risk assessment, this chapter deals only with the quantitative aspect. For information on qualitative application to the operations of nuclear power plants, the reader should consult the Handbook.

4.1.3 LIMITATIONS AND UNCERTAINTIES

For a complete human-reliability analysis, the risk-assessment team should include a person who is, by professional training and experience, competent in applying the techniques of human-performance analysis to complex

systems. Such a person is usually known as a human-factors specialist, an engineering psychologist, or an ergonomist. (For a more detailed description of the qualifications of a human-factors specialist, see pages 8 and 9 of NUREG-0801 (USNRC, 1981a).) To carry out the procedure described in Section 4.5 of this chapter, he must be thoroughly familiar with, and have a good understanding of, this document as well as the Handbook. For a less complete analysis (e.g., a bounding analysis) the only requirement in this respect is that the HRA analyst be familiar with this chapter and the Handbook; he need not necessarily be a human-factors specialist.

In all cases, it is presumed that the human-reliability analysis will be an integral part of the PRA project. There will be considerable and continuing interaction between those responsible for the human-reliability analysis and those working in fault-tree and system-reliability analysis. In no case should the human-reliability analyst work in isolation from the rest of the PRA team. The structure of the team should in itself facilitate the interaction necessary among the several analysts.

The major source of uncertainty in human-reliability analysis is the dearth of actuarial data on human-error probabilities (HEPs). For the most part, the Handbook presents the best available data on human performance in carrying out the tasks performed in nuclear power plants. Most of the estimates of human-error probabilities in the Handbook represent extrapolations from human-error data based on tasks performed outside, but behaviorally similar to those performed in, nuclear power plants. The tasks are behaviorally similar because they may involve the same types of cues, interpretations, response requirements, and responsibilities as those performed in nuclear power plants. Therefore, in those cases for which an analyst can find better human-performance data than those presented in the Handbook, he should use them.

It is expected that the uncertainty and speculation involved in estimating human-error probabilities for nuclear power plants will be reduced considerably in the not too distant future. Under the sponsorship of the NRC's Office of Nuclear Regulatory Research, a program plan for a human-performance data bank is being developed, and efforts are under way to collect HEP data from realistic simulator exercises for control-room tasks and from maintenance and other tasks performed outside the control room.

As explained in the Handbook, nearly all of the tabled human-error probabilities relate to routine human actions. For some operations, cognitive errors are critical (e.g., errors in evaluating display indications). There is very little information on errors of interpretation or decisionmaking (i.e., errors in the thought process). A later section (4.5.7.1) gives a general guideline for the judgments required to estimate error probabilities for post-accident decisionmaking.

The Handbook presents nominal values for the probabilities of given human actions as well as uncertainty bounds. The nominal values reflect the best estimate (based on available data and on judgment) of the probability of a particular error in a generic sense. The uncertainty bounds are considered to approximate the middle 90-percent range of the human-error probabilities to be expected under all possible scenarios for a particular action. These uncertainty bounds are based on subjective judgment rather than on actuarial data and are not meant to represent statistical confidence limits.

As discussed in the Handbook, there are several sources of uncertainty in the generic HEP values. The variability of human performance is reflected in the differences among plant personnel--differences in skill, experience, and other personal characteristics. There can be wide variation in specific environmental situations and in other physical aspects of the tasks to be performed or in the response requirements under which the operator must act. Only some of this variation in such performance-shaping factors is accounted for in the Handbook data by providing different estimates of human-error probabilities for different sets of influencing factors. The width of the uncertainty bounds surrounding each estimated nominal probability represents an attempt to account for the residual uncertainty.

Unless specifically stated otherwise, all of the probability estimates in the Handbook are based on a set of common assumptions that limit or restrict the use of the data as stated. Exceptions to these assumptions are clearly indicated. These data apply to situations in which the following hold true:

1. The plant is operating under normal conditions. There is no emergency or other state that would produce in the operators a level of stress other than the optimal.*

2. In performing the operations, the operator does not need to wear protective clothing.

3. A level of administrative control roughly equal to the average of those employed industry-wide is in effect.

4. The tasks are performed by licensed, qualified plant personnel, such as operators, maintainers, or technicians. They are assumed to be experienced--to have functioned in their present positions for at least 6 months.

5. The environment in the control room is not adverse. The levels of illumination and sound and the provisions for physical comfort are adequate even if not optimal.

The above-mentioned factors must be evaluated qualitatively for each situation being analyzed. The finding that a situation is similar to, or significantly different from, these assumed scenarios is highly judgmental. There are no absolute guidelines for establishing a plant's conformance to what is "normal" for the rest of the industry. Only with experience and exposure to several operating plants can a human-reliability analyst develop the skills necessary for performing these discriminations successfully and reliably.

---

*Most of the human-error probabilities estimated in the Handbook apply to routine human actions, often referred to as "rule-based behavior." The method for estimating the probability of human error under nonroutine (stressful) situations is unproved. Therefore, such estimates in the Handbook are characterized by wide uncertainty bounds.

It is mainly the level of detail that will differ for human-reliability analyses performed at different stages in the life cycle of a nuclear power plant. The level of detail of the procedure presented in this chapter is aimed at analyses performed for plants that are already operating. If the analysis is performed earlier (e.g., at the construction-permit stage), some of the information necessary for a detailed task analysis will not be available. Nevertheless, the procedure can still be applied as discussed in Chapter 4 of the Handbook. For analyses performed very early, much of the information needed to determine the potential for human error will have to be derived from human-reliability analyses conducted for similar plants that are already operating.

4.1.4 PRODUCT

The main result of the human-reliability analysis is, for each iteration of the analysis, a set of estimated plant- and situation-specific human-error probabilities. During quantification of the risk-significant events, these estimated human-error probabilities can be grouped into sets for incorporation into the total PRA on the basis of their effects on the reliability of a component, a whole system, or the entire response scenario required by an initiating event. The assumptions on which these sets of estimates are based are also presented to the system analysts.

4.2 OVERVIEW

Figure 4-1 shows the four phases of HRA: familiarization, qualitative assessment, quantitative assessment, and incorporation. Most HRA methods follow this general format. A block diagram illustrating the application of these phases to the procedure followed in performing a human-reliability analysis by Handbook methods is shown in Figure 4-2. The sequence of activities shown in this figure may, however, be different from that of an analysis performed in another context. Moreover, since this is a block diagram and not a flow chart of actual activities, most of the interactions between the human-reliability analyst and the rest of the PRA team are left out. This is not to suggest that they do not exist, but Figure 4-2 is meant simply to provide a schematic of the major tasks to be performed by the human-reliability analyst himself. In reading the description of these activities, it is necessary to keep in mind that the order of the various HRA activities is not a fixed one, with each activity being performed only once: the entire process is highly iterative and its parts recursive.

It is necessary to begin preparation for this analysis concurrently with the rest of the probabilistic risk assessment. Otherwise, there will not be sufficient time to perform all the activities required for an accurate assessment of the effects of human errors.

As already mentioned, a human-reliability analysis is an iterative process; various steps will be repeated as additional plant-specific or other

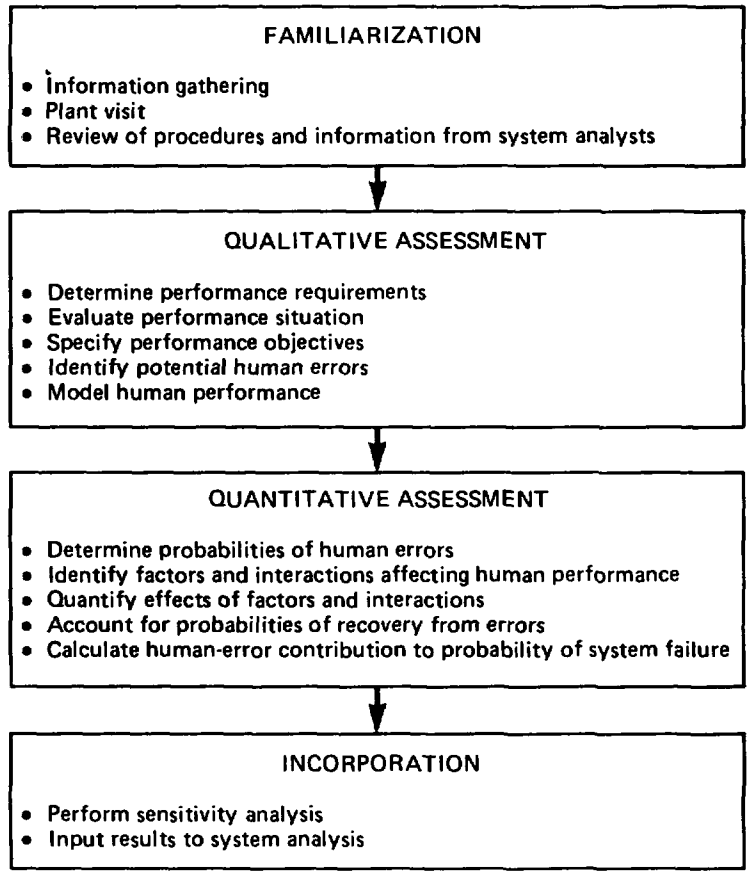

Figure 4-1. The phases of a human-reliability analysis.

information becomes available. Figure 4-2 is a block diagram for a complete analysis; for less detailed studies, such as a bounding analysis, some of the steps can be modified to reflect the appropriate level of detail and some of the steps can be eliminated. Obviously, the less plant-specific information the analyst has, the more uncertain his estimates. In a sense, the degree of uncertainty drives the level of analysis that is possible. The more uncertain an analyst's estimates, the closer his analysis is to being qualitative. A bounding analysis is more appropriate than a strictly quantitative assessment of the likelihood of any set of human errors when the information leading to the estimation of such errors is suspect.

4.2.1 PLANT VISIT

A survey of the control room during a general plant visit is an essential preliminary to the performance of a plant-specific HRA. This is to allow the analyst to become familiar with the operation of the plant. The purpose of the visit is not necessarily to evaluate the design of the control room, but rather to identify the aspects of the control room, the general plant layout, and the plant's administrative control system that affect generic human performance. No evaluation of any individual's performance is to be done. This point must be clearly understood by plant personnel if accurate and complete information is to be obtained.

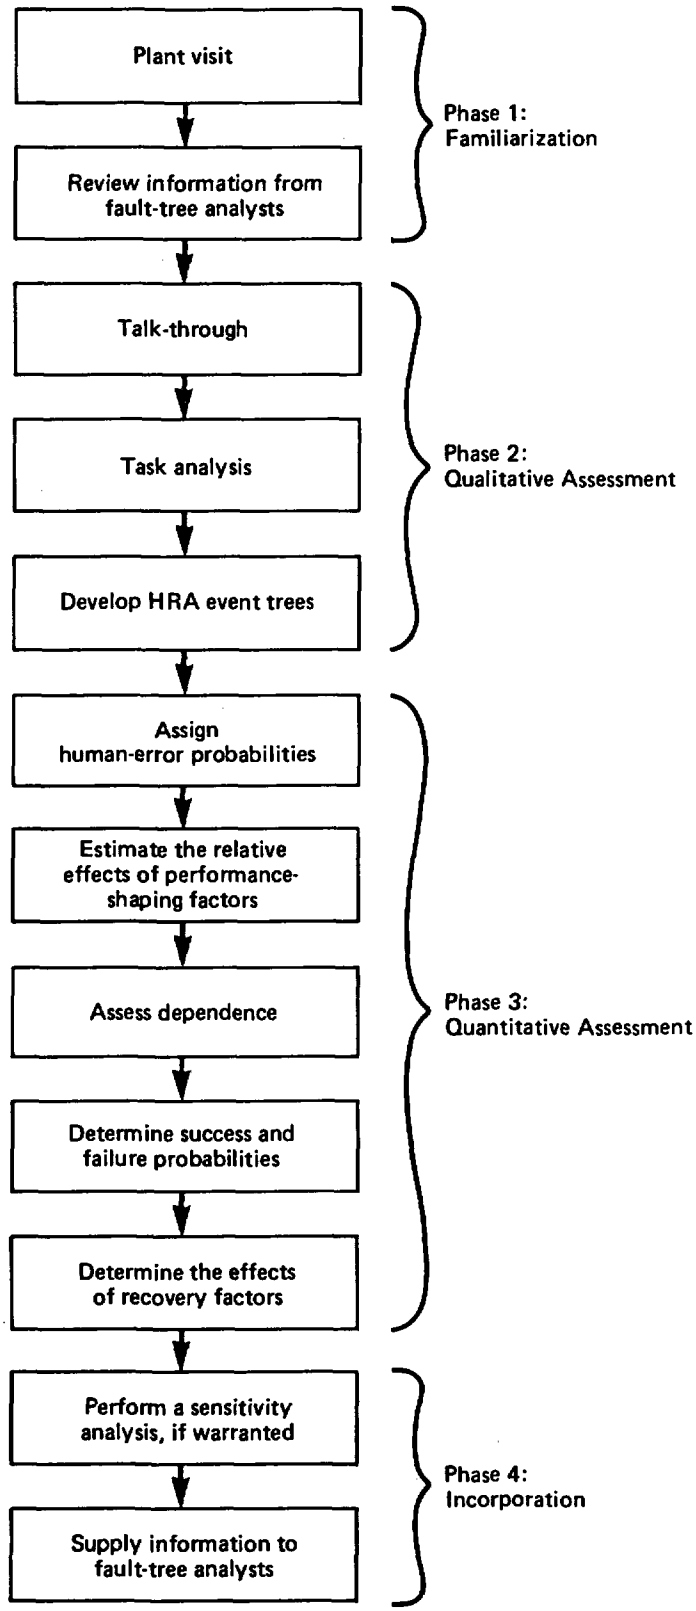

Figure 4-2. Overview of a human-reliability analysis.

4.2.2 REVIEW OF INFORMATION FROM SYSTEM ANALYSTS

For a given scenario or sequence of events, the system analysts identify human actions that directly affect the system-critical components. In light of the information obtained from the plant visit, the human-reliability analyst must review these actions in the context of their actual performance; the objective is to determine whether these actions can be affected by factors that may have been overlooked by the system analysts. For example, if performance on a noncritical element subsequently affects performance on a system-critical element, this effect must be considered, even though that task in itself is not important to the reliability of the system as defined by the system analysts.

4.2.3 TALK-THROUGH OF PROCEDURES

Sometimes performed in conjunction with the survey of the control room and sometimes at a later date during interviews with operations personnel, talk-throughs of the procedures in question are an important part of any human-reliability analysis. They are conducted by the human-reliability analyst and performed by plant operations personnel. The analyst questions the operator on points of the procedure until his understanding of the task is such that he could perform it himself or at least be able to understand fully the performance of the task. Performance specifics are identified along with any time requirements, personnel assignments, skill-of-the-craft requirements, alerting cues, and recovery factors. (The talk-through can also be performed for activities not defined by a specific plant procedure, but the effort required of the human-reliability analyst for such an analysis is greatly increased.)

The information obtained in a talk-through should enable the analyst to account for the effects of a situation's performance-shaping factors. (See Chapter 3 of the Handbook for a discussion of these factors.) Modifications made to the nominal HEP values from the Handbook will be based on information gathered here.

4.2.4 TASK ANALYSIS

At this point, a task analysis should be performed, as described in Chapter 4 of the Handbook. A "task" is defined as a quantity of activity or performance that the operator sees as a unit either because of its performance characteristics or because that activity unit is required as a whole to achieve some part of the system goal. Only the tasks that are relevant to the safety of the system are considered. A task analysis involves breaking down each task into individual units of behavior. Usually, this breakdown is done by tabulating information about each specific human action. The format of such a table is not rigid: any style that allows the information to be retrieved easily can be used. The format will reflect the level of detail as well as the type of task analysis to be performed. The analysis itself and the information it yields can be either qualitative or quantitative. Examples of task-analysis formats are presented later in this chapter.

Specific potential errors should now be identified for each unit of behavior. For every human action appearing in the task-analysis table, likely errors of omission and commission should be identified. A human action (or its absence) constitutes an error only if it has at least the potential for reducing the probability of some desired event or condition. The existence of this potential should be identified in conjunction with the system analysts. For every human action appearing in the task-analysis table, likely errors of omission and commission should be pinpointed. As mentioned earlier, extraneous acts are seldom considered. For example, the analyst may determine that, because of the control-panel layout, a selection error is possible during the manipulation of a specific switch, but his analysis will not usually predict which other element will be chosen, nor will it deal with the consequences of selecting a specific incorrect switch.

The analyst must also evaluate errors that may affect the probabilities of system success and failure but do not appear in the task analysis. Some of these can be disregarded by assuming for the entire analysis that a certain condition does or does not exist. For example, in the case of a post-maintenance test, if we are interested in the conduct of the test itself, we may arbitrarily assume that the supervisor has ordered the test. In determining which of these assumptions may be made, great care must be taken, however. In analyzing actual plant conditions, it is inappropriate to assume that something that should be done will always be done.

4.2.5 DEVELOPMENT OF HRA EVENT TREES

Each of the errors defined above should be entered as a binary branch on an HRA event tree, as described in Chapter 5 of the Handbook. The possible error events should appear on the tree in the order in which they might occur if such order is relevant. The suggested format for HRA event trees will be presented later. The product of the HRA event tree is a probabilistic statement as to the likelihood of a given sequence of events. Some PRAs deal only with the probability of successful completion of all human actions, while others take a more global approach, considering all system interactions and reactions that may contribute to the probability of system success. In either case, recovery factors usually are not included at this time. This is simply a time-saving feature of this HRA procedure. If, in a preliminary system analysis, the probability of an unrecovered human error is found not to impact system safety significantly, there is no need to expend additional time and effort on identifying and quantifying the effects of recovery factors acting on the situation.

4.2.6 ASSIGNMENT OF NOMINAL HUMAN-ERROR PROBABILITIES

An estimate of the probability of each human-error event on the HRA event tree must be derived from the data tables in the Handbook or from other sources. Tables of human-error probabilities (and the associated uncertainty bounds) for generic task descriptions are found in Chapter 20 of the Handbook.

One of the reasons the analyst should become familiar with the Handbook is the need for a thorough understanding of the assumptions and limitations of these tables. If there is no exact match between the description of a task in the Handbook and that defined by the task analysis, the estimated error probability for a similar task can be used as is, or it can be extrapolated, depending on the degree of similarity between the descriptions. "Similarity" in this context refers to the likeness of required operator behaviors. There can be a high degree of similarity between the performance of two tasks even though the equipment is dissimilar. The experience of a human-factors specialist is very valuable for this kind of judgment.

4.2.7 ESTIMATING THE RELATIVE EFFECTS OF PERFORMANCE-SHAPING FACTORS

The human-error probabilities estimated in the Handbook for a given task must now be modified to reflect the actual performance situation. For example, if the labeling scheme at a particular plant is very poor, in comparison with those described in Military Standard 1472C (U.S. Department of Defense, 1981) or NUREG-0700 (USNRC, 1981b), the probability should be increased toward the upper bound of its uncertainty bounds. If the tagging control system at a plant is particularly good, perhaps the probability for certain errors should be decreased.

Some of the performance-shaping factors (PSFs) affect a whole task or the whole procedure, whereas others affect certain types of errors, regardless of the tasks in which they occur. Still other PSFs have an overriding influence on the probability of all types of error in all conditions. Familiarity with the Handbook's treatment of PSF effects is necessary for the performance of these procedures.

4.2.8 ASSESSMENT OF DEPENDENCE

In any given situation, there may be different levels of dependence between an operator's performance on one task and on another because of the characteristics of the tasks themselves or because of the manner in which the operator was cued to perform the tasks. Dependence levels between the performances of two (or more) operators may differ, also. The analyst should keep in mind that the effect of dependence on human-error probabilities is always highly situation-specific. The concepts presented in the Handbook (the chapter on dependence) should be followed precisely.

4.2.9 ESTIMATING SUCCESS AND FAILURE PROBABILITIES

The criteria for system success and failure will be supplied by the system analysts. These criteria are used as the basis for labeling the end point of each path through an HRA event tree as a success or a failure. Multiplying the probabilities assigned to each limb in a success or failure path through the HRA event tree provides a set of success and failure

probabilities that can then be combined to estimate the total system success and failure probabilities.

4.2.10 DETERMINING THE EFFECTS OF RECOVERY FACTORS

It is often convenient to postpone consideration of the effects of recovery factors until after the total system success and failure probabilities have been determined. The estimated probabilities for a given task sequence may be sufficiently low without considering the effects of recovery factors so that the sequence does not appear as a potentially dominant failure mode. In this case, it can be dropped from further consideration.

4.2.11 PERFORMING A SENSITIVITY ANALYSIS, IF WARRANTED

To determine the effect of a single parameter on the total system-success probability, a sensitivity analysis can be performed. In this exercise, the value of a given parameter is manipulated and the resulting system-success probabilities are compared to judge the impacts of different magnitudes of change. This is not a necessary part of a human-reliability analysis in all cases, but it is extremely helpful in identifying the elements of the system that have relatively large or small effects on system safety.

4.2.12 SUPPLYING INFORMATION TO SYSTEM ANALYSTS

A copy of each HRA event tree along with a synopsis of the results, a copy of the task-analysis table, and a list of the underlying assumptions should be presented to the system analyst. The system analyst, the human-reliability analyst, and someone familiar with the actual performance of the operation should then go over the HRA event tree and the associated assumptions very carefully. This ensures that the human-reliability analyst has correctly defined the success of the system and that the system analyst does not apply the results of the HRA event tree outside the scope of its stated limitations.

4.3 METHOD

The theory, models, and data presented in this chapter are taken from the Handbook. Original sources for some of the methods (e.g., task analysis) can be found there.

The basic components of a human-reliability analysis are the task analysis and the Technique for Human Error Rate Prediction (THERP). Task analysis involves breaking down system-required human actions (or tasks) into

small units of physical or mental performance (steps) as well as identifying to the extent possible likely human actions not required by the system but having the potential for degrading certain system functions. These small units are then fully described and analyzed in terms of the PSFs that affect each function and combinations of them. The performance models and theories that make up THERP are then applied to these steps. Possible human errors are identified, and estimates of the probability of each error are derived. The end product of a human-reliability analysis is a set of system success and failure probabilities that reflect the probable effects of human errors. These system-based probabilities are in a form suitable for entering into the system fault trees by task or component.

Alternatives to THERP are discussed in Section 4.9 as well as in reports by Meister (1971), Embrey (1976), and Pew et al. (1977).

For cases in which it is necessary to use expert judgment to derive estimates of the probabilities of human error in nuclear power plants, there are a number of psychological scaling methods available. For a recent review, see Stillwell et al. (1982). In addition, the NRC, the Institute of Nuclear Power Operations, and the British National Centre of Systems Reliability (Embrey, 1981) are developing methods for psychological scaling specifically addressing nuclear power plant tasks. At present, no one method can be recommended since these studies are still under way.

4.4 INFORMATION REQUIREMENTS

After the system-analysis team has determined which system-critical events or components are to be evaluated, the human-reliability analyst should double-check to ensure that no potential human contributions have been overlooked. Procedures for performing each of the tasks involved in these events must therefore be evaluated. These procedures can be written, oral, or in the form of known standard shop practice or skill of the craft. In the case of written procedures, a copy of the procedure itself should be supplied to the human-reliability analyst; in the other two cases, the specifics required of the performance must be determined in the course of interviews with, and observations of, plant personnel.

The human-reliability analyst must become familiar with the plant, especially with the layout of the control room, and with the plant's general operating standards and administrative controls. The analyst who is not familiar with these aspects of a particular plant should make at least one visit (and preferably several) to the plant specifically for surveying the control room. Blueprints, drawings, or photographs of the consoles and control boards should be available for later reference. Personnel familiar with all phases of plant operations should be on call to provide information about control-room specifics and other features peculiar to the plant.

Human-reliability analysts need not have a thorough understanding of plant systems and functions--they need not have the same understanding of

these systems and functions as other specialists on the risk-assessment team. (Ideally each member of the PRA team would have at least a working knowledge of PRA fields other than his own; however, such people are not usually available in numbers large enough to support a full-scale PRA.) They should concern themselves primarily with actual human performance--system causes and effects are of no interest except in that they may influence an operator's perception of the urgency of a particular task. The system analysts and plant representatives are chiefly responsible for defining the impacts of human errors on the systems and functions of the plant. Their close interaction with the human-reliability analyst will ensure that the modeling of the effects of human errors is correct. In quantifying these effects, the underlying assumptions and limitations that apply to the models and data presented in the Handbook must be understood and not contradicted in their applications to a PRA.

4.5 PROCEDURE

4.5.1 INTRODUCTION

The purpose of performing a human-reliability analysis as part of the PRA described in this document is to determine the contribution of human errors to predetermined significant system failures. The object of such an analysis is to treat the relevant human actions as components in system operation and to identify error probabilities that could significantly affect system status. This section outlines an approach to be used in deriving relevant human-error probabilities along the guidelines established in the Handbook.

As already stated, the human-reliability analysis should be performed by a human-factors specialist who is familiar with the theory and techniques presented in the Handbook. For a complete human-reliability analysis, he must have an understanding of the plant's administrative-control network, some familiarity with the layout and the operating characteristics of the control room, and frequent access to plant personnel who can provide information on specific aspects of performance situations. Without sufficient plant-specific information, he will be unable to perform a human-reliability analysis that models the actual plant situation adequately in that he will not have defined all the potential human errors--nor will he have accounted for all the likely recovery factors.

This section discusses each of the major HRA tasks outlined in Section 4.2. An example of a human-reliability analysis is presented in tandem with these discussions. The description of each task is supported by an example of application to an actual human-reliability analysis.

There are several possible sequences for the elements of a human-reliability analysis. The sequence presented here is by no means absolute, but it is a sequence that served well for the Interim Reliability Evaluation Program and other PRAs. The elements themselves were derived from THERP and should be included in all complete human-reliability analyses. The

recording and reporting formats described here can be modified for the convenience of the analyst, but he should keep in mind the type and level of detail of information necessary for someone else to understand his analysis. The analysis can be used for qualitative as well as quantitative assessments, with the level of detail of the information collected reflecting that of the analysis itself. Of necessity, human-reliability analysis must depend largely on data that are extrapolations from tasks not directly related to nuclear power plants and on models that have not been verified in the strictest sense of the word. Nevertheless, this application of the theory, data, and models presented in the Handbook represents an attempt at standardizing the approach to performing human-reliability analyses for the probabilistic risk assessments of nuclear power plants.

4.5.2 PLANT VISIT

4.5.2.1 Discussion

At least one plant visit, specifically including a detailed survey of the control room, should be made at the onset of the analysis. Before this visit, the analyst should make arrangements with the plant as to the plant areas to be visited, the requirements for access, and the types of personnel to be made available for interviews. Every attempt should be made to minimize impact on the plant and on the utility as well as the disruption of plant operations.

When possible, the human-reliability analyst should meet with representatives of the plant and/or utility before visiting the plant. The objective of this meeting is to advise the plant and utility representatives about the purpose of the evaluation. More cooperation at all levels of involvement will be afforded if the concerned parties understand that the role of the human-reliability analysts is not condemnatory or judgmental. The main purpose of the visit should be stressed: the observation of plant conditions in order to provide accurate descriptions of actual performance for the analysis. The observations are to be expressed only in descriptive terms. No "solutions" to plant problems or inadequacies are to be offered.

In the initial visit to the plant, the human-reliability analyst will make notes on relevant performance-shaping factors, especially those pertinent to control-room operations. If the system analysts have already identified the plant subsystems or procedures that are of interest, these can be examined closely at this time. This visit should provide general information about the plant's operating characteristics and a "feel" for the effectiveness of the plant's administrative controls.

In surveying the control room, specifics relating to the layout of controls and displays should be noted. Copious notes should be taken on the characteristics of critical controls and displays, noting any factors that would influence their use--anything that would aid or hinder the operators in either locating, manipulating, or interpreting them. Deviations from good human-factors engineering practices, such as those noted in the previously cited military standard (U.S. Department of Defense, 1981) and NRC guidelines

(USNRC, 1981b), should be noted. Any specifics related to the operation of critical subsystems that have been pinpointed for observation by the system analysts should be recorded. If the system analysts have identified the plant procedures of interest, the time at the plant should also be used for a talk-through of these procedures (Section 4.5.4).

4.5.2.2 Example

Listed below is a set of notes similar to those that would be collected during an actual plant visit.

1. On some chart recorders the indications are hazy because of the use of nonglare glass. The operations superintendent says they are all being changed to regular glass. (The nonglare glass had been recommended by the manufacturer.)

2. Some labels for two-channel switches are sideways because of space restrictions. (Later note: When these sideways labels appear between displays, some confusion in relating a label to a display may result.)

3. Each annunciator panel is numbered, with the numbers increasing from right to left (so do the numbers for the control board and panels).

4. On the fronts of control panels CB1 and CB2, there are rows of J-handle switches, the first of which are turned inward to prevent inadvertent manipulation. This is not true for panel CB4, but its J-handle switches are not critical to plant operation. Those on panels CB1 and CB2 are for oil pumps and turbines, and their movement would cause a trip. The direction of manipulation for the reversed J-handles is the same as for the outward-facing ones.

5. Some J-handles have arrows at their bases that indicate the direction of operation; some do not. (Note: Different manufacturers?) Handles, other than the J-handles, have arrows at their bases, especially knurled or symmetrical handles. The size of these shape-coded handles is such that the arrows cannot be seen easily, especially when viewed at eye level straight on.

6. At the alarm cathode-ray tube (CRT), there are three display modes: a flashing dark-green display indicates a new, unacknowledged alarm; a steady dark-green display indicates an uncorrected but acknowledged alarm; and a steady light-green display indicates a cleared alarm (it remains on for reference only).

7. For the engineered-safety-feature (ESF) panels in the cabinets in the back (as well as other indications in the control room), display status and some parameter readings must be recorded at various intervals. (Note: Need to request a copy of "Procedures for Conducting Plant Operations" to review the checklist used versus the frequency of its use and the location of all controls checked.)

8. On the ESF panels in the control room, the color of the label for a particular item is the color of the indicator light during actuation

of the automatic safety equipment. During system response to an emergency, the operator can scan the ESF panel quickly to see whether the lights that are on are the same color as the labels for those items. A disagreement between the colors indicates that some safety system has malfunctioned or has been overridden manually for some reason.

9. Stubs from yellow tags for valve-change operations are tossed into a drawer; no record of them is in evidence. (Note: Check this out.)

10. The labels on locally operated valves are impression-printed on metal tags and, because of poor lighting, are difficult to read. No indication that designates their normal positions is present at these valves.

Obviously, there are other observations that could be made during a survey, but they have been omitted here for the sake of brevity. The levels of detail for the control-room survey and the inspection tour of the plant are at the discretion of the human-reliability analyst and should reflect the level of detail required by the risk assessment being performed. Specific information about the conduct of certain procedures identified later in the program can be supplied by plant personnel during a talk-through, with the human-reliability analyst interpreting that information in the light of knowledge gained during the plant visit.

4.5.3 REVIEW OF INFORMATION FROM SYSTEM ANALYSTS

4.5.3.1 Discussion

After the screening process the system analysts will present the human-reliability analysts with a set of scenarios to be analyzed. These will usually take the form of operator performance on a critical system element during the course of following a set of plant procedures. The system analysts will have identified system-critical components and the circumstances under which they will be manipulated. The human-reliability analysts must then determine the probability of human errors in dealing with these components. They must also determine whether human performance on other elements or in the conduct of the plant's administrative controls will affect the probability of error in operating the system-critical components.

Often, the system analysts will present the human-reliability analysts with a set of plant procedures from which they have pinpointed the steps that they feel deal directly with the operation of system-critical components. In other cases, they may have identified entire systems for which human errors must be identified and quantified. In either case, the human-reliability analyst must examine all of the plant procedures associated with these elements to determine whether they require performances on other elements that might affect the probability of error on the critical components or systems. At times, these determinations will have to be made in conjunction with the talk-throughs of the procedures (Section 4.5.4).

During this review of the information received, the critical task of the human-reliability analyst is to ensure that all human actions are

analyzed in the context of actual performance. Human actions in a nuclear power plant should not be treated as isolated entities, unaffected by other factors. There are many interactions in a nuclear power plant--between personnel and between tasks--that must be identified. Some of them will affect the assessment of levels of dependence between certain behaviors (Section 4.5.9); some of them will have a global effect on the performance of all tasks in a given procedure. The system analysts will have identified the interfaces between critical equipment items and associated human tasks. However, the interactions between these and other system elements should be identified by the human-reliability analyst, who has been trained to spot them. This extra investigative effort on the part of the human-reliability analyst must ensure that they are all identified.

In some cases, a single plant procedure will cover several sets of tasks involving critical components. For example, in restoring items of equipment after maintenance, the operators may follow a general plant procedure governing the application and removal of tags. This administrative control may apply to all tasks in which tags are used. In this case, it is the conduct of the administrative-control procedure that is analyzed, as well as the restoration act itself. The operator is actually following the control procedure rather than a set restoration procedure for a specific component. Here the human-reliability analyst can examine one procedure (the administrative-control procedure) and apply the results to all tasks involving restoration after maintenance. He must take care, however, to determine that the administrative-control procedure applies to every task he analyzes.

As he reviews the information received from the system analysts, the human-reliability analyst should search for deviations from, or inconsistencies with, the assumptions underlying the theories and models in the Handbook. The human-error probability estimates in Chapter 20 of the Handbook are based on limitations on their use--limitations that must not be contradicted. The human-reliability analyst must examine a given procedure in the context of its performance to assess its conformance to these limitations.

4.5.3.2 Example

A set of hypothetical plant procedures dealing with response to a small loss-of-coolant accident (LOCA) is presented in Figure 4-3. Only part of the procedure is given, and the steps identified by the system analysts as being critical are indicated with a double asterisk. The system analysts have assumed that the situation has been diagnosed correctly and that the operators have correctly completed the immediate actions required by the situation. These assumptions limit the nature of the human-reliability analysis because, given them, the human-reliability analyst does not have to account for errors in diagnosis or for the fact that the level of stress experienced by the operators might be higher because of their having made mistakes in the immediate actions. However, those systems that have been judged to have the potential of being degraded by human errors are those involved in the "Subsequent Activities" section of the procedures. These, therefore, are the only ones to be considered in this example. (The treatment of diagnosis errors will be discussed in a later section.)

D. SUBSEQUENT ACTIVITIES

Note: Reverify asterisked parameters in all sections, using alternative indications if available. Select proper computer functions to monitor incore thermocouples.

*If FW and RCPs are available (manual HPI actuation, no automatic actuation), proceed through Section D.
*If no FW is available, proceed to Section E.
*If FW is available but RCPs are not, proceed to Section F.

D.1 Stop all but one RCP in each loop.

Note: If ES actuation occurs before HPI can be manually established and the RCS pressure recovers, do not reset ES analog channels, since this would delay restart of actuated equipment in the event of a loss of offsite power as pressure would have to fall again to the actuation setpoint.

**D.2 Monitor RCS pressures and temperatures; maintain at least 50°F margin to saturation by holding RCS pressure near the maximum allowable pressure within the cooldown pressure-temperature curve (Figure B).

Note: If RCS pressure is not restored before the pressurizer goes solid, or if the RCS relief valve alarm remains in, the leak may be in the pressurizer steam space, and the pressurizer must be taken solid to regain RCS pressure. If such is the case, reopen ERV block valve MOV-1300 to allow ERV operation before pressurizer code safeties.

**Caution: HPI components are not to be overridden unless the following criteria are met:

1. The HPI system has been in operation for 20 min, and all hot- and cold-leg temperatures are at least 50°F below saturation temperature for the existing RCS pressure, or

2. The RCS is >50°F subcooled, and throttling of HPI is necessary or

3. The RCS is 50°F subcooled, and HPI throttling is necessary to remain within the plant cooldown pressure-temperature curve limits, or

4. DH or LPI has been operating for >20 min with total flow rates of ≥2000 gpm.

If margin to saturation drops below 50°F after HPI override, reinitiate maximum HPI until >50°F subcooled. UNDER NO CIRCUMSTANCES IS HPI TO BE OVERRIDDEN IF RCS IS NOT SUBCOOLED.

D.3 Monitor RB pressure; if pressure reaches 4 psig, verify reactor building isolation and cooling actuation (ES channels 5 & 6) and HPI & LPI actuation (channels 1, 2, 3, and 4).

Note: Proper ES actuation is verified by noting that the colors of components' indicating lamps on the ES panels ES-16 and ES-18 and CB-26 correspond to the colors of the switch nameplates. Proper flow ranges for HPI, LPI, and RB spray are marked on the meter faces. Proper penetration room ventilation is verified by noting all room isolation damper lights out, flow indicated, and negative penetration room pressure indicated....

**D.4 If RCPs and FW are available, and RCS margin to saturation is >50°F, override and throttle HPI MOVs to control system pressure if pressurizer is solid or to hold pressurizer level at setpoint while using pressurizer heaters and spray for RCS pressure control; initiate plant cooldown per Plant Procedure 12 at a rate that allows RCS pressure to be maintained within the cooldown pressure-temperature envelope.

D.5 If RCS pressure falls to within 50°F of saturation or if low margin to saturation temperature alarms are received, maintain maximum HPI flow until 50°F margin is restored.

D.6 If RCS pressure falls below secondary pressure, reduce and maintain secondary pressure at 20 lb/in. less than primary pressure and maintain maximum HPI flow until subcooled, then initiate a cooldown by decreasing secondary pressure per Plant Procedure 23.

Figure 4-3. Excerpt from the procedures for responding to a small LOCA. The critical steps are indicated by a double asterisk.

**D.7 Prepare for LPI boost to MU pump suction and RB sump recirc as follows:

D.7.1 Verify MU tank outlet MU-13 closed.

D.7.2 Open DH-7A and DH-7B, LPI discharge to MU pumps suction, verify MU pump suction crossover valves MU-14, MU-15, MU-16, and MU-17 open, and verify MU pump discharge crossover valves MU-23, MU-24, MU-25, and MU-26 open.

D.7.3 Isolate the DH rooms by closing both DH room floor drain valves, ABS-13 and ABS-14, securing room purge dampers CV-7621, CV-7622, CV-7637, and CV-7638 from ventilation control panel (east wall of 404-foot ventilation room) and closing watertight doors.

D.7.4 Verify both DH pumps operating and both LPI MOVs open (MOV-1400 and MOV-1401).

D.8 Once a 50°F margin to saturation is attained.....

**D.9 Monitor BWST level; when BWST level has fallen to 6-foot indicated level or when the corresponding BWST lo-lo-level alarm is received, transfer suction to RB sump by verifying RB sump suction valves inside containment MOV-1414 and MOV-1415 open, opening RB sump suction valves outside containment MOV-1405 and MOV-1406 (a slight upward perturbation should be noted on pump flows indicating suction transfer) then close both BWST outlets MOV-1407 and MOV-1408 (refer to Plant Procedure 23 for RCS temperature control methods). Close NaOH tank outlets MOV-1616 and MOV-1617. MANUAL OVERRIDE PUSHBUTTONS MUST BE DEPRESSED FOR ALL VALVE MANIPULATIONS IF ES ACTUATION HAS OCCURRED.

Figure 4-3 (continued). Excerpt from the procedures for responding to a small LOCA. The critical steps are indicated by a double asterisk.

Given the above assumptions and following a detailed reading of the procedures, everything seems to be in order for a straightforward use of the theories and models in the Handbook, with one exception: the performance of these tasks occurs about an hour after the onset of the small LOCA. The Handbook chapter on stress states that there will be three operators in the control room at this time. However, some of the actions required by this procedure take place outside the control room. Because of the response time involved in donning the protective clothing required for these tasks, it is assumed here that only two qualified operators will be in the control room. Of course, during an incident of this type several people will probably be present in the control room. However, the shift supervisor is still in charge of operations, and personnel working for him are likely to follow his instructions and line of thought. Therefore, it is conservatively assumed that the presence of several people would be no more beneficial than the presence of only three licensed operators.

4.5.4 TALK-THROUGH

4.5.4.1 Discussion

In a talk-through of a set of procedures for which safety-critical events have been identified, the human-reliability analyst questions someone familiar with the performance of that procedure on specific points of the procedure until the analyst is so familiar with the tasks that he could perform them himself or at least understand fully the performance of an operator. The talk-through can be performed on sets of written or oral plant procedures,

standard shop practice, or training methods. It could take place at a simulator instead of at the plant itself, but the human-reliability analyst must take great care in noting which of the characteristics of the simulator are unlike those of the plant.

During the talk-through, the analyst must determine the performance-shaping factors that influence behavior, such as the location and the physical and operating characteristics of specific controls, the type and location of alarms and annunciated indicators, control-room manning and task allocation, time requirements, and limits for alarm indications and responses. He must also "translate" the written procedures into English as he speaks it; that is, he must determine the meaning of the specific instruction resulting from each command given in the set of procedures in the language of that particular plant. The analyst must specify in language he can understand the exact interpretation the operators will make from the sometimes vague wording of plant procedures. At times, these interpretations are based on the operator's knowledge of system operation rather than on a standardized plant definition of the term in question. When this is the case, the analyst must ascertain whether all the operators define that term in the same way.

In performing a talk-through, the human-reliability analyst conducts an interview with a plant employee who is familiar with the performance of the procedure in question. (In the case of a new plant, the person most familiar with the development of the procedures should be interviewed.) To obtain more familiarity with the performance characteristics of the procedure, the analyst should ask general questions about the performance-shaping factors acting at the time of performance and specific questions about the factors affecting the performance of the critical steps.

A talk-through can be performed as part of the control-room survey. In this case, the operator and the human-reliability analyst actually follow the path taken by the operators in performing the procedure. When the procedures call for the manipulation of a specific control or for the monitoring of a specific set of displays, the operator and the analyst approach them at the control panels, and the operator points out the controls and displays in question. The procedure is followed in sequence, and the analyst could generate a link analysis at this time. (Link analysis is discussed in Chapter 3 of the Handbook.)

Careful notes recording the outcome of the talk-through must be taken. Much of the information from these activities will be entered directly into the task-analysis tables (Section 4.5.5) for later use.

4.5.4.2 Example

In the talk-through of the procedures in Figure 4-3, some general information was gathered that relates to the performance of all the steps in the procedure. They are listed below.

1. The plant is following an emergency procedure. (Note for later reference: There will be some level of stress for the operators.)

2. The "Subsequent Activities" section of the procedures will be performed approximately 1 to 1.5 hours after the start of the accident.

3. At least three licensed operators will be available to deal with the situation. One of them will probably be the shift supervisor.

4. At this plant, "verify" means to check and, if necessary, to correct the status of a given item of equipment. For example, if the operator must verify that a valve is open and, on checking its status, finds it closed, he must open it manually.

5. The asterisked note at the beginning of the section indicates that the performance of the procedures in Section D is to be reverified (double-checked) after the procedures have been completed. This constitutes a recovery factor and, as such, will not be included in the HRA event tree at this time.

6. The "caution" in Figure 4-3 stems from actions taken during the incident at Three Mile Island Unit 2. Because of the special implications of performing them incorrectly, these actions will be considered separately.

7. Steps D.2, D.4, D.9, and D.7.4 are performed in the control room. They will be diagrammed separately from steps D.7.1, D.7.2, and D.7.3, which take place outside the control room.

Specifics relating to the performance of individual steps will now be given in the order of the steps.

Step D.2. The pressures of the reactor-coolant system (RCS) are found on a chart recorder; RCS temperatures can be read from digital indicators; both are on a front control board. A copy of the pressure-temperature curve is taped to the side of the computer terminal, adjacent to these other indicators. To manipulate the pressure and temperature values, the heater switches found on the same front control board will be used.

Step D.4. There are four switches for four motor-operated valves (MOVs) for high-pressure injection on the vertical ESF panels. A sketch of the layout of the controls is shown in Figure 4-4. Cooldown is initiated by following another procedure. The operator says that this other procedure is so well known that he cannot think of any situation in which it would actually be necessary to refer to it.

Step D.7.1. Valve MU-13 is a manual valve in the stairwell outside the main unit pump room. This stairwell is two levels down from the control room.

Step D.7.2. The layout of these valves is shown in Figure 4-5, with the channels they represent. One channel should always be completely open so that the operator should only have to open one low-pressure-injection (LPI) discharge valve, two makeup-pump-suction crossover valves, and two makeup-pump-discharge crossover valves. The operators view this entire series of tasks as one unit task: in their interpretation, all these steps are

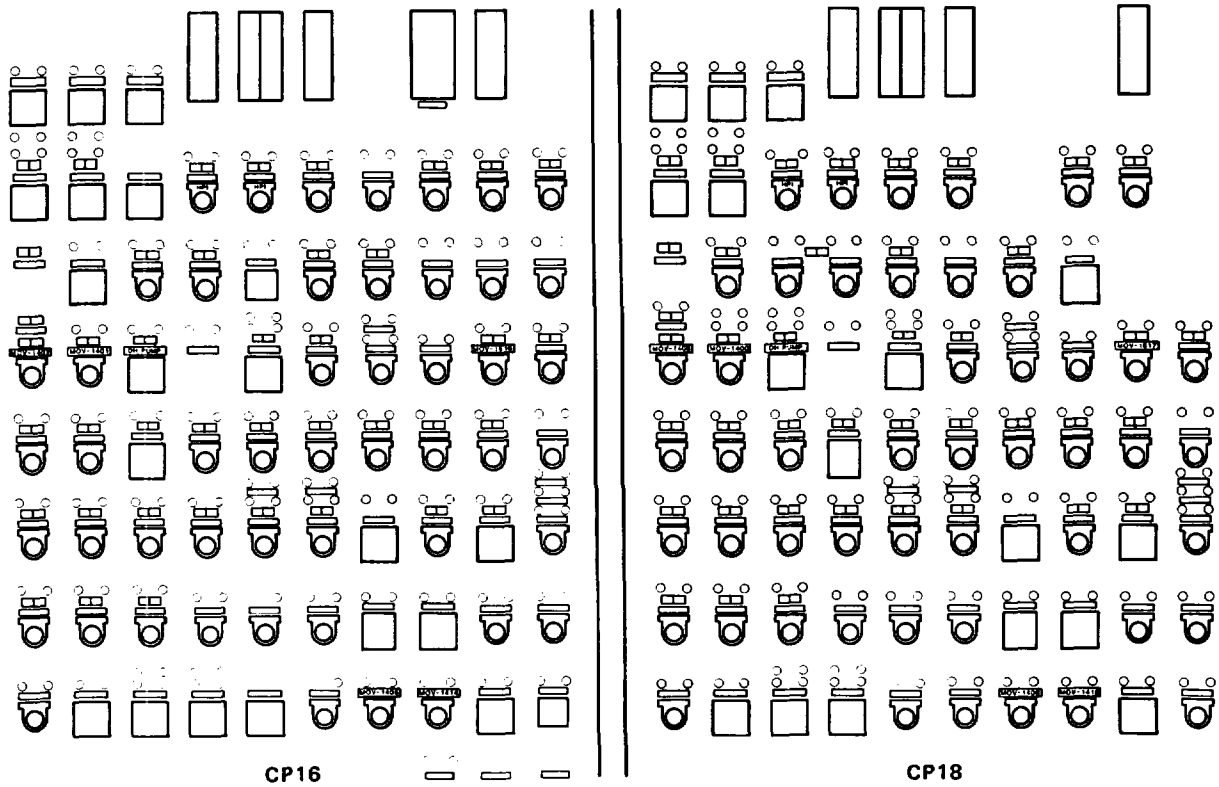

Figure 4-4. Layout of controls on the ESF panels.

performed to satisfy a major system function. These valves are located one level below the makeup-pump room.

Step D.7.3. Valves ABS-13 and ABS-14 are large, locally operated valves located outside the decay-heat (DH) rooms, one level below the decay-heat pump rooms. They are large valves situated under the grating outside the watertight doors. There are no other valves under the grating. The ventilation room is two levels above the control room. The switches for CV-7621, 22, 37, and 38 are on the wall there, in the midst of dozens of other similar switches. They are grouped near each other and near other switches that control equipment in the same physical area of the plant, but there are no location cues on the wall to indicate where this grouping can be found among other groups.

Step D.7.4. Indicators for the decay-heat pumps and for the LPI MOVs are on the vertical ESF panels in the control room. (See Figure 4-4 for the layout of the panels.)

Step D.9. The level indicator is on a panel adjacent to the vertical ESF panels in the control room. The low-low-level alarm sounds when the 6-foot level is reached. During a small LOCA, this should happen no sooner than 1.5 hours after the start of the event. All the MOVs are on the ESF panels.

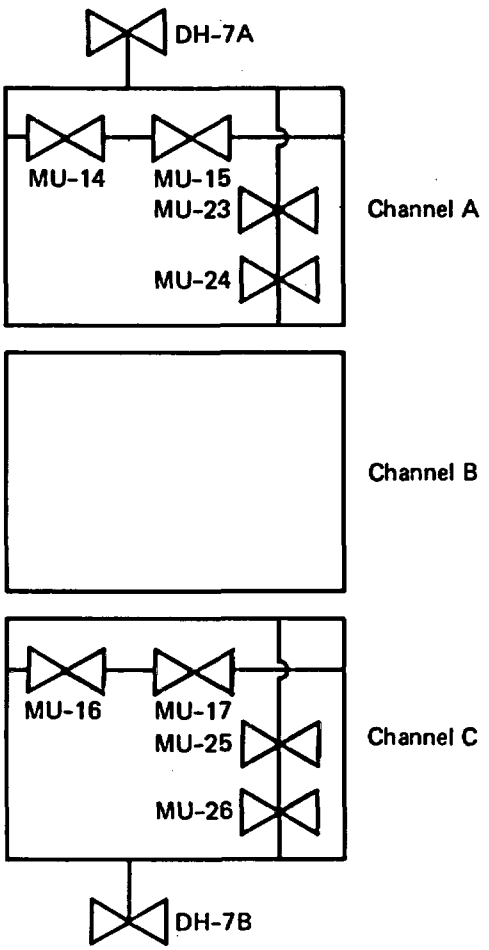

Figure 4-5. Layout of valves in DH pump rooms.

4.5.5 TASK ANALYSIS

4.5.5.1 Discussion

At this point, the procedure should be formally broken into tasks or smaller units of behavior; that is, for each step in the procedure, individual units of operator performance must be identified, along with other information germane to the performance. These individual units of performance constitute elements of behavior for which potential errors can be identified. In other words, a large task consisting of a set of steps should be broken down to allow the identification of errors associated with each step. All of this information must then be entered into a task-analysis table.

The format of this table is not specified, but the table must contain all the information necessary for later parts of the analysis. In most cases, the necessary information will consist of such items as the piece of equipment on which an action is performed, the action required of the operator, the limits

of his performance, the locations of the controls and displays, and explanatory notes. If different tasks are to be performed by different operators, the allocation of tasks to personnel can be indicated in the task-analysis table, or separate task-analysis tables can be made for each operator. The example in this section takes the latter approach.

The level of detail in a task analysis and the amount of information recorded should reflect the level of detail (qualitative or quantitative) of the risk assessment and are obviously determined judgmentally. The guiding rule for this determination is that one should be able at a later date (perhaps when the results of the human-reliability analysis are compared with those of another analysis) to recapitulate the rationale for the human-error probabilities that were used in the analysis.

Once the task steps have been broken down, potential errors must be identified for each step. The analyst must decide whether, for any given step, he should consider an error of omission or the various errors of commission (selection, reversal, sequence, etc.) that are likely for that step. This decision must be made based on the relevant performance-shaping factors and the task analysis. The steps should be listed chronologically.* Considering the characteristics of the actual performance situation, the human-reliability analyst must determine and record which types of errors the operator is likely to make and which he is not. For example, if an operator is directed by a set of written procedures to manipulate a valve and that valve is fairly well isolated on the panel, differs in shape from other valves on the same panel, and is very well labeled, the analyst may decide that errors of selection are not to be considered in this case. He should also have determined that, in following the written procedures, the operator might make an error of omission.

Extreme care should be exercised in deciding which errors, if any, are to be completely discounted. In comparison with tasks in other industries, most of the tasks performed in nuclear power plants have very low human-error probabilities, on the order of 10^{-3}. Although one error in a thousand opportunities seems quite low, a human-error probability of 10^{-3} may contribute substantially to the frequency of system failure. Rather than discounting a "questionable" error that he thinks may be unlikely, the human-reliability analyst should consider it and perform a sensitivity analysis to ascertain its impact on the probability of system success (Section 4.5.12). If the impact is found to be negligible, an appropriate indication can be made in the fault-tree block for the error.

---

*In some cases, it may be discovered that the order of the steps in the procedure is not necessarily the one followed by the operators. The task analysis and the resulting HRA event tree can easily reflect any performance sequence. However, the order of the steps in the procedure is usually assumed to be the most likely order of task performance. Recordkeeping is simplified by following the same task sequence from procedures to task analysis to HRA event trees.

Once he has identified the errors likely to be made in each unit of performance, the analyst must look for other factors that may affect performance. The entire performance scenario must be considered in this search. The analyst is looking for elements that are usually outside the scope of the procedures followed by the operator. For example, if something is to be done at the discretion of the shift supervisor, the supervisor's remembering to order the task will determine whether the task is performed by the operator. These extraneous factors that affect the probability of human error usually involve some sort of failure in the plant's administrative-control system. The quality of the plant's personnel-communication system and the potential for the disruption of communications during a particular performance sequence will also have to be examined in these cases.

Events other than human actions that affect subsequent performance must also be taken into account. If an operator's cue to initiate a task involves some signal from the equipment or an order from a supervisor, the probability of that signal's being generated or that order's being given must be considered. Many times, these equipment-failure probabilities are not provided by the system analysts or are not considered in the analysis on the basis of assumptions provided by the system analysts. The human-reliability analyst should not assume that the supervisor's order will always be given when it should be unless direct evidence supports such a conclusion.

The task analysis is usually designed and performed to agree with the level, dictated by the system analysts, at which the human-reliability analysis is incorporated into the system analysis. However, the level of incorporation--system event trees, a high (subsystem) level of the system fault trees, a low (component unavailability) level of the system fault trees, or any other level--affects only the format of the HRA results. It has no effect on the actual performance of the human-reliability analysis: all tasks are to be analyzed in the contexts of their performances. It is also of little consequence to the human-reliability analyst whether the information about task performances is considered in part or as a whole in another section of the PRA. The results of his analysis can be parceled for inclusion at the component level in the system fault trees or taken as a whole for inclusion at the subsystem level. The format used in the example can accommodate either.

4.5.5.2 Example

The task analysis for the procedures in Figure 4-3 has been done in two consecutive steps: (1) the tasks performed by the operators in the control room and (2) the tasks performed by an auxiliary operator outside the control room.

The table format used for this example is shown in Figures 4-6 and 4-7. The format used for the task analysis is relatively unimportant; it can be modified to reflect the type and the amount of information needed in later phases of the risk assessment. The step number from the written procedures is included for easy reference to the procedures should any questions arise. The actual items of equipment to be manipulated, read, or otherwise dealt with are listed in the "equipment" column. The "action" column contains

the commands given to the operator; they are usually the action verbs contained in the procedure. In the "indication" column, the analyst notes the cues (usually from visual displays) that inform the operator whether the action has been performed correctly and any restrictions on the operator's actions. In the sample task analyses of Figures 4-6 and 4-7, many of the indications are so obvious (e.g., turn switch to ON position) that no entry has been made. The physical positions of the equipment items are given in the "location" column. The "notes" column contains any information the human-reliability analyst believes will be useful in later parts of the analysis. In Figures 4-6 and 4-7, these columns indicate whether the equipment items of interest in the control room are on the ESF panel and whether locally operated valves are isolated or part of a group. The "errors" column lists the errors deemed likely for each task. They are discussed in detail for each step, beginning with those in Figure 4-6.

In Figure 4-6, dashed lines are drawn between sets of actions that apply to specific plant functions. They help the system analysts to keep track of which portion of the HRA event tree should be excerpted for insertion at the subsystem level of the system fault trees. In this case, step D.2 involves the operator's diagnosis of plant status. This step should be excerpted for inclusion with all others since its correct performance affects the probability of correct performance on the rest of the steps. Once this diagnosis has been made correctly, the operator will move to effect cooldown after verifying that saturation is adequate per step D.4. Step D.7.3 involves isolating the decay-heat pump rooms. Step D.7.4 calls for the operator's verifying the initiation of the decay-heat-removal function. Then, from the water level indicated for the borated-water storage tank (BWST), he must diagnose the need for switching to recirculation. This involves the first part of step D.9 (monitoring the BWST level) and must be excerpted along with any of the other errors from step D.9 (effecting recirculation) for inclusion in the system analysis of the recirculation system.

Step D.2. Monitoring and maintaining RCS pressure and temperature within the curve is considered to be a unit task of three steps: (1) reading the pressure chart, (2) reading the temperature from the digital indicator, and (3) manipulating the heater switches to keep the above values within the acceptable range on the pressure-temperature curve. As such, the probability of an error of omission applies to the entire task: only by forgetting to perform the task itself will the operator forget to perform any element of it. The possible commission errors are those made in reading the pressure from the chart recorder, the temperature from the digital readout, and the curve, which is in the form of a graph. The feedback from manipulating the heater switches incorrectly is almost immediate, and therefore the probability of making a reversal error in their operation is not considered. The pressure chart, the digital indicator, and the heater switches are located on one of the front control boards; a graph of the pressure-temperature curve hangs off the CRT console immediately adjacent. This unit task is performed several times per shift under normal and emergency operating conditions. The heater switches are functionally grouped and well labeled. Under these circumstances, errors of selection were not considered. These steps are considered dynamic in that they involve the continuous monitoring of the displays and the operation of the heater switches.

Step	Equipment	Action	Indication	Location	Notes	Errors	HRA event tree
D.2	RCS pressure	Monitor		CB4		1. Omission (all)	1
						2. Reading	2
	RCS temperature	Monitor		CB4		Reading	3
	heater switches	Maintain pressure and temperature	Within curve on chart	CB4		Reading	4
D.4	4 HPI MOVs	Override and throttle		CP16, CP18	ESF	1. Omission (all)	5
						2. Selection (1)	6
		Initiate cooldown	Procedure 12			Omission	7
D.7.3	CV-7621,22,37,38 (room-purge dampers)	Secure	Close switches	Ventilation room		1. Omission (all)	8
						2. Selection (each)	9,10,11,12
D.7.4	Decay-heat pumps	Verify on	Indicator lamps	CP16, CP18	ESF	1. Omission (for MOVs too)	13
						2. Selection	14
						3. Interpretation	15
	MOV-1400, 1401	Verify open	Indicator lamps	CP16, CP18	ESF	1. Selection	16
						2. Interpretation	17
D.9	Borated-water storage tank	Monitor level	>6 feet	CP14		1. Omission	18
						2. Reading	19
	MOV-1414, 1415	Verify open	Indicator lamps	CP16, CP18	ESF	1. Selection	20
						2. Interpretation	21
	MOV-1405, 1406	Open	MOV switches	CP16, CP18	ESF	1. Selection	22
						2. Reversal	23
	MOV-1407, 1408	Close	MOV switches	CP16, CP18	ESF	1. Selection	24
						2. Reversal	25
	MOV-1616, 1617	Close	MOV switches	CP16, CP18	ESF	1. Selection	26
						2. Reversal	27

Figure 4-6. Task-analysis table for actions by operators assigned to the control room. The column labeled "HRA event tree" does not usually appear in a task analysis; it has been included for the reader's convenience. The numbers in this column refer to the error event numbers in the HRA event trees starting with Figure 4-9.

Step	Equipment	Action	Indication	Location	Notes	Errors	HRA event tree
D.7.1	MU-13	Verify closed	Position	Stairwell outside makeup-pump room	Only valve	Omission	2
D.7.2	DH-7A, 7B	Open	Position	Outside decay-heat pump rooms		Omission (for all D.7.2)	3
	MU-14, 15, 16, and 17	Verify open	Position	Decay-heat pump rooms			
	MU-23, 24, 25, and 26	Verify open	Position	Decay-heat pump rooms			
D.7.3	ABS-13, 14	Close	Position	Outside decay-heat pump rooms	Only valve	Omission (for all D.7.3 here)	4
	Watertight doors	Close	Locks in place	Decay-heat pump rooms			

Figure 4-7. Task-analysis table for actions by auxiliary operator outside the control room. The column labeled "HRA event tree"does not usually appear in a task analysis; it has been included for the reader's convenience. The numbers in this column refer to the error event numbers in the HRA event trees starting with Figure 4-10.

Step D.4. Because their manipulations are called out in the same procedural step and because of their close proximity (see Figure 4-4), the operator views the throttling of the four HPI MOVs as a unit action. Therefore, the probability of an error of omission applies to them all. Because on the actual panel they are delineated with colored tape, a selection error for the group is very unlikely. However, as Figure 4-4 shows, a similar switch is next to the last HPI MOV control in the group. A selection error for that control is likely: instead of MOVs 1, 2, 3, and 4, the operator may throttle MOVs 2, 3, and 4 and the other control. The operators have stated that, in initiating cooldown, they probably would not refer to the other set of procedures. For this reason, an error of omission is assigned to the entire task of performing that other procedure.

Step D.7.3. We have assumed that at this time three licensed operators are available to deal with the accident. One of them is performing the activities shown in Figure 4-7. Of the two operators remaining in the control room, one will have to go two levels above the control room to secure (close the switches) the purge dampers for the decay-heat pump rooms. If he performs this task, he will manipulate four MOV switches. (An error of omission is assigned to the manipulation of all four switches because they are all in the same procedural step.) Because of the poor layout of the ventilation room (no cues are provided as to the location of functional groups), selection errors for each of the four switches are assigned.

Step D.7.4. Verifying that the decay-heat pumps are on and verifying that the LPI MOVs are open are called out in the same procedural step. The equipment items are all located on the ESF panel. An error of omission is assigned for forgetting the task entirely. For the decay-heat pumps, the wrong items of equipment could be chosen or the indications on the correct items could be interpreted incorrectly. For the LPI MOVs, the wrong switches could be selected, or their indications could be interpreted incorrectly. Two errors of commission have been assigned to each item.

Step D.9. Monitoring the level of the borated-water storage tank, a dynamic task, cues the operator to perform the rest of this step. If he fails to monitor or if he monitors incorrectly, the other activities in this step will not be performed. An error of omission is assigned to the monitoring task only. A reading error is also assigned to the monitoring task. For the manipulation of the valves, errors of selection and interpretation or reversal are possible.

The errors assigned for the operations outlined in Figure 4-7 were determined in a slightly different manner. First, consider the fact that the auxiliary operator performs these actions in response to an order from the senior control-room operator. If the senior operator fails to order these tasks, they will not be performed. In developing the HRA event tree for this set of tasks (Section 4.5.6), this probable error will have to be considered. Regarding the rest of these tasks, the auxiliary operator must perform them on three different levels of the plant. He views his job at each level as a unit task; therefore, errors of omission apply to each of these unit tasks. If he remembers to stop at a given level, it is assumed that the operator will attempt all the tasks required at that level. Errors of commission are discussed below.

Step D.7.1. Manual valve MU-13 is the only valve located in the stairwell outside the makeup-pump room. No selection error is possible. It is not deemed likely that the operator will make a reversal error on a manual valve in this situation.

Step D.7.2. Valves DH-7A and 7B are outside the decay-heat pump rooms, one on each end of the hall. They are very large valves, and the only other valves in that area are too small to be confused with them. Of all the valves inside the decay-heat pump rooms, these are the ones that are located high on the walls of the rooms; the only other valves in the rooms are on piping lines that run along the floor. In none of these cases are errors of selection deemed likely.

Step D.7.3. Valves ABS-13 and 14 are located under the grating outside the watertight doors. They are the only valves there; likewise, there is only one set of watertight doors at this location. Again, selection errors are not considered likely.

4.5.6 DEVELOPMENT OF HRA EVENT TREES

4.5.6.1 Discussion

In making a probabilistic statement as to the likelihood of human-error events, each error defined as likely in the task analysis is entered as the right limb in a binary branch of the HRA event tree. Chronologically, in the order of their potential occurrence, these binary branches form the limbs of the HRA event tree, with the first potential error starting from the highest point on the tree at the top of the page. An example of an HRA event tree is shown in Figure 4-8.

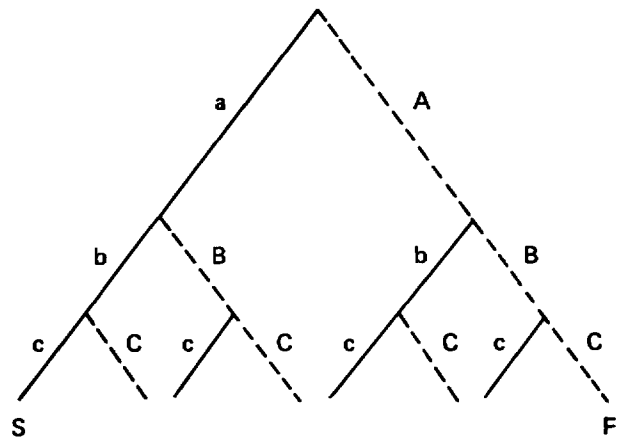

Figure 4-8. An example of HRA event-tree diagramming. Here A, B, and C are the first, second, and third tasks that are performed. Solid lines represent success; broken lines, error.

Any given task appears as a two-limb branch, with each left limb representing the probability of success and each right limb representing the probability of failure. (In a later phase of the human-reliability analysis, the human-error probabilities from the Handbook will be entered into the tree. See Section 4.5.7.) Once a task is diagrammed as having been completed successfully (or unsuccessfully), another task is considered; the binary branch describing the probability of the success (or the failure) of the second event extends from the left (or the right) limb of the first branch. Thus, every limb following the initial branching depicts a conditional probability. The initial branching also represents a conditional probability in that the probabilities for that branch are based on the existence of a given situation. However, it is defined as the starting point for the analysis, not as a conditional probability, since the analysis does not investigate the probabilities of occurrence of the circumstances of the basic situation. (As described in Chapter 5 of the Handbook, the conditional probabilities are understood in the labeling scheme shown in Figure 4-8; for example, a limb labeled b actually means b|a.)

Each limb of the HRA event tree is described or labeled, usually in a form of shorthand. Capital letters in quotation marks ("A") represent certain tasks themselves. Capital letters (A) represent failure or the probability of failure on given tasks. Lowercase letters (a) represent success or the probability of success on certain tasks. The same convention applies to Greek letters, which represent non-human-error events, such as equipment failures. The letters S and F are exceptions to this rule, in that they represent system success and failure, respectively. In actual practice, the limbs are sometimes labeled with a short description of the error itself. This eliminates the need for a legend at the bottom of the page that defines the alphabetic code for each event. The labeling format that is used is unimportant: the critical task in developing HRA event trees is the definition of the events themselves and their translation onto the trees. (Examples of labeling formats are shown in Figures 4-9 and 4-10.)

All the limbs of an HRA event tree are heavy solid lines in the diagram. For illustration only, the limbs representing failure in Figure 4-8 are shown as broken lines. (See Chapter 5 of the Handbook for a more complete discussion of the basics of HRA-event-tree diagramming.)

In a probabilistic risk assessment, the analyst is usually interested in determining the probability of error on a single task or the probability that, for a set of tasks, none or all will be performed incorrectly. For the first case, no HRA event tree need be developed unless performance on that task is affected by other factors whose probabilities should be diagrammed. A description of the task and knowledge of the performance-shaping factors are sufficient for entering Chapter 20 of the Handbook to determine the probability of a single human error.

For the second case, in which we want to know the probability of all tasks being performed without error, a complete-success path through the HRA event tree is followed (as discussed in Chapter 7 of the Handbook). Once an error has been made on any task, a criterion for system failure has been met. Given such a failure, no further analysis along that limb is necessary at

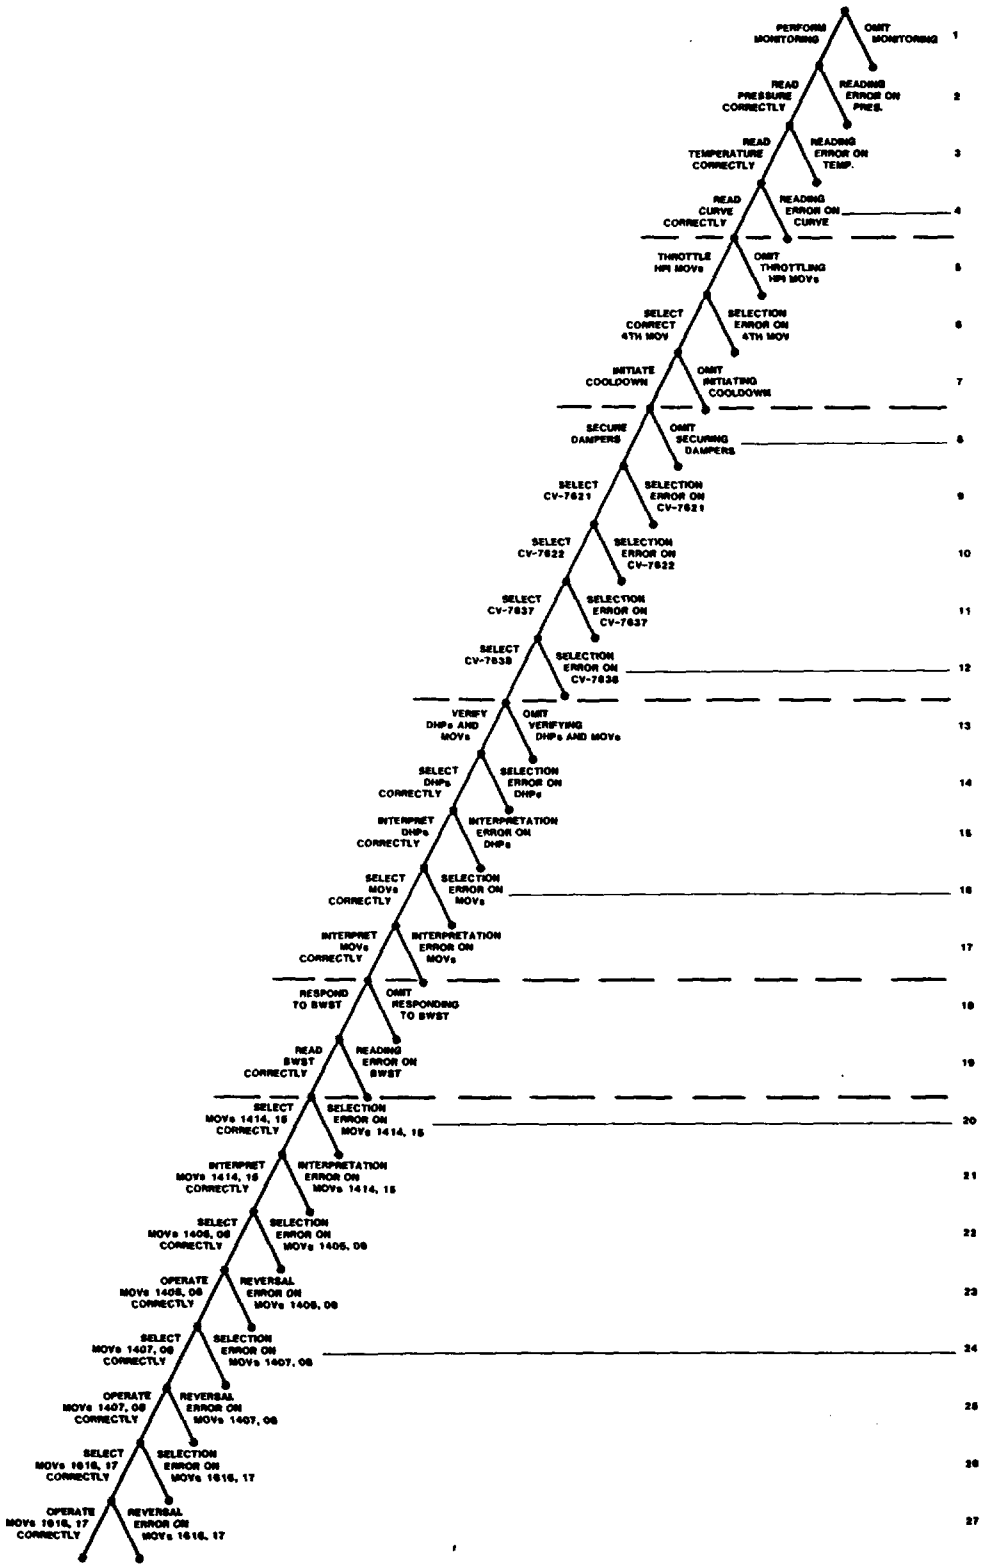

Figure 4-9. HRA event tree for actions by operators assigned to the control room.

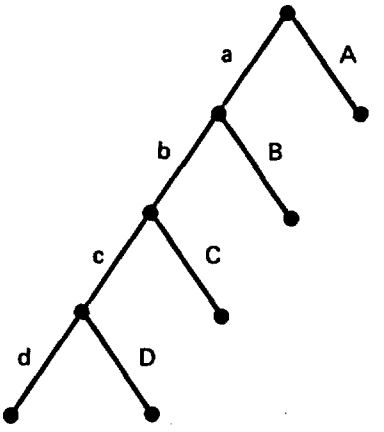

Event

A = Control-room operator omits ordering the following tasks

B = Operator omits verifying the position of MU-13

C = Operator omits verifying/opening the DH valves

D = Operator omits isolating the DH pump rooms

Figure 4-10. HRA event tree for actions performed outside the control room.

this point. In effect, the probabilities of event success that follow a failure and still end in a system-success probability constitute recovery factors and should be analyzed later, if at all. Thus, as shown in Figures 4-9 and 4-10, there are HRA event trees that are developed along the complete-success path only. This does not mean that we think this is the only possible combination of events; it means only that, in the initial analysis, we go no further once a system-failure criterion has been met.

The development of the HRA event tree is the most critical part of the quantification of human-error probabilities. If the task analysis has listed the possible human-error events in the order of their potential occurrence, the transfer of this information onto the HRA event tree is much easier. Each potential error and success is represented as a binary branch on the HRA event tree, with subsequent errors and successes following directly from the immediately preceding ones. Care should be taken not to omit the errors that are not included in the task-analysis table but might affect the probabilities listed in the table. For example, administrative-control errors that affect a task's being performed may not appear in the task-analysis table but must be included in the HRA event tree.

4.5.6.2 Example

The HRA event trees shown in Figures 4-9 and 4-10 represent the task analyses shown in Figures 4-6 and 4-7, respectively. Figure 4-9 (HRA event tree for actions by operators assigned to the control room) uses a labeling format that incorporates a short description of each event for its corresponding limb. Such a format is very convenient for analyses in which large numbers of events are diagrammed; referring back and forth to a descriptive legend would be inconvenient. The lines in Figure 4-9 are placed according to those found in the corresponding task-analysis table (Figure 4-6). Again, they are included to aid the system analyst in extracting information from the HRA event tree for inclusion in the system analysis. Figure 4-10 (HRA event tree for actions performed outside the control room) demonstrates that a format consisting of alphabetic labels and a descriptive legend can be used very effectively when a small number of events are involved. The legend format has the advantage of allowing a more complete description of the error events than does the short-label format. As already stated, however, the actual labeling format is of little importance as long as it is helpful to the analyst. Combinations of these two styles can be used, or entirely new formats can be developed by the analyst.

Both of the HRA event trees shown here reflect the technique described above and in Chapters 4 and 5 of the Handbook. The possible errors listed in the respective task-analysis tables have been put directly onto the right limbs of the branches. Only the complete-success paths are shown, as previously explained. The first branch of Figure 4-10 represents the administrative control error identified in the discussion of that set of tasks. In the HRA event tree itself, no distinction is made between the error events that appeared in the task-analysis table and those that were identified during other parts of the analysis.

4.5.7 ASSIGNMENT OF NOMINAL HUMAN-ERROR PROBABILITIES

4.5.7.1 Discussion

When the human errors have been identified, defined, and diagrammed, the analyst must estimate the probability of occurrence for each error. Since the analyst should be familiar with the theories, models, and limitations presented in the Handbook, he will be able to use Chapter 20 of that document for most of these estimates.

First, the task itself must be categorized. The analyst determines whether he is dealing with an operator manipulating valves, checking another's work, using a written procedure, or attempting some other type of task. Errors are then considered on the basis of their being of the omission or the commission type. In the tables in Chapter 20 of the Handbook, human-error probabilities (HEPs) are grouped by the type of error (omission or commission) that may occur in the performance of a certain type of task.

The analyst should become familiar with the organization of the HEPs in Chapter 20 of the Handbook. Some of the tabular data are duplicates of data presented in the subject chapters of the Handbook; others are condensations

of data found in several chapters. An analyst who becomes familiar with the organization of Chapter 20 before trying to use it as a source document will save a considerable amount of time. Furthermore, he will be able to establish beforehand the cases in which he will need to estimate HEPs directly from the task analysis because no such task is described in Chapter 20.

A description of each error identified for every task in the task analysis should be looked up in Chapter 20; that is, the description that most closely approximates the situation under consideration should be identified. In some cases, the description in Chapter 20 will detail a scenario that differs slightly from the one in the analysis. If the differences in specifics are not great, the analyst may decide that they are too minor to affect materially the use of the HEP as is. In other cases, the actual situation and the one described in Chapter 20 may reflect tasks that are basically the same but are performed under different circumstances. The HEP must then be modified to reflect the conditions of actual performance. Usually, this is done during the assessment of the performance-shaping factors acting on the task (Section 4.5.8).

If an HEP entered into the HRA event tree was not obtained from the Handbook, its source should be recorded, along with the assumptions made in its derivation. If Chapter 20 is the source of the HEP, the table number and item number should be recorded. If an HEP from the Handbook was used as a reference point for the derivation of an estimated HEP, its specific source and the reasoning behind its modification should be noted. For easy reference, this information can be added to the task-analysis tables in new columns. This documentation is necessary for many reasons. Other analysts may want to check the similarity of their solutions to other problems. Given that the estimates of many of the HEPs in the Handbook are numerically identical, these other analysts must have some method for tracing the original analysis. The assumptions should be recorded to prevent the analyst's needing to reinvestigate a situation should he need to refer to that analysis again. Also, in the course of performing a series of analyses on a single plant, some sections of an analysis may be used several times. The analyst must, however, be able to demonstrate that the situations are indeed identical before reproducing part of one analysis without modification in another.

In the HRA event tree of Figure 4-11 and in subsequent discussions and figures, results are shown to several decimal places merely to illustrate the arithmetic. In practice, final answers are subjected to judicious rounding.

As mentioned in Section 4.1.3, one of the limitations of the HEPs tabled in the Handbook is that nearly all of them apply to rule-based human actions. For cognitive errors related to the evaluation of display indications, the following interim guideline that should be used as a supplement to the 1980 issue of the Handbook is suggested: A generic estimate of .1 (.01 to .5) per operator should be used for the failure to evaluate an accident properly unless there is plant-specific information to the contrary--unless there is evidence that such errors are not likely to be characterized by an HEP of .1. (In applying this rule, appropriate estimates of the levels of dependence must be made to account for the presence of more than one operator in the control room.) It will be a matter of judgment as to whether modification of the generic HEP of .1 is necessary. For some kinds of abnormal conditions, there are

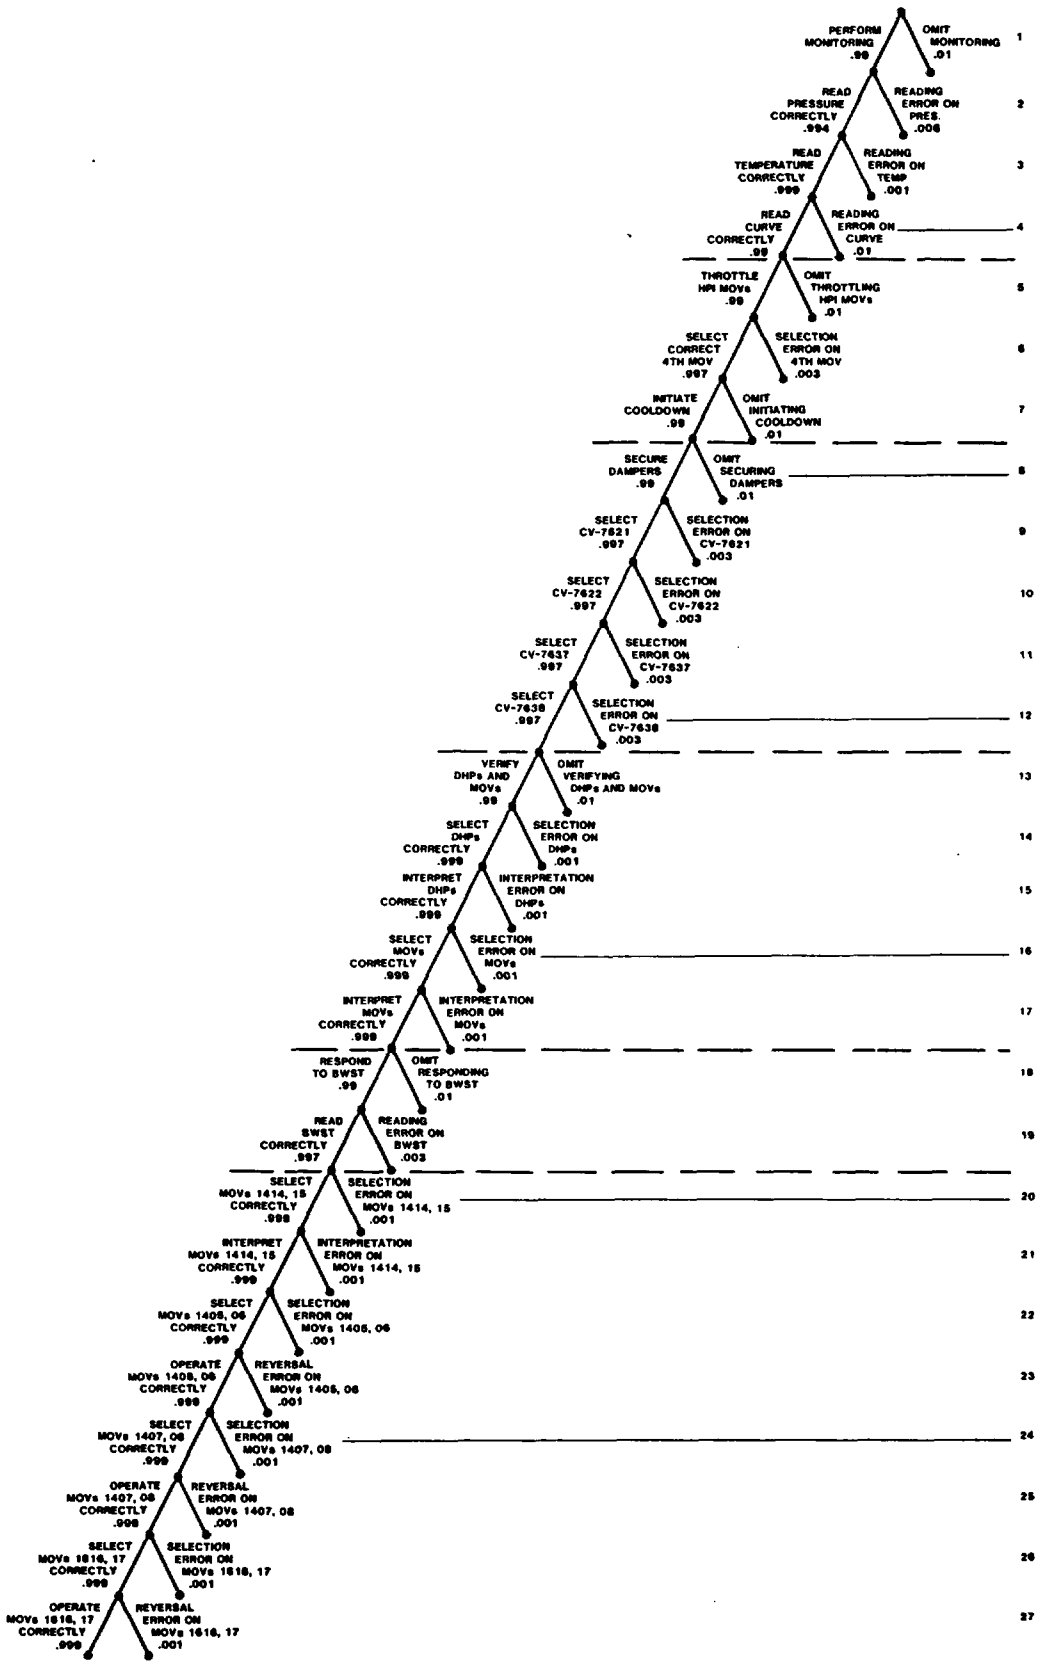

Figure 4–11. HRA event tree for actions by operators assigned to the control room, with estimates of nominal human-error probabilities.

plant-specific operating rules that, if rehearsed properly, will effectively eliminate any initial indecision on the part of the operator when an accident occurs.* In such a case, the main effort of the human-reliability analyst will be to estimate the effectiveness of the provisions for in-plant rehearsal of these operating rules. This type of treatment reflects the state of the art in human-reliability analysis and points to the need for studies of the type mentioned earlier in Section 4.1.3. (See also Section 4.9, "Alternative Methods," for discussions of other approaches to estimating the likelihood of such errors in the cognitive process.)

4.5.7.2 Example

In studying this example, it is necessary to keep in mind the situational characteristics that affect the performance of the tasks in question: the actions of operators who are following a set of written procedures. Any errors are made in the context of using those procedures. Recovery factors are not to be considered at this time. Even though there will be three licensed operators in the control room, this first analysis considers only the actions of one operator.

In the first part of this example, each error and the source of its estimated HEP are discussed in detail. Later in the example, only the source HEPs are given for errors that have already been discussed. Figures 4-11 and 4-12 are the HRA event trees diagrammed in Figures 4-9 and 4-10, but they include the HEP estimates for each error. As shown, this can be done by adding the HEP as part of the label for each limb or by including the HEPs in the legend for the HRA event tree. Again, the method employed for displaying the HEPs on the HRA event tree is unimportant.

The first error† on the HRA event tree is the operator's failure to perform the monitoring of RCS pressure and temperature. This is the first part of step D.2. If the operator fails to do this part of step D.2, it is presumed that he will fail to carry out the remainder of the step. The failure to maintain RCS temperature and pressure was designated a system failure by the system analysts. Since we are dealing with the operator's following a set of written procedures, we use an estimate of the error from Table 20-20 in the Handbook.‡ This table presents estimates of errors of omission made by operators using written procedures. In other words, these estimates reflect the probability, under the conditions stated, of an

---

*For an example, see the case study described on pages 21-11ff in the Handbook.

†References to error numbers correspond to the numbered events in all related HRA event trees and to like-numbered entries in the task analysis.

‡All cited table and item numbers are from the October 1980 draft of the Handbook.

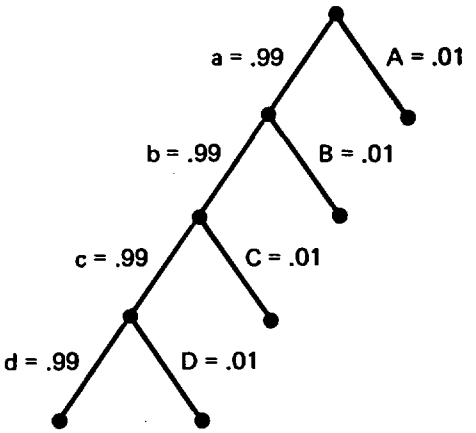

Event	HEP	Source
A = Control-room operator omits ordering the following tasks	.01 (.005 to .05)	Table 20-22, item 1 (p. 20-31)
B = Operator omits verifying the position of MU-13	.01 (.005 to .05)	Table 20-18, item 3 (p. 20-28)
C = Operator omits verifying/opening the DH valves	.01 (.005 to .05)	Table 20-18, item 3 (p. 20-28)
D = Operator omits isolating the DH pump rooms	.01 (.005 to .05)	Table 20-18, item 3 (p. 20-28)

Figure 4-12. HRA event tree for actions performed outside the control room, with estimates of nominal human-error probabilities.

operator's omitting any one item from a set of written procedures. Since the procedures in this example are emergency procedures that do not require any checkoff of steps by the operator, we use the section of Table 20-20 that deals with procedures having no checkoff provision. Looking at the procedures in Figure 4-3, we see that more than 10 steps must be performed by the operator. This analysis deals with fewer than 10 procedural steps, but the steps must be considered in the context of their performance. The fact that only a few steps are analyzed has no effect on the operator as he follows the set of procedures. Given that this error occurs in using a long list of written procedures that does not require a checkoff, its estimated HEP is .01 (.005 to .05), as given in item 5 of Table 20-20. At this point in the analysis, the nominal value of the HEP is entered into the HRA event tree.

The second error shown in Figure 4-11 is the operator's error in reading the indicator for RCS pressure. This indicator is a chart recorder. Reading errors are errors of commission and are grouped in Chapter 20 according to the type of information that is displayed and to the type of indicator that makes up the display. In this instance, the operator is reading a numerical value from the chart recorder. Table 20-5 presents estimated HEPs for errors made in reading quantitative information from different types of display.

For the chart recorder in question, item 3 from that table is used, .006 (.002 to .02).

The third error also involves reading an exact value from a display. In this case the display is a digital readout; therefore, item 2 from Table 20-5 is used, .001 (.0005 to .005).

The fourth error is also a reading error, this time involving the pressure-temperature curve. Since the curve is presented as a graph, the HEP for errors made in reading quantitative information from a graph is used, item 5 from Table 20-5, .01 (.005 to .05).

Another error of omission appears as the fifth error limb on the HRA event tree in Figure 4-11: the operator's not throttling the HPI MOVs. For errors of omission, the nature of the task does not affect the probability of the error. Therefore, the same HEP that was used for the first error, .01 (.005 to .05), is used again here.

A switch-selection error for the fourth of the HPI MOVs was identified as likely in the task analysis. It is the sixth of the errors on the HRA event tree. Figure 4-4, which shows the layout of the control panels containing the switches for the HPI MOVs, demonstrates that the HPI MOV switches are in similar positions on control panels CP16 and CP18. Surrounding them are several similar switches, one of which (to the immediate right of the switches for HPI MOVs on CP18) is the switch most likely to be the target of the selection error. An estimate of this error of commission is found by looking in the tables in Chapter 20 that deal with errors made in the manipulation of valves. Table 20-14 contains HEPs for errors of commission in changing or restoring valves. Since item 7 most closely approximates the situation described here, the HEP of .003 (.001 to .01) is used as the estimate for this error.

The seventh error involves an omission on the part of the operator to initiate cooldown by following another set of written procedures. As far as we are concerned here, this is a case of his omitting a single step of this procedure, so .01 (.005 to .05) is used again. It is also used for the eighth error, that of omitting to secure the purge dampers for the decay-heat pump rooms.

The 9th, 10th, 11th, and 12th errors are selection errors involving the manipulation of the switches for four MOVs. The switches are probably close to each other on a wall of the ventilation room, but we have no specific information about the ease or difficulty of locating the group. Since it is not known whether the layout and the labeling of the switches in the ventilation room help or hinder the operator in his search for the controls, we take the conservative position of assuming them to be among similar-appearing items. We use the same HEP as that used for the selection error associated with the fourth HPI MOV (error 6), .003 (.001 to .01), for each of these MOVs.

The 13th error is one of omitting a procedural step. The HEP of .01 (.005 to .05), discussed earlier, was used. If this procedural step is performed (is not omitted), errors of selection for both types of components

mentioned (the decay-heat-removal pumps and the LPI MOVs) are possible. These selection errors appear as the 14th and the 16th errors on the HRA event tree. We know from Figure 4-4 that both of these sets of controls are part of groups that have been arranged functionally on the control panels. They are very well delineated and can be identified more easily than can most of the switches in the control room. Since there is no entry in Table 20-14 (commission errors in changing or restoring valves) that accurately reflects this situation, an HEP from Table 20-13 is used. This table consists of HEPs for commission errors in manipulating manual controls (e.g., the hand switch for an MOV). Item 2 in this table involves a selection error in choosing a control from a functionally grouped set of controls; its HEP is .001 (.0005 to .005). (Note: On page 20-19 of the Handbook, please insert the words "locally operated" before the word "valves" in the second sentence. It is intended that the estimated HEPs in this table apply to switches of all kinds, including the control-room switches used to operate MOVs.)

Errors of interpretation are also possible for the decay-heat pumps and the LPI MOVs. Given that the operator has located the correct switches, there is a possibility that he might fail to notice their being in an incorrect state. In effect, this constitutes a reading error, one made in "reading" (or checking) the state of an indicator lamp. No quantitative information is involved, so Table 20-7, which deals with commission errors in checkreading displays, is used. The last item on that page describes an error of interpretation made on an indicator lamp, so .001 (.0005 to .005) is used. The 15th and 17th errors on the HRA event tree represent these interpretation errors.

The HRA event tree's 18th error is defined as the operator's omitting to respond to the level of the borated-water storage tank. The same omission HEP used previously, .01 (.005 to .05), is repeated here. Given no such omission error, a reading error (19 on the event tree) could be made on the BWST meter. Going back to Table 20-5 for commission errors made in reading quantitative information, the HEP to use in considering an analog meter is .003 (.001 to .01), the first term in the table.

Errors 20, 22, 24, and 26 involve selecting the wrong set of MOV switches from sets of functionally grouped switches. As above, this HEP is from item 2 of Table 20-13, .001 (.0005 to .005).

The 21st error (interpretation) is made while checking the status of an indicator lamp. An HEP of .001 (.0005 to .005) (as cited for the 15th error above) is assigned.

The 23rd, 25th, and 27th errors represent reversals made by the operator: instead of opening valves, he closes them, or vice versa. Since errors of commission for valve-switch manipulations are involved, Table 20-13 is used. Item 7 most closely describes this error; hence, the HEP of .001 (.0001 to .01) is used.

For the HRA event tree in Figure 4-12, we are analyzing actions performed outside the control room. The first error diagrammed is one of administrative

control and did not show up in the task analysis: the control-room operator omits ordering another operator to perform this set of tasks. Since the ordering of the tasks is his responsibility, this constitutes a failure to carry out plant policy. An HEP of .01 (.005 to .05) from Table 20-22, item 1, is used.

The second, third, and fourth errors shown in Figure 4-12 are errors of omission by the operator who actually performs the tasks. These tasks call for the manipulation of valves located on levels of the plant under the control room. We assumed that the operator will not be working from a set of written procedures (he will not take a copy of the procedures with him) but from an oral instruction by the control-room operator. The model accounting for errors of omission made in following a set of oral instructions will be followed. The data for this model are found in Table 20-18. It was stated in the discussion of the talk-through (Section 4.5.4) that the operator sees these as three distinct unit tasks, one to be performed on each of the three levels he must visit. We therefore assume that he must recall three tasks and use item 3 in the table, which shows an HEP of .01 (.005 to .05) for each of the tasks.

4.5.8 ESTIMATING THE RELATIVE EFFECTS OF PERFORMANCE-SHAPING FACTORS

4.5.8.1 Discussion

A primary consideration in conducting a human-reliability analysis is the variability of human performance. This variability is exhibited by any given individual in the performance of tasks over time (from day to day, from week to week, etc.). Variability also results from the performances of different personnel (from man to man, shift to shift, or from plant to plant). Variability is caused by performance-shaping factors (PSFs) acting within the individual or on the environment in which the task is performed. Because of this variability, the reliability of human performance usually is not predicted solely as a point estimate but is determined to lie within a range of uncertainty. A point value HEP for the PRA can be estimated by considering the effects of relevant PSFs for the task in question. The estimates provided so far in this chapter apply to nonstressful, normal working conditions. Modifications of these basic estimates can be made on the basis of guidelines provided in the Handbook.

The nominal HEPs are to be used when the scenario outlined in the Handbook reflects the situation being analyzed. If the plant situation is worse in terms of the PSFs or the response requirements than the one described in the Handbook, the HEP for that task should be higher than the nominal value. That is, if the analyst judges that the situation under study is more likely to result in error than the one outlined in the Handbook, he should use an HEP that is closer to the upper bound than the nominal is. Likewise, if a plant's situation is judged to be less likely to result in a human error, the analyst should use an HEP that is closer to the lower bound than the nominal is. However, in a safety analysis, one should generally avoid the optimism that results from using a lower HEP.

In judging these effects, the analyst should first consider the error events individually. For each error probability, a judgment must be made as to whether the nominal HEP should be used. The analyst should examine the performance situation for the factors that might affect each event. For errors of omission, for example, the analyst should search for cues or reminders that would make forgetting any item less likely or for poorly written procedures that would make forgetting an item more likely. For errors of commission, it is necessary to identify the elements of the performance situation that might affect the actions themselves or the operator as he performs them. For example, if the face of a display is such that reading it is unusually difficult, an HEP higher than the nominal value for reading errors for such a display should be assigned.

Next, the analyst should consider the influence of PSFs that have a global effect--those that affect the probability of error for all or most of the events in the analysis. Some models presented in the Handbook reflect the influences of these overriding PSFs. The most commonly encountered ones deal with stress and the operator's level of experience.

The data in the Handbook reflect by their organization the effects of some PSFs. For example, for errors of omission in using a written procedure, the distinction based on the availability of a checkoff provision is really based on the quality of the procedure as a PSF. Whether an available checklist is used properly is an example of the PSF of administrative control. Reading errors for displays are related to the difficulty of the reading task. In these cases, the effects (to some extent) of the PSFs have been already determined for the analyst.

4.5.8.2 Example

For evaluating the effects of PSFs on the individual error events, in each case the scenario described in the Handbook is appropriate for the imaginary plant of these examples, and therefore no modification of the nominal HEPs is necessary.

Now we must consider the effects of overriding PSFs--those that will affect all of the HEPs. It was stated in the original assumptions that the operators are experienced. Since they are following an emergency procedure, we will consider them to be under a moderately high level of stress. We see from Table 20-23 that the HEPs for experienced personnel operating under a moderately high level of stress should be doubled for discrete tasks and multiplied by 5 for dynamic tasks. Discrete tasks are defined as the tasks that require essentially one well-defined action by the operator. Dynamic tasks are those requiring a series of connected (continuous) subtasks; an example is monitoring an indicator over a period of time.

Figure 4-13 shows the HRA event tree for control-room actions with the nominal HEPs of Figure 4-11 modified to reflect the effects of a moderately high stress level. The only dynamic tasks in Figure 4-13 are those calling for monitoring activities: the monitoring of the RCS temperature and pressure indicators (tasks 2 and 3) and the interpolation of these values onto the

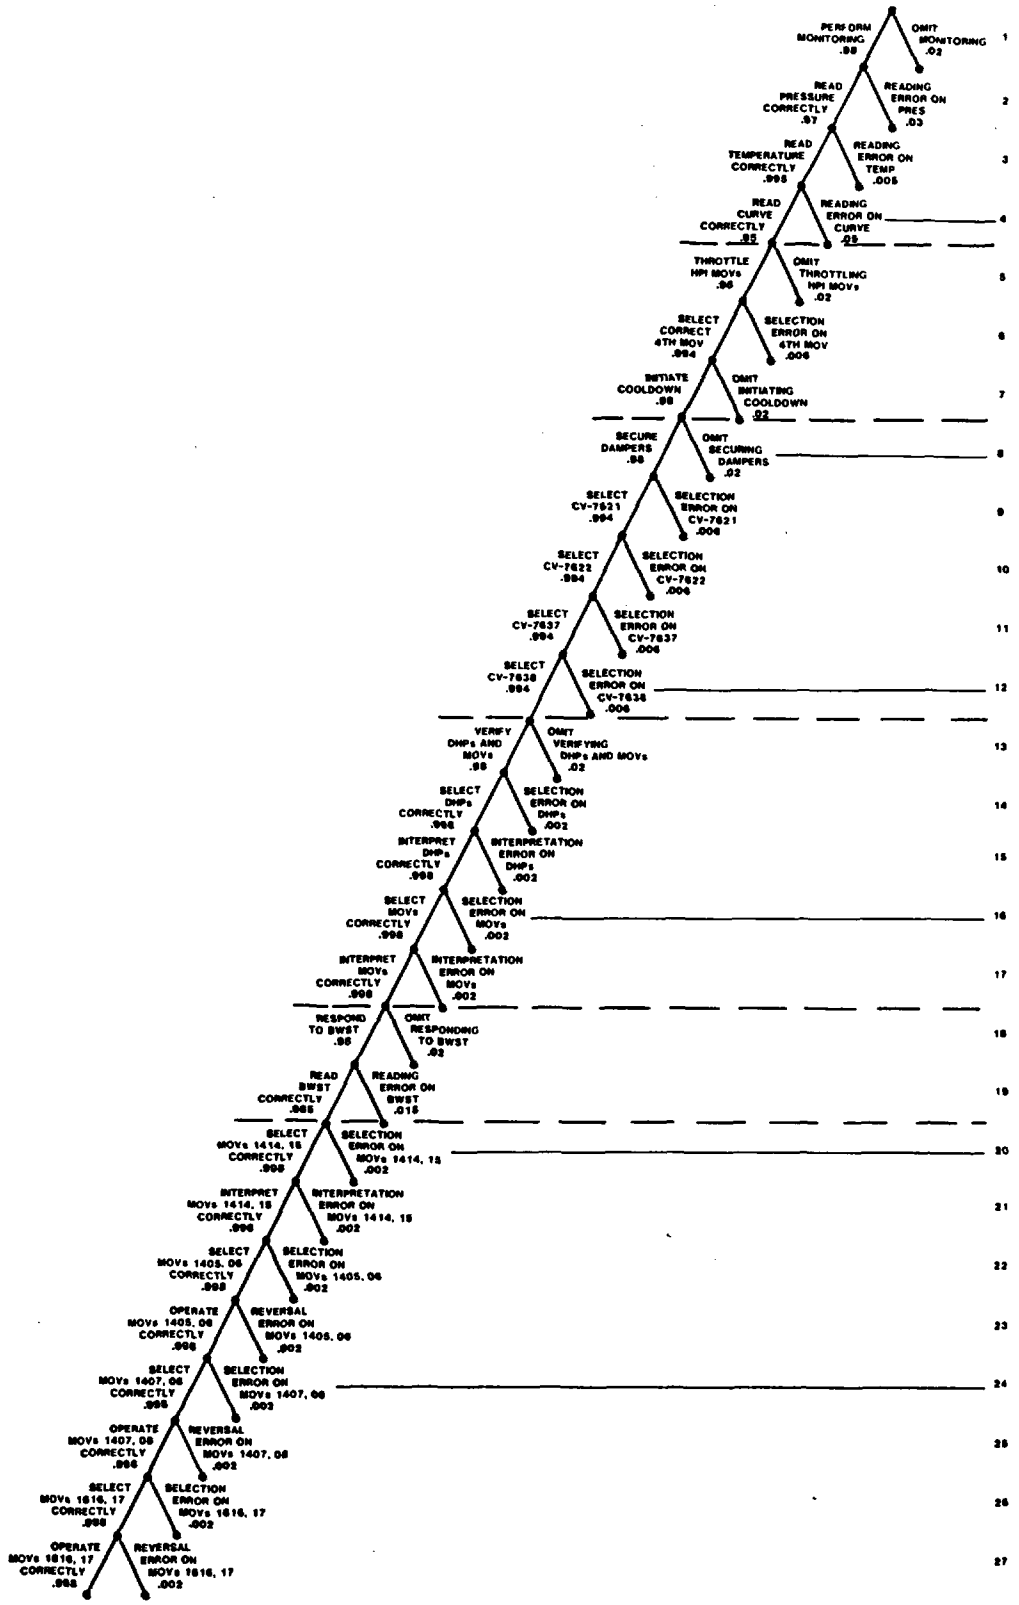

Figure 4-13. HRA event tree for actions performed by operators assigned to the control room, with human-error probabilities modified to reflect performance-shaping factors.

cooldown curve (task 4) and the monitoring of the BWST level (task 19). The nominal HEPs for these tasks have been multiplied by 5; those for the other events in this figure have been doubled.

Another overriding PSF that must be considered, this time for the tasks performed outside the control room, is the effect of the operator's having to wear protective clothing. If protective clothing is necessary, we assume that the operator is highly motivated to complete the task quickly because of the heat in the working environment, his isolation, and the general discomfort caused by the protective clothing. These factors combine to increase the HEPs for tasks performed by operators wearing such clothing. This is discussed on pages 3-8 and 17-7 of the Handbook. On the latter page, it is stated that the HEPs for such tasks should be doubled.

Figure 4-14 shows the events taking place outside the control room, with their HEPs modified to reflect these PSFs. The first error (failure of administrative control) takes place in the control room. The HEPs for this and for the other events have been doubled to reflect the effects of the moderately high stress level. The HEPs for the three tasks that actually take place outside the control room have been doubled again to reflect the effects of the operator's wearing protective clothing.

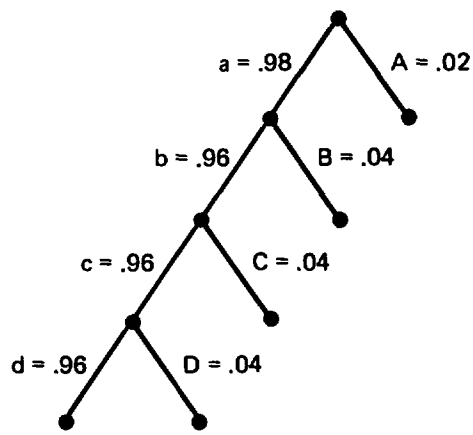

Event	HEP	Source
A = Control-room operator omits ordering the following tasks	.02 (.01 to .1)	Table 20-22, item 1 (p. 20-31)
B = Operator omits verifying the position of MU-13	.04 (.02 to .2)	Table 20-18, item 3 (p. 20-28)
C = Operator omits verifying/opening the DH valves	.04 (.02 to .2)	Table 20-18, item 3 (p. 20-28)
D = Operator omits isolating the DH pump rooms	.04 (.02 to .2)	Table 20-18, item 3 (p. 20-28)

Figure 4-14. HRA event tree for actions performed outside the control room, with human-error probabilities modified to reflect PSFs. The HEP for event A has been modified to reflect the effects of moderately high stress and dependence; the HEPs for events B, C, and D have been modified to reflect the effects of moderately high stress and protective clothing.

4.5.9 ASSESSMENT OF DEPENDENCE

4.5.9.1 Discussion

It has been stated earlier that, except for the first branch of an HRA event tree, all branches represent conditional probabilities of success and failure. Dependence between events directly affects these conditional probabilities. Some cases of dependence will be spotted during the talk-through, which is a good time to make note of equipment similarities that contribute to the level of dependence between actions performed on like items.

Dependence can occur between two performances with respect to errors of omission, errors of commission, or both. If dependence is assessed because two actions are called for in the same procedural step, dependence is likely to affect HEPs for errors of omission. If components are to be manipulated at different times in a given procedure, the dependence is likely to affect the HEPs for errors of commission, especially for selection errors. Common-cause dependence is likely to affect the HEPs for all types of errors. In effect, the overriding PSFs discussed in the preceding section are sources of common-cause dependence in that they result in modifications to all HEPs.

Guidelines for assigning the level of dependence are found in the dependence chapter of the Handbook. There are no cut-and-dried rules for this kind of assessment, but it must be made only after a carefully detailed study of the performance situation since it is highly situation-specific. The dependence level should be assessed for every task performed in every procedure targeted for human-reliability analysis. This is necessary because dependence may exist between one task considered during the analysis and one that is not. Given the performance context of each analysis, the effects of such dependence must still be quantified.

A decision as to whether complete dependence or complete independence applies to a given case can be made relatively easily. That is, it should be obvious that one action is the causal factor for another or that two actions are totally unrelated. Distinctions between the three intermediate levels of dependence are more difficult to make. First, we must decide whether there is any dependence at all--whether the actions are completely independent. If dependence does exist, we must decide whether complete dependence is appropriate and, if so, under what circumstances it applies. If we decide that the dependence is greater than zero but less than complete, an intermediate level must be assigned. This judgment can be based on the relation of the actual situation to zero and complete dependence. If we decide that the dependence is much closer to zero than to complete dependence, a low level of dependence is assigned. If, on the other hand, we decide that the situation exhibits a degree of dependence that is very close, but not equal, to complete dependence, a high level of dependence is assigned. If we cannot make a definitive statement to the effect that either of the above is true, moderate level of dependence is to be assigned.

Another method of assigning an intermediate level of dependence is to make a precise estimate as to the percentage of time the effects of zero or complete dependence will be seen. That estimate is used to assign the intermediate dependence level that most closely approximates it. For example, if we make a judgment (perhaps on the basis of a frequency count from actual

data or from our knowledge of the work situation) that task B will be performed correctly half of the time, given that task A has already been performed correctly, we have assigned a conditional probability of b|a = .5.

It should be remembered that the dependence model in the Handbook deals only with the effects and the quantification of positive dependence. If negative dependence is found to be appropriate to a situation, its effects will have to be determined directly rather than by using the dependence model. Furthermore, dependence is not necessarily symmetrical. The level of dependence may not be the same for the success and the failure paths of an HRA event tree.

The model presents some point estimates that can be used in lieu of the exact equations to determine the conditional probabilities of dependent events. These point estimates should be used only when the basic human-error probability (BHEP) is less than or equal to .01. In other cases, the equations should be used.

4.5.9.2 Example

In the sample problem, several cases of dependence have already been accounted for. For example, in the case of the four HPI MOV switches, their physical similarity, their positions in the procedure, and their location in relatively identical positions on the control panel led to our assumption that, for errors of omission, they are completely dependent. In considering dependence for the selection errors that could be made on these MOV switches, the same factors plus the layout of the rest of this control board led us to decide that the first three are completely dependent for selection errors (none are considered likely), and the fourth is susceptible to such an error. The nature of the tasks performed outside the control room and the operator's perception of them (from interviews with plant operators we determined that the operator typically views each set of tasks performed on a plant level as a single unit task) led to our considering them to be completely dependent with respect to errors of omission.

The presence of more than one operator in a given location constitutes a recovery factor. If we determine the effects of having more than one operator in the control room during the performance of this procedure, we are in fact quantifying a recovery factor for the procedure. However, since we will show that there is some level of dependence among the operators in the control room, we will quantify these effects now as an illustration of dependence.

According to Chapter 17 of the Handbook, one can assume that, after 20 minutes into an incident, three operators are present in the control room, with a moderate to high level of dependence between the two senior operators present and a high to complete level of dependence between the most junior operator and each of the two others. We have modified these assumptions to reflect the actual situation.

Since this procedure calls for the performance of several tasks outside the control room and since these tasks require the wearing of protective

clothing, we assume that one of the three operators will leave the control room during the entire procedure to prepare for and then perform these tasks. We assume that this will be the most junior operator in the control room since the other two are more capable of handling the incident from the control room. Responding to the nature of the control-room tasks, we assumed high dependence between the operators there. This assumption is based on the fact that, at this time in the incident, one of the operators will be involved mainly in directing the actions of the junior operator as he changes the positions of locally operated valves. Telephone communication between the two will call for most of this operator's concentration as he describes the necessary operations. The other control-room operator will be involved with monitoring the displays and performing the manipulations necessary at the ESF panels. High dependence is assumed because we judge that the operator on the telephone will, for the most part, rely on the operator at the ESF panels to perform those tasks correctly. Nevertheless, we judge that despite his primary task of coordinating the junior operator's tasks by telephone, this operator will catch errors made by the other control-room operator about half the time.

Figure 4-15 shows the HRA event tree of the actions performed by the control-room operators, with the HEPs (already modified to reflect the effects of performance-shaping factors) modified to reflect the effects of dependence. The probabilities of error for both the available operators have been collapsed onto a single limb for each type of error. The numbers in parentheses (shown for illustration only) are the conditional HEPs for the second operator's making the same error as the first. The other numbers are the products of these conditional HEPs and the basic HEPs of the first operators, and thus they represent the probability of both operators committing each error. The actions taking place in the ventilation room do not demonstrate any dependence between operators since we assume that one operator will be performing them. The only event in Figure 4-16 that is affected by dependence is the first. If the senior control-room operator forgets to order those tasks, the other senior operator or the junior operator himself may remind him of the necessity to do this.

4.5.10 DETERMINATION OF SUCCESS AND FAILURE PROBABILITIES

4.5.10.1 Discussion

Once the human-error events have been identified and quantified individually, their contribution to the probabilities of system success and failure must be determined. All paths in an HRA event tree should be defined as resulting in system success or failure in terms of their possible system consequences, not in terms of the specific human errors leading to these consequences. The system analysts will have identified the human-system interfaces to be analyzed in the human-reliability analysis, but errors made in operating at these interfaces may not significantly degrade system reliability or safety. For example, an error made in manipulating a system-critical component may not result in system failure as defined by the system analysts. The human-reliability analyst must point out potential human errors for a given set of tasks and then must quantify the probability of these errors; he does not,

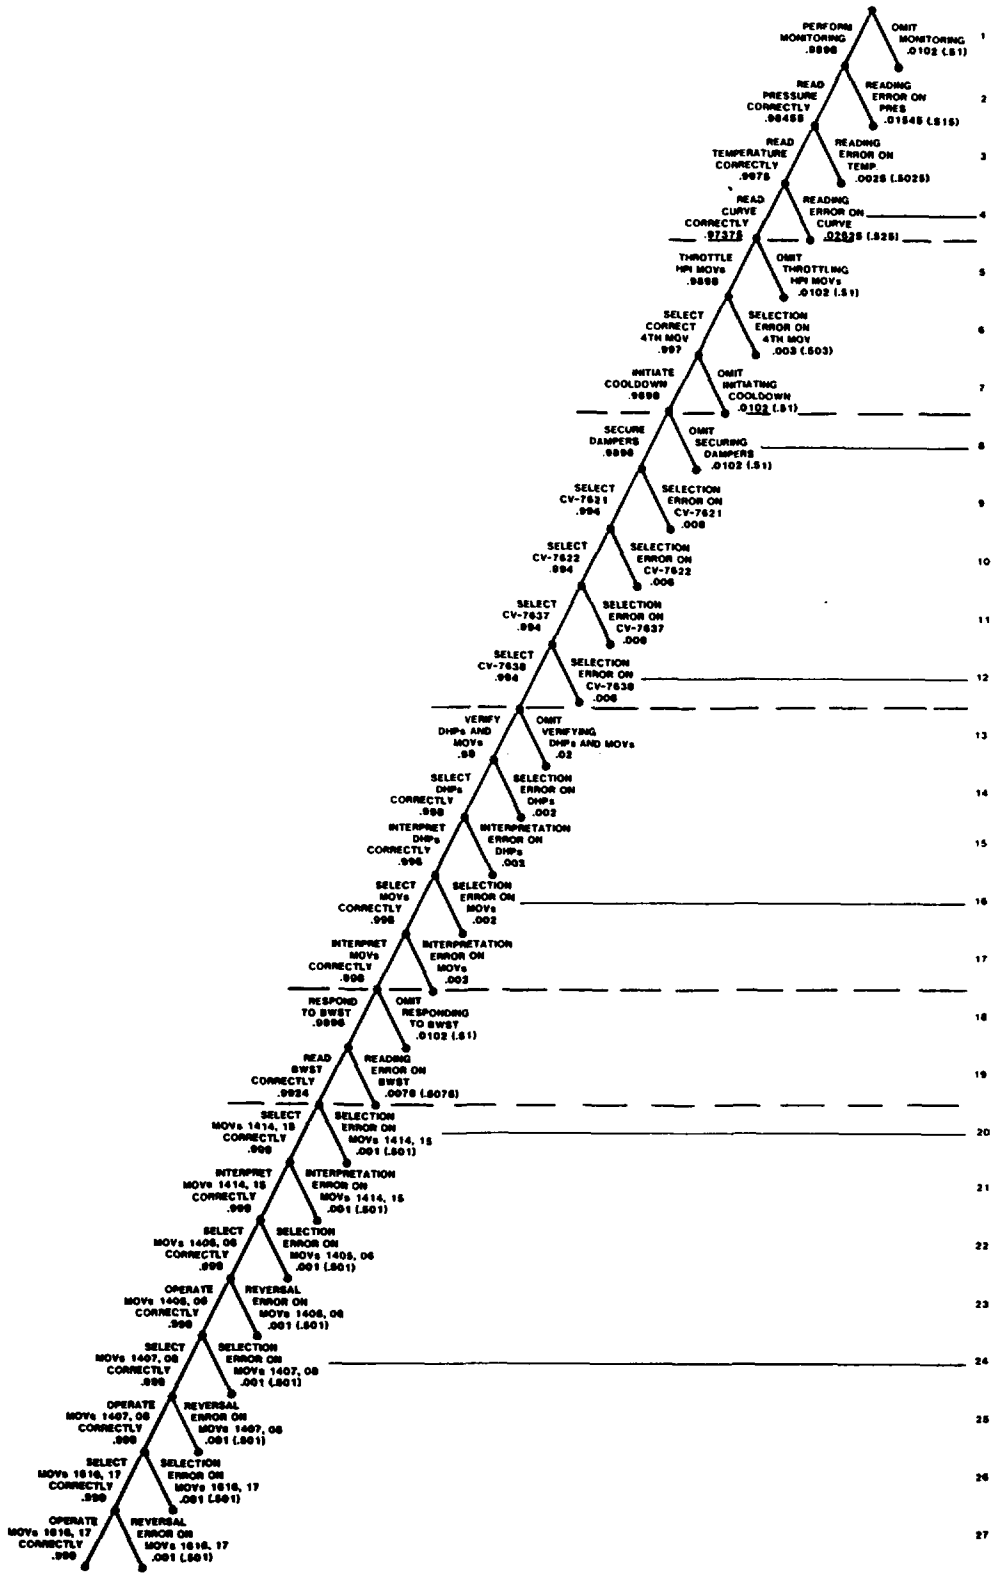

Figure 4–15. HRA event tree for actions by operators assigned to the control room, with human-error probabilities modified to reflect dependence. (Refer to page 4–47 for an explanation of the numbers in parentheses.)

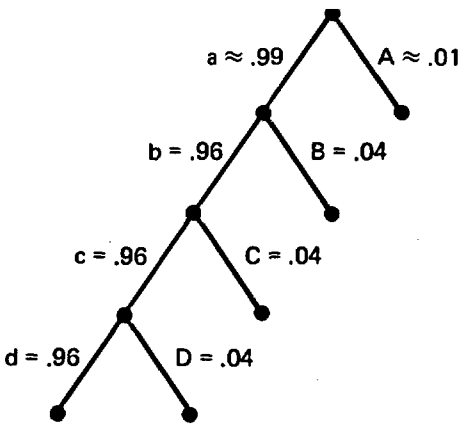

Event	HEP	Source
A = Control-room operator omits ordering the following tasks	.01 (.005 to .05)	Table 20-22, item 1 (p. 20-31)
B = Operator omits verifying the position of MU-13	.04 (.02 to .2)	Table 20-18, item 3 (p. 20-28)
C = Operator omits verifying/opening the DH valves	.04 (.02 to .2)	Table 20-18, item 3 (p. 20-28)
D = Operator omits isolating the DH pump rooms	.04 (.02 to .2)	Table 20-18, item 3 (p. 20-28)

Figure 4-16. HRA event tree for actions performed outside the control room, with human-error probabilities modified to reflect dependence. The HEP for event A has been modified to reflect the effects of moderately high stress and dependence; the HEPs for events B, C, and D have been modified to reflect the effects of moderately high stress and protective clothing.

however, decide whether a given sequence through the HRA event tree will contribute to system success or failure.

At this point in the human-reliability analysis, the system analyst should examine the HRA event tree for discrepancies between his understanding of the system and the human-reliability analyst's representation of it. He should consider the implications of each path through the HRA event tree, and then he should label each end point of the tree as a system success or failure. These end points should be quantified as probabilistic statements; the statements will be combined to formulate total system success and failure probabilities. This examination of the HRA event tree by the system analysts could be performed during the early stages of the human-reliability analysis or during the initial screening of the system. It is done here for illustrative purposes.

4.5.10.2 Example

After deciding which errors contribute to system failure probabilities, the system analyst made the following adjustments for Figure 4-15 (the final analysis to this point of the actions performed by the control-room operator): he defined the paths ending in error events 1, 2, 3, 4, 7, 18, 19, 22, and 23 as system failure and those ending in error events 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 20, 21, 24, 25, 26, and 27 as system success. Since the implications of the accident at Three Mile Island Unit 2 have great potential impact on error events 5 and 6, these error events were removed from the analysis at this point, to be considered separately.

For the HRA event tree of Figure 4-16, a similar decision was made by the system analyst. He decided that all of the paths terminating in a human error constituted contributions to system failure.

Once the paths that result in system failure have been determined, total system success and failure probabilities can be quantified in either of two ways. The first method is the simpler, requiring no redrawing of the HRA event trees. In it, the end points of the limbs on the existing HRA event tree are simply labeled as success or failure. All of the terminal success probabilities are summed to reach the total system success probability. The failure probabilities are obtained by the same method or by subtracting the total system success probability from 1.

The second method is more complex and requires that the HRA event tree be redrawn. When error on a human task does not contribute to system failure, both limbs representing this task on the HRA event tree contribute to the probability of system success. Algebraically, a probability of 1 is being multiplied by the system success probability since the results of paths going through both limbs are combined into the system success probability. In effect, that error has no influence on system failure. Therefore, we need not even consider it since we are concerned with estimating the probability of system failure in a risk assessment. The branches that represent events whose outcomes do not contribute to total system failure probabilities can be deleted from the HRA event tree altogether. The tree should be redrawn, diagramming only the events that have some effect on the probability of system failure. Figure 4-17 shows how the HRA event tree for actions performed by the control-room operators is changed when this second method for quantifying total system success and failure probabilities is used.

4.5.11 DETERMINING THE EFFECTS OF RECOVERY FACTORS

4.5.11.1 Discussion

Complete analyses are performed for the dominant sequences that show up in the computer modeling of the fault trees. To save time and effort in the human-reliability analysis, the effects of recovery factors are not considered until it is determined that a given analysis is part of a potentially dominant sequence. The probability of system failure due to human error will certainly be higher when recovery factors are ignored than when they are included. If the situation being analyzed does not appear as a potentially dominant sequence

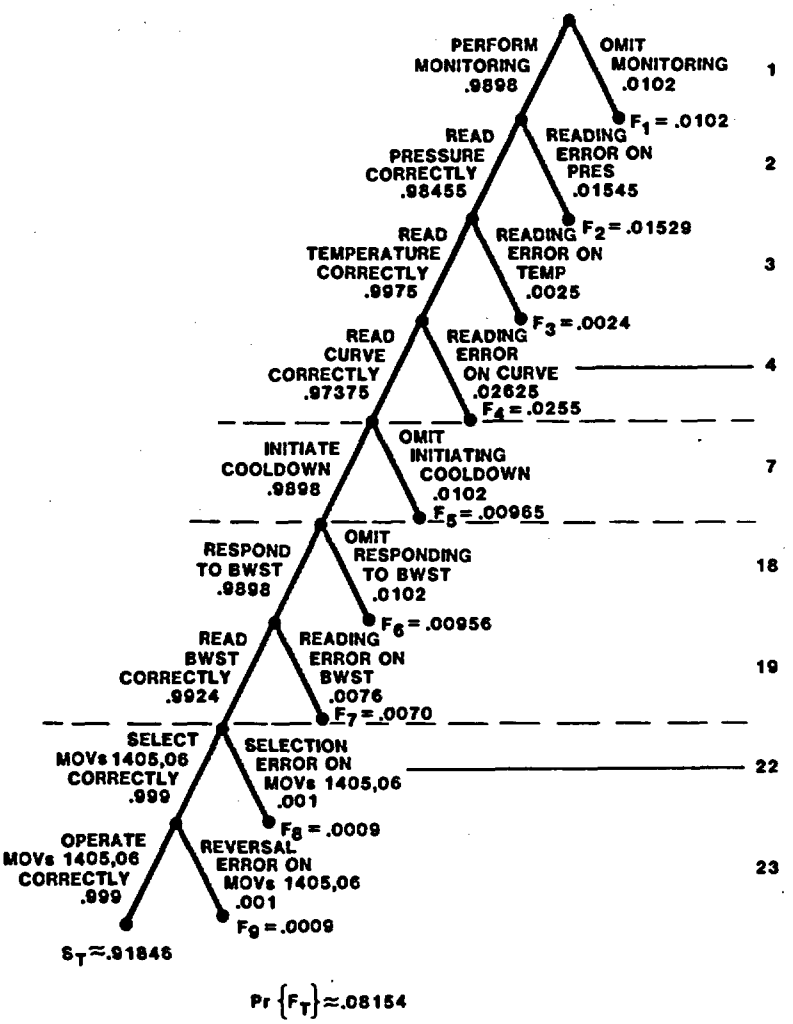

Figure 4-17. HRA event tree for actions by operators assigned to the control room, modified by second method for quantifying system success and failure probabilities.

when this inflated system failure probability is used, there is no need to analyze it further. In fault-tree terms, the frequency of an accident sequence can only be decreased by considering recovery factors.

To decrease the actual number of human-reliability analyses that must be performed for each plant, it is recommended that recovery factors not be included in the preliminary analyses. Once potentially dominant sequences have been identified, recovery factors for each can be added to see whether a complete representation of the system as it operates will eliminate the potential dominance. The incorporation of recovery factors can be done in stages, the purpose being to decrease the amount of time required for each analysis. If there are five recovery factors for a given scenario, the human-reliability analyst may choose to model only two of them at first. If the inclusion of these results in that sequence's ceasing to be potentially dominant, no more work need be done at this time. If this scenario still shows up as one of the system's potentially dominant sequences, the other three recovery factors should be analyzed.

Some recovery factors are highly situation-specific, while others can be applied generically. Alerting cues for recovery actions for any given incident will always depend on the specifics of response requirements for that incident. However, when analyzing recovery factors operating after maintenance activities it will sometimes be possible to generate HRA generic event trees that can be applied without modification to every such case for that plant. This is possible because, in many plants, a single procedure dictates the steps to be followed in restoring components after maintenance. In either case, the recovery factor can take the form of a point value (an HEP) or of a separate HRA event tree. The point value or the total success probability of the recovery HRA event tree should be inserted onto the associated error limb of the main HRA event tree. The probability of error for that limb is then multiplied by the success probability of the recovery HRA event tree and by the probabilities of the other events in that path to obtain the probability of recovery from the error. The end point of the original system failure path for that error is multiplied by the failure probability for the recovery factor to obtain the probability of an unrecovered error.

4.5.11.2 Example

As mentioned earlier, human redundancy as a recovery factor has already been analyzed for this problem to demonstrate the quantification of the effects of dependence. We can now consider situations in which the operator could catch his own errors or in which another operator working at a later date could catch his errors. An example would be an inspection process like the walk-around (see Chapter 8 of the Handbook). Since this problem deals with responding to an emergency, however, it is not appropriate to use the walk-around as a recovery factor. It is also possible for the operator to catch his own errors when the situation provides some additional alerting cue either to the action that should be taken or to the error itself.

In this problem and from the procedures in Figure 4-3, we see that the operator should respond to the BWST level's falling to 6 feet. His response is cued from two sources: if he is following the written procedures correctly, he will be monitoring the meter indicator of the BWST level; if he is not using the written procedures, there is still a possibility that the low-low-level alarm (annunciator) will remind him that he needs to perform the follow-up actions. We will treat the alarm as an additional alerting cue and analyze its effect as a recovery factor. From Chapter 20 of the Handbook, we need to find an estimate of an HEP for response to an annunciator. Table 20-4 lists HEPs for failing to respond to one of any number of annunciating indicators. We have no exact information on this, but assume that at this time into the incident 10 annunciators are alarming. The probability of the operator's failing to respond to any one of these 10 is .05 (.005 to .5). Figure 4-18 shows the diagramming for this recovery factor. Note that its inclusion in the analysis increased the unrounded probability of total system success from .91846 to .92746. If this is an adequate increase (if the sequence does not prove to be potentially dominant when the success probability is .92746), no more recovery factors need be analyzed.

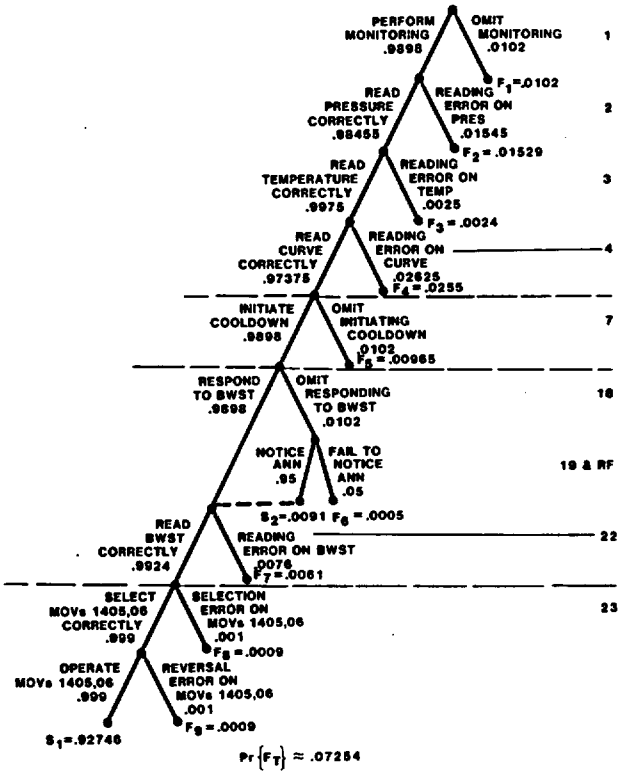

Figure 4-18. HRA event tree for actions by operators assigned to the control room, including one recovery factor.

4.5.12 SENSITIVITY ANALYSIS

4.5.12.1 Discussion

At times during the course of a human-reliability analysis, the analyst will want to determine the effects of manipulating the values of one or more of the elements analyzed. He may do this because he has some reservations about the assumptions he made, because the data he used are very uncertain (e.g., estimates of diagnosis errors by control-room personnel), or because he has not been able to obtain detailed information about some set of performance-shaping factors he judges are important determiners of the reliability of a task he has to analyze. Changing the assumptions of the analysis or changing the values of certain parameters may affect the probabilities of system success and failure. It may be of interest to manipulate these values to determine the effects of changes in design or procedures before such changes are made.

If the probabilities of some errors in an analysis stand out with respect to those of others, the analyst may want to see what effect lower probabilities for these errors would have on total system success and failure probabilities. The HEPs can be decreased by the action of recovery factors (see Section 4.5.11) or by changing the characteristics of the task to reflect a situation in which an error is less likely. These changes can be

accomplished by improving man-system interfaces, by increasing feedback adequacy, or by upgrading the quality of associated procedural steps. The new, lower HEPs can be entered onto the HRA event tree, and the resulting differences in total system success and failure probabilities evaluated. Sensitivity analyses are extremely useful in tradeoff analyses of proposed design changes and in pinpointing areas of potential system improvement.

In performing best- and worst-case analyses for a PRA, a bounding analysis can be executed, as described in detail in the appendix to NUREG/CR-2254 (Bell and Swain, 1981). For this exercise, two sets of HEPs are used and the results of the two analyses compared. The upper and lower bounds of the nominal HEPs for a given situation can be used, or two sets of assumptions and PSFs relating to the situation can be defined. The results of these two analyses can be evaluated by entering them onto the appropriate fault tree to see how sensitive some part of the PRA is to the two sets of HEPs. For PRA, the criterion for evaluating the sets of results should be risk significance. If there is very little difference in outcome, the analyst may decide to select the more conservative set for inclusion in the final PRA, at least as a temporary measure. If the difference in outcome is considerable, he should take steps to obtain better data.

4.5.12.2 Example

In this problem, the two most important errors, in terms of their probabilities, are errors 2 and 4, reading errors on the RCS pressure chart recorder and the graph of the pressure-temperature curve. Suppose we want to find out, as a design tradeoff comparison, whether changing either or both of these tasks to result in lower task HEPs is worthwhile in terms of system success probability. The simplest change involves changing the nature of the displays themselves to make reading errors less likely. For RCS pressure, the display could be a digital meter instead of a chart recorder. From Table 20-5 in the Handbook, we see that this would change the basic HEP for that task from .006 (.002 to .02) to .001 (.0005 to .005). This new HEP of .001 must be modified to .005 (.0025 to .025) to reflect the effects of stress and then modified again to reflect the effects of dependence, becoming .0025 (.001 to .01). Using the .0025 instead of the .01545 for this HEP results in a total system success probability of .9396 as opposed to .927.

If we make the same sort of adjustment for error 4, we might redesign the graph so that it is comparatively easy to read. If we now use the lower bound of the HEP in Table 20-5, item 5, instead of the nominal value, we have .005 (.002 to .02). This becomes .025 when modified for stress and .0128125 when modified for human redundancy. Modifying only this graph results in a total system success probability of .9402.

For a larger increase in the total system success probability, we could analyze the effects of both changes. An HRA event tree with these new values is shown in Figure 4-19. The total system success probability becomes .95262. Whether the new estimate of the probability of system success is large enough to warrant the incorporation of both changes is, of course, a management decision.

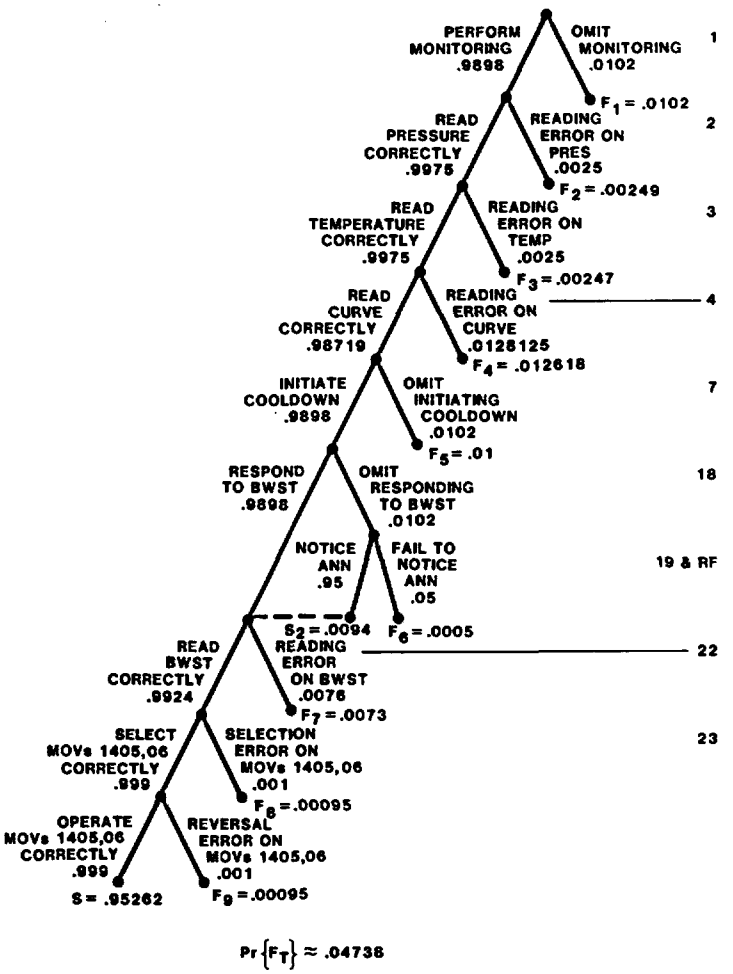

Figure 4-19. HRA event tree for actions by operators assigned to the control room, with tasks 2 and 4 modified.

4.5.13 SUPPLYING INFORMATION TO SYSTEM ANALYSTS

4.5.13.1 Discussion

All of the information used in performing the human-reliability analysis, especially the assumptions made and the modified HRA event trees, should be presented to the system analysts. The human-reliability analyst should then go over his analysis with them to ensure that there are no misunderstandings--no unresolved conflicts between the two concepts of the operating system. The system analyst should be familiar enough with the basic principles of HRA event-tree diagramming that he can use the HRA event tree itself to obtain the necessary inputs for his analyses. He should be able to use the total system success and failure probabilities or an HEP for a single item of equipment or for a single error for a given piece of equipment. These values can be entered directly into the human-error blocks of the system fault trees. The sources of the HEPs may be of interest to the system analysts, but are not strictly necessary. Section 4.6 discusses the method for formatting this information so that it is usable.

Any dependence found by the human-reliability analyst should be specifically indicated to the system analysts, especially in the case of dependence between different items of equipment. When dependence exists because of two operators performing the same task, combined HEPs representing the performances of both are entered into the human-error block of the fault tree--no change in the system fault-tree model is necessary. When dependence exists between performances on different items of equipment, the fault trees must be modified to reflect this common-mode failure. Identifying where and between which system elements the dependence exists will enable the system analyst to modify his models accordingly.

4.5.13.2 Example

If the system analyst needs an HEP for the entire procedure outlined in Figure 4-19, he should use the total system success probability, .962. If he needs a value for all possible human errors made in operating MOVs 1405 and 1406, he must consider all three of those diagrammed: the error of omission for the entire step (18), the selection error (22), and the reversal error (23). In effect, the combination of these errors represents a small HRA event tree. The system analyst must use the product of the success probabilities for each error event, .988, as the probability of success on those components. If the system analyst were only interested in the likelihood of an error of omission when dealing with MOVs 1405 and 1406, he would use the HEP for that specific error, .0102.

The human-reliability analyst should point out to the system analyst that MOVs 1405 and 1406 are completely dependent for all errors considered in the analysis. They (as a single item of equipment) are also dependent on the monitoring task (18): an equipment failure of the BWST meter would result in an error on MOVs 1405 and 1406.

4.6 METHODS OF DOCUMENTATION

The results of the human-reliability analysis go directly into the system analyses as probability statements. The only HRA data that are used in the rest of the risk assessment are the HEPs for given error events or for total system success and failure probabilities, and the information on dependence (where and what kind). The most important part of any final HRA report is the cataloging of the HEPs by item (of equipment) or by procedure, depending on the level of detail in the system fault trees and the system event trees, and the pinpointing of existing dependence. Other information included in the final report is not necessary as an input to the analysis itself, but is instead necessary as a reference on the performance of any particular human-reliability analysis.

Other human-reliability analysts must be able to trace through the analyses and to understand them fully. To obtain the necessary information, they must have access to the material on which the analysis was based. The

analyst should therefore provide in the final HRA report a set of the written procedures analyzed or his written version of the "standard operating procedure," along with the assumptions made in defining the situation under which the procedure would be performed. These assumptions will have been made during the visit to the plant and during the talk-through of the procedures with plant personnel. A copy of the final HRA event tree resulting from the analysis should be included. The basic HEP for each limb of the tree and its source as well as the source for any modifications (performance-shaping factors, dependence) should be included. This information can be added to the table of the task analysis; this is a clear, concise method for presenting a definition of the error events found in the HRA event tree. If recovery factors were considered or a sensitivity analysis was performed, the outcomes of these should be included.

In short, the final report should include all information necessary for the system analyst to check his assumptions about the performance situation against the human-reliability analyst's. It should also include sufficient information so that another human-reliability analyst could analyze the same scenario and arrive at a similar result.

4.7 DISPLAY OF FINAL RESULTS

As mentioned in Section 4.6, the most efficient method for displaying the results of a human-reliability analysis is to use the task-analysis format shown in Figures 4-6 and 4-7. These tables can be expanded to include the other information necessary for a complete documentation, as shown in Figures 4-20 and 4-21 for the example that was worked in this chapter. With these tables and copies of the HRA event trees, the system analysts should be able to take information in any form or at any level needed for input into the fault trees. The expanded task-analysis tables, HRA event trees, list of assumptions, and copy of the procedure should provide sufficient documentation for a human-reliability analysis.

This type of complete documentation of a human-reliability analysis is important for PRAs to be performed at various times in the life of a plant. As the plant equipment, manning, or operations change over time, the PRAs reflecting the different assumptions become points of comparison for the effects of these changes.

4.8 UNCERTAINTY AND VARIABILITY IN HUMAN-RELIABILITY ANALYSIS

Each estimate of a human-error probability for the performance of a task or activity is associated with some degree of uncertainty. Therefore, each such estimate is bounded by some range of values that is judged to have a high probability of encompassing the actual value of any given performance. This section discusses various sources of this uncertainty and

Step	Equipment	Action	Indication	Location	Notes	Errors	HRA event tree	HEP	T,I[a]	Final[b]
D.2	RCS pressure	Monitor		CB4		1. Omission (all)	1	.01	20, 5	.0102
						2. Reading	2	.006	5, 3	.01545
	RCS temperature	Monitor		CB4		Reading	3	.001	5, 2	.0025
	heater switches	Maintain pressure and temperature	Within curve on chart	CB4		Reading	4	.01	5, 5	.02625
D.4	4 HPI MOVs	Override and throttle		CP16, CP18	ESF	1. Omission (all)	5	.01	20, 5	.0102
						2. Selection (1)	6	.003	14, 7	.003
		Initiate cooldown	Procedure 12			Omission	7	.01	20, 5	.0102
D.7.3	CV-7621,22,37,38 (room-purge dampers)	Secure	Close switches	Ventilation room		1. Omission (all)	8	.01	20, 5	.0102
						2. Selection (each)	9,10,11,12			
D.7.4	DH pumps	Verify on	Indicator lamps	CP16, CP18	ESF	1. Omission (for MOVs too)	13	.01	20, 5	.02
						2. Selection	14	.001	13, 2	.002
						3. Interpretation	15	.001	7, 9	.002
	MOV-1400, 1401	Verify open	Indicator lamps	CP16, CP18	ESF	1. Selection	16	.001	13, 2	.002
						2. Interpretation				
D.9	Borated-water storage tank	Monitor level	>6 feet	CP14		1. Omission	18	.01	20, 5	.0102
						2. Reading	19	.003	5, 1	.0076
	MOV-1414, 1415	Verify open	Indicator lamps	CP16, CP18	ESF	1. Selection	20	.001	13, 2	.001
						2. Interpretation	21	.001	7, 9	.001
	MOV-1405, 1406	Open	MOV switches	CP16, CP18	ESF	1. Selection	22	.001	13, 2	.001
						2. Reversal	23	.001	13, 7	.001
	MOV-1407, 1408	Close	Switches	CP16, CP18	ESF	1. Selection	24	.001	13, 2	.001
						2. Reversal	25	.001	13, 7	.001
	MOV-1616, 1617	Close	Switches	CP16, CP18	ESF	1. Selection	26	.001	13, 2	.001
						2. Reversal				

[a]These numbers refer to table and item numbers in Chapter 20 of the Handbook.

[b]The nominal HEPs have been modified to reflect the effects of a moderately high stress level and (in some cases) high dependence between two operators.

Figure 4–20. Display of final results in a task-analysis table for actions by operators assigned to the control room. The column labeled "HRA event tree" does not usually appear in a task analysis; it has been included for the reader's convenience. The numbers in this column refer to the error event numbers appearing in HRA event trees starting with Figure 4–9.

Step	Equipment	Action	Indication	Location	Notes	Errors	HRA event tree	HEP	T,I[a]	Final[b]
D.7.1	MU-13	Verify closed	Position	Stairwell outside makeup pump room	Only valve	Omission	2	.01	18, 3	.04
D.7.2	DH-7A, 7B	Open	Position	Outside DH pump rooms		Omission (for all D.7.2)	3	.01	18, 3	.04
	MU-14, 15, 16, and 17	Verify open	Position	DH pump rooms						
	MU-23, 24, 25, and 26	Verify open	Position	DH pump rooms						
D.7.3	ABS-13, 14	Close	Position	Outside DH pump rooms	Only valve	Omission (for all D.7.3 here)	4	.01	18, 3	.04
	Watertight doors	Close	Locks in place	DH pump rooms						

[a]These numbers refer to table and item numbers in Chapter 20 of the Handbook.
[b]The nominal HEPs have been modified to reflect the effects of a moderately high stress level and (in some cases) high dependence between two operators.

Figure 4-21. Display of final results in a task-analysis table for operations by an auxiliary operator outside the control room. The column labeled "HRA event tree" does not usually appear in a task analysis; it has been included for the reader's convenience. The numbers in this column refer to the error event numbers appearing in HRA event trees starting with Figure 4-10.

describes some methods for assigning uncertainties in a human-reliability analysis. (A detailed discussion of measures of uncertainty and their propagation is found in Chapter 12.)

4.8.1 SOURCES OF UNCERTAINTY

There are five major sources of uncertainty in estimating the probabilities of human errors in the operation of nuclear power plants:

1. The dearth of data on human performance in nuclear power plants.

2. The inexactness of models of human performance that purport to describe how people act in various situations and conditions.

3. The identification of all relevant performance-shaping factors and their interactions and effects.

4. The skill and knowledge of the human-reliability analyst.

5. The variability in the performance of a given individual and among the performances of different individuals.

The first source, the shortage of human-performance data specific for nuclear power plants, is the most critical. Historically, such data have not been collected on a scale large enough to establish a data base for operations in nuclear power plants. There are, however, some data sources that have been used for human-reliability analysis. The licensee event reports include descriptions of incidents involving human error, but no information on human-error rates or probabilities is given. Furthermore, the determination of what constitutes human error in these reports is frequently questionable.

Although programs to collect data useful for human-reliability analysis are under way, there is at present no single source of data collected from the measurement of human performance in nuclear power plants. Therefore, most estimates of human-error probabilities must involve extrapolation from other sources of information. These sources include (1) the collective judgment of experts (i.e., people with expertise on the performance of the tasks being evaluated) who may directly or indirectly assess error probabilities, (2) the human-performance models and the associated derived data from sources like the Handbook, and (3) data gathered on operationally similar tasks. For example, the actions involved in closing a valve, as specified in a set of procedures, often will be very similar whether the actions are performed in a chemical processing plant or in a nuclear power plant. Such data from similar tasks can be extrapolated or modified to account for dissimilarities in the situations. This extrapolation is subject to error itself, but represents the best approximation available. Many of the estimated human-error probabilities in the Handbook represent this type of extrapolation.

In those cases for which data from operationally similar situations or even derived data are not available, various methods for the use of expert

judgment can be applied. These methods, however, vary greatly in their consistency and validity (Stillwell et al., 1982). (The NRC is sponsoring programs at Sandia and Brookhaven National Laboratories to develop recommended methods and procedures for given nuclear-power-plant applications.) The use of expert judgment as a substitute for actuarial data represents an extreme in the extrapolation process.

The second source of uncertainty is the modeling of human performance. The state of the art of human-reliability analysis is such that the modeling of human behavior can qualitatively account for its variability and for discrepancies in response situations, but there are definite limitations in quantifying such models. There are many models of human performance, but few can be used to estimate the probability of correct or incorrect human performance in applied situations. Furthermore, all models, even those that can be applied to a human-reliability analysis (e.g., the models in the Handbook) are themselves abstractions of real-world circumstances. As such, they only partially represent the situations they simulate. In some cases, experimental data have provided strong support for the general form of the models (e.g., the usual curvilinear form of the performance-under-stress curve), but in others the forms are still speculative (although based on sound psychological concepts).

The third source of uncertainty, the identification of the performance-shaping factors associated with a task, also involves some abstraction and is subject to some interpretation on the part of the analyst. This is probably the biggest source of error in extrapolating data from other sources to the nuclear power plant. Unless the tasks required in both situations are analyzed in sufficient detail, data from other sources may be misapplied to the tasks performed in a nuclear power plant. For example, a valve-restoration task in a chemical processing plant may be superficially similar to an equivalent task in a nuclear power plant, but the HEP from the chemical plant may be based on errors made by people using well-designed checklists, whereas the valve-restoration procedures carried out in the nuclear power plant may be performed from memory only. Using the HEP from the chemical plant to estimate the HEP for the nuclear power plant would obviously result in a gross underestimation of the true HEP.

The above difficulties will be exacerbated if there is little interaction between the human-reliability analyst and other members of the PRA team. Unless the human-reliability analyst is a real working member of the team, his identification of relevant performance-shaping factors and his estimates of the effects of these factors in the human-reliability analysis may ignore important influences of certain plant-specific factors. His estimates of nominal HEP values may be too low or too high. In such cases, the assignment of large uncertainty bounds will not compensate for his lack of knowledge.

The analyst himself is the fourth source of uncertainty; that is, the PRA team may include an HRA analyst who is not fully qualified. He may not be able to perform the necessary extrapolations or to use the human-performance models correctly. The less the PRA team knows about the operations and human activities in a given plant, and the less the team (or at

least the designated person) knows about the underlying psychology, physiology, and sociology of human behavior in general, the less accurate their estimates of human-error probabilities will be. That is obviously a form of uncertainty, but the untutored analyst may not recognize it as such. An independent, qualified observer, however, would want to increase the uncertainty bounds around the estimates made by less qualified analysts. It must be reiterated, however, that merely increasing the uncertainty bounds will not compensate for large errors in estimating the nominal values of the HEPs around which the bounds are placed.

Finally, in the prediction of human behavior, there is an uncertainty that results from the inherent variability of human performance due to individual differences, both within and between the people whose performances are being assessed in the human-reliability analysis. Even if one had a large amount of excellent-quality human-performance data collected for years on all nuclear-power-plants tasks, this variability would contribute to the uncertainty in a human-reliability analysis. A human-reliability analysis does not attempt to estimate the performance of one known person; instead, the analyst's estimates have to account for the fact that any given task may be performed by any one of many individuals, each of whom may vary somewhat in his reliability from day to day or even within a day.

The amount of uncertainty resulting from intra- and inter-individual differences is judged to be considerably less than that resulting from the combination of all the other sources of uncertainty. Some data on individual differences in a wide variety of industrial tasks were collected by Wechsler (1952). These data indicate that for routine and very well defined tasks the ratio of the performance scores of skilled performers near the top of a distribution for some measure of ability to the scores of performers near the bottom of the distribution is about 3:1. In these measures, the upper and lower one-tenth of 1 percent of the distribution was ignored, and thus the 3:1 range ratio includes about 99.9 percent of the scores. In the Handbook, a more conservative range ratio of 4:1 was assigned for individual differences per se, excluding the upper and lower 5 percent of the distribution of HEPs on routine tasks performed by skilled personnel. Thus, it is presumed that the 4:1 range ratio includes the middle 90 percent of the HEPs due to individual differences alone.

In the Reactor Safety Study (USNRC, 1975), to account for the variability in modeling human performance in general and the occurrence of a given error in particular, the Handbook's 4:1 range ratio was increased to 10:1 for most tasks and to 100:1 for tasks whose nature could not be well defined and for tasks performed under conditions that were ill defined or judged to be highly stressful. The Handbook has adopted and refined this concept of larger uncertainty bounds for "more uncertain task behavior." For routine tasks the typical range ratio is 10:1. For tasks involving interpretation or decision-making, a 20:1 ratio is not uncommon (in the revised draft in press), and a high 25:1 range ratio is used for performance under high stress. Each range reflects the uncertainty due to human variability, the lack of representative data, the imprecision of the modeling process, and the identification of relevant performance-shaping factors, but excludes the uncertainty attributable to analysts untrained in HRA techniques.

For applications of the Handbook HEPs and uncertainty bounds to human-reliability analysis, it is assumed, as noted earlier in this chapter, that the PRA team has the necessary expertise not only in HRA techniques but also in the other areas relevant to probabilistic risk assessments.

To summarize, the most significant contributors to uncertainty in the human-reliability analysis of nuclear-power-plant operations can be ranked by importance. Assuming the necessary analytical skills, the lack of data from actual human performance in nuclear power plants is the most important contributor. Naturally, if we had sufficient data on human-error probabilities for each task being analyzed, it would not be necessary to model each task. The second most important contributor to uncertainty is the inexactness of the models. No abstraction can fully define or account for all the variables in response situations as complex as those found in a nuclear power plant. Furthermore, it is unrealistic to suppose that each model will be applied consistently across all analyses. This lack of consistency is related to the difficulties in performing the necessary analyses of human inputs, mediating processes, and responses so that the relevant performance-shaping factors can be identified and assessed correctly (the third most important contributor to uncertainty). The fourth most substantial contributor to uncertainty is the variability of human performance. The uncertainty bounds associated with the estimates of human-error probability are almost certainly very conservative in accounting for the range of possible human performance on the various tasks modeled by various human-reliability analysts.

4.8.2 METHODS FOR HANDLING UNCERTAINTIES IN A HUMAN-RELIABILITY ANALYSIS

A human-reliability analysis consists of combining, in some fashion, HEPs for many different tasks or activities. For some PRA purposes, the use of uncertainty bounds may not be necessary. Instead, it may be sufficient to use single-point estimates as illustrated earlier in this chapter. When it is necessary to assign uncertainty bounds, there are two general approaches that have been used. The first is to propagate uncertainty bounds throughout the HRA portions of the PRA, using the methods discussed in Chapter 12. The second approach is to proceed with the usual propagation of point estimates through the HRA portion and then to assign uncertainty bounds about the final point estimate (i.e., the total human-error term for each portion of the human-reliability analysis). These methods can result in uncertainty bounds that are quite different, and it is up to the PRA team to select and justify the method it employs.

With regard to the first approach, the propagation of uncertainty bounds for each HEP, a commonly accepted method is that of using a Monte Carlo procedure to sample values from the distribution of each error probability in the analysis. Generally, in applying a Monte Carlo procedure, random sampling from each distribution in the analysis is used. In actual fact this procedure will not reflect the true response situation in that a dependence over tasks could exist. If an operator's skill level is fairly constant with respect to those of other operators for any of the tasks he

undertakes, his error probabilities are likely to fall close to the same relative position on each of the distributions being analyzed. Therefore, if the same operator performs each of the tasks being analyzed, there is very little likelihood that his performance will correspond to a set of randomly sampled HEP. To avoid this problem, one could set up a sampling procedure to reflect the above or other sources of dependence.

An alternative is the discrete probability distribution (DPD) method, also discussed in Chapter 12, in which the distribution of each HEP is graphed as a discrete histogram. In essence this method represents each continuous distribution with some finite number of points. To evaluate the uncertainty associated with combinations of human actions and other events, histographs representing the distributions of each can be combined to derive an uncertainty distribution associated with the combined failure probabilities of interest. The above-stated cautions about sources of dependence also apply to the DPD method.

If the robustness of a Monte Carlo or a DPD procedure is deemed unnecessary or inappropriate in view of the lack of actual data on human-error distributions in the performance of nuclear-power-plant tasks, the second approach to the treatment of uncertainties can be used. This approach avoids the necessity of propagating uncertainty bounds through the HRA portion of the PRA. Instead, uncertainty bounds are assigned to the total human-error probability obtained from each HRA portion of the PRA. For example, one would assign uncertainty bounds to the total error probability obtained from an HRA event tree like the one shown in Figure 4-18. In the remainder of this discussion on uncertainties, the HRA-event-tree method from the Technique for Human Error Rate Prediction is used to explain some methods used in this second approach to the treatment of uncertainties. However, the discussion pertains to any other HRA method as well.

In discussing the second approach, it is useful to define some terms. An HEP and uncertainty bounds are given in the form of

$$\text{HEP}\left(\frac{1}{k_1} \times \text{HEP},\ k_2 \times \text{HEP}\right)$$

where the first term in parentheses is the lower bound and the second term in parentheses is the upper bound. For example, as in the tables from the Handbook, if the estimates are

$$.005\ (.001 \text{ to } .05)$$

then

$$\text{HEP} = .005, \quad k_1 = 5, \quad k_2 = 10$$

If $k_1 = k_2$, the bounds are said to be symmetrical and the "error factor" is used to denote both k values. The uncertainty range (UR) for asymmetrical uncertainty bounds is $UR = k_1 k_2$, and for symmetrical bounds it is the square of the error factor.

Discussed briefly below are three methods, or approximations, that involve the usual propagation of point estimates through the HRA event tree, with the assignment of uncertainty bounds about the final point estimate (i.e., the total failure term for the tree). The output--that is, the final failure term and the associated uncertainty bounds--is then entered into the appropriate places in the system event or fault trees. What the point estimate represents (for instance, whether it is the mean or the median of some distributions) depends on the analyst's interpretation and understanding. However, if point estimates are taken from the Handbook, the usual practice is to consider them as medians of a lognormal distribution.

The simplest of the three methods is to assign some arbitrary set of uncertainty bounds to the total failure probability obtained from the HRA event tree. In some PRAs, once this total failure probability was determined as a point estimate, uncertainty bounds of a factor of 10 on each side of the point estimate were assigned. It is important to note that this error factor of 10 is considerably larger than the typical error factors for the individual HEPs that were used to calculate the total failure probability. For a lengthy and interactive HRA event tree, especially one that represents the performance of more than one person, some analysts might judge that an error factor of 10 is not sufficiently conservative.

Another method for assigning uncertainty bounds to the total failure term of an HRA event tree is to take the largest error factor (the square root of the uncertainty range about an HEP) found for any HEP in the tree and to apply it as the error factor for that total failure term. This method should be employed only where the distribution of the uncertainty bounds about the total failure probability is to be symmetrical.

The third method, a variant of the second, does not require symmetrical uncertainty bounds. The largest uncertainty range about an HEP is used as the uncertainty range for the resulting probability of total failure in the human-reliability analysis.

In following either the second or the third method, we say that the uncertainty associated with the entire analysis is no greater than that associated with the most uncertain element of the analysis. In some cases, this assumption may not be sufficiently conservative.

Some of these methods have been documented, as they were used in PRAs that have already been completed. In view of the different viewpoints as to how uncertainties should be propagated in a PRA, no recommendation can be made here as to the best method for assigning uncertainty bounds in the human-reliability analysis per se. Furthermore, because most uncertainty bounds for individual HEPs are not determined from data collected in nuclear power plants, the method employed may not be very critical in a PRA so long as the uncertainty bounds for terms entered into the system analysis are not unrealistically narrow. It is apparent that a sensitivity analysis can be very useful to ascertain the impact on the system analysis of assuming different uncertainty bounds for the human-error terms to be incorporated into the system event or fault trees.

4.9 ALTERNATIVE METHODS OF HUMAN-RELIABILITY ANALYSIS

While other methods for estimating the human-error contribution to system reliability have been developed and documented, it is important that the reader keep in mind the state of the art of human-reliability analysis in considering them for use in a probabilistic risk assessment. Several of the newer methods were developed specifically for use in PRAs, while others are the result of modifications made to models of human performance that were initially developed for quite different purposes. Some human-performance models can be used to estimate the likelihood of human errors, but many of them may not be useful for a PRA in that they cannot be applied to all situations modeled in a risk assessment. Some models that have been documented are very limited in scope; they model human performance at a level so detailed that it cannot be realistically observed and thus cannot be verified. Other models deal with human performance in contexts that are largely covered by other portions of the PRA. For example, human errors made in conducting maintenance operations (rather than in restoring equipment after such operations) will usually be detected in the equipment-failure rates. The inclusion of such errors in the system models constitutes a double accounting: the impact of human errors made in maintaining equipment will be incorporated into the system fault trees twice. Some of the alternative methods simply represent restatements or reorganizations of the material in the Handbook or other sources and should be used if their presentation formats fit in better with the overall scheme of a particular PRA. Extreme care should be taken in employing these or any HRA methods since the potential for error in using them is high given the context of the PRA.

4.9.1 HUMAN-RELIABILITY ANALYSIS IN THE OCONEE PRA

In the human-reliability analysis performed for the Oconee PRA, human errors were classified into two types, latent and dynamic (Dougherty, 1981). Latent errors are made by maintainers or operators who fail to restore components or systems to their proper states after testing, maintenance, or calibration. These errors result in component or system unavailabilities and occur before a transient (during which, it is assumed, the component or system would be required). Dynamic errors are made by operators during the course of an accident. The circumstances under which any error is made are usually of less interest than are the system effects of that error. In other words, whether a valve is unavailable because of an error in restoration after testing or because an operator locked it while responding to a transient is irrelevant in terms of the system effects, which are that the valve is unavailable. The causes of the unavailability are important to the estimation of the probability of the underlying error, but not to the estimation of the system effects of the error itself. The distinction between latent and dynamic errors is, however, supported by the different classes of recovery factors that apply to each case. Also, this classification fits in well with the scheme of the Oconee study for incorporating the results of the human-reliability analysis into the entire PRA, as discussed below.

In the Oconee PRA, estimates of human errors were incorporated at three levels (Dougherty, 1982):

1. Above the system level (in the system event trees or in the logic connecting the system fault trees to the system event trees).

2. At the system level of the system fault trees.

3. At the component level of the system fault trees.

At the first level, the Oconee PRA took into account the effects of several factors that have the potential for affecting the probability of human error in responding to a transient. These include the operator's perception of the severity of the situation, the timing of the accident sequence, the amount and the quality of direct indications of plant status in the control room, the success options available to the operator, and the training and/or procedures available to the operator that would support his successful completion of the proper response to the transient.

The general criteria for estimating the probabilities of human errors and the effects on these probabilities of the above-mentioned factors were obtained from the Handbook (NUREG/CR-1278) an the subjective judgment of the HRA team for the Oconee study. A Delphi method was used to solicit estimates of the basic human-error probabilities and the relevant factors. The group sampled included members of the HRA team and former plant operators. The HRA team was interested in obtaining order-of-magnitude best estimates of human-error probabilities.

The human errors that were included in the first level of incorporation were grouped according to four general types (Dougherty, 1982):

1. Situations where the actions of the operator represent an immediate redundancy to system performance.

2. Situations where the operator acts to find alternative success paths.

3. High-stress situations where the operator has little time to succeed or must leave the control room to succeed.

4. Low-stress situations where the operator has long times to succeed but must make significant repairs to plant systems.

At the second level of incorporation, the system level, the estimates of human-error probabilities were input at the top of the system fault trees. At this level, human errors that could affect the availability of an entire system were considered. For example, if an operator misdiagnoses an accident, he can disable an entire system required to respond correctly to the accident. The probabilities of these misdiagnoses were determined by using a "confusion matrix" developed for the Oconee study. This matrix is the result of interviews with PWR operators who estimated the likelihood that different initiators would be mistaken for each other. The time available to the operator for making a diagnosis--that is, the interval between

the initiation of the accident and the time at which system reliability would be degraded--was taken into account in estimating these errors. Errors in calibrating safety systems that could result in out-of-tolerance system performance were also included at this level.

At the third level of incorporation, the component level, three types of errors were identified: errors made in restoring items of equipment after testing, maintenance, or calibration; violations of technical specifications in concurrently performing maintenance or testing on parallel systems, thus rendering them unavailable; and procedure-based errors in which the operator, in trying to respond successfully to an accident or a transient, causes the unavailability of some component. The probability of concurrent maintenance was judged by the Oconee HRA team to be negligible because the plant has a very good administrative-control system. These errors were not included in the analysis. Neither was the last type of error defined at this third level of incorporation--the errors made by the operator in attempting to follow the correct set of procedures in responding to an accident--included at this point in the analysis. The Oconee HRA team judged that several different operator errors at this point would result in the same system effects, and these errors were therefore grouped with others for inclusion at a higher level in the system models.

4.9.2 THE OPERATOR-ACTION TREE

The operator-action tree (OAT) has been used in the PRA for the Susquehanna nuclear plant. In general, it involves a higher-level human-reliability analysis than that described in the Handbook because the OAT format provides for the incorporation of the HRA results at the system-event-tree level and because, in modeling the response to a transient, it emphasizes the importance of units of team performance over those of the individual. (This level of incorporation of the human-reliability analysis into the PRA can conceivably be accomplished with the results of a Handbook human-reliability analysis, but the Handbook method is not specifically designed for this level of incorporation.)

The OAT method uses a horizontal event-tree format to model the probability of occurrence of the initiating event and the following human behaviors: monitoring indicators, interpreting the problem correctly, and taking timely correct action (Wreathall, 1981). Monitoring indicators involves the operators' taking notice of any displays that give information as to the type of event that has occurred. Interpreting the problem correctly calls for the operators' correctly assessing the state of the reactor from the available displays. This ability is very strongly influenced by the amount and the type of training the operators have received and by their familiarity with that particular event. Taking timely correct action depends almost entirely on the operators' correct interpretation of the event. It involves their correcting errors made in preparing the plant for the proper automatic response and taking appropriate steps to mitigate the effects of the event. (It is possible that this step could be performed correctly (at least for a time) when an incorrect interpretation was made. This might happen if the operators mistook for the true initiating event an event with similar response requirements. It is assumed that correct

response while reacting to an incorrect model of plant status would not be possible for the entire course of the accident sequence.)

Data for the monitoring activities, for taking correct action, and for taking recovery action can be obtained from the Handbook or from a similar source of human-performance data. Data for the correct interpretation of plant status can be derived from the OAT time-reliability curve (Wreathall, 1982).

Since the time available for making a correct diagnosis and correctly responding is the major variable affecting performance, it is the factor used to characterize the operators' response behavior. The time-reliability curve plots the probability of failure against the time available for the operator to make a correct diagnosis. The available time is defined as the interval between the initiation of the accident and the time at which response activities would come too late to avoid undesirable system consequences. The curve ignores the first few minutes after a transient as involving behavior that is too uncertain to model. It deals with team behavior; that is, it plots the probability of the entire control-room team's failing to diagnose the event correctly. This allows implicit consideration of the types of team interaction considered in some of the Handbook's models, such as the dependence model.

The data points for the time-reliability curve are obtained from the expertise of the analysis team. The members of the analysis team use their familiarity with the specific plant being analyzed and their knowledge of the principles of human behavior to estimate the probability of the operating team's performance in diagnosing transients correctly. In the Susquehanna study, the analysis team included persons with expertise in engineering psychology, systems engineering, and nuclear plant operations.

To account for the uncertainty in the data-gathering process and for the variability of human performance, the time-reliability curve is characterized by an uncertainty range consisting of an order-of-magnitude spread on either side of the best-estimate predictions. This uncertainty range is not meant to imply statistical confidence limits, but only to reflect the predicted middle 80 percent of the actual performance distribution for the operating team. This uncertainty range is also used to accommodate the effects of "reluctance" factors, which are similar in effect to the performance-shaping factors described in the Handbook. For example, if an operator is required by the plant condition to take an action he would normally avoid because of his training, he is less likely to perceive the requirement for this action in comparison with an action that is in agreement with his training. In this case, the probability of a failure in diagnosis at any given point in time on the OAT time-reliability curve would be increased by some factor, usually 2 to 5.

In incorporating the results of the analysis into the system fault trees, the OAT method accounts for dependence among events by assigning dependent events the same fault designator. Thus, when unrelated components are affected by behaviorally related activities, these activities are linked by giving them the same label in the fault tree. That fault-tree event will appear as the developed set of potential human errors. In this way, the dependence can be included in the fault tree for any component.

4.9.3 ACCIDENT INITIATION AND PROGRESSION ANALYSIS

In the accident initiation and progression analysis (AIPA) performed for a high-temperature gas-cooled reactor (HTGR), an operator-response model was developed to "provide a consistent basis for evaluating both the time and likelihood of a proper operator response for the accident sequence under consideration" (Fleming et al., 1978). The model is essentially an input/output model for the operators of the HTGR, with the inputs being any incoming information presented to the operators, such as alarms or other signals, and the outputs being the set of possible operator responses.

These possible operator responses were grouped into two categories: mitigating activities and nonmitigating activities. In general, mitigating activities involve an operator's responding to abnormal plant conditions by reducing power or initiating plant shutdown. Nonmitigating activities involve an operator's responding to abnormal plant conditions by taking inappropriate action or by taking no action, either of which would degrade system reliability (Raabe et al., 1977). Human-factors methods were developed during the AIPA study to treat both the beneficial and the detrimental actions of operators and maintenance crews (Hannaman, 1981).

The characteristics of an HTGR are such that extremely rapid responses on the part of the operators are rarely, if ever, required. Under most abnormal plant conditions, the operators are allowed sufficient time to make and reevaluate decisions about the nature of the occurrence, which makes it likely that they will take at least some corrective action. Because of this, in the first phase of the study, the effect of the operators' taking inappropriate or uncorrected action was modeled as taking no action to simplify the analysis. In the second phase, inappropriate actions or errors of commission were incorporated on a case-by-case basis.

The AIPA approach to modeling the impact of human errors consisted of several activities. Event trees and fault trees were used to define the explicit human interactions that could change the course of a given accident sequence and to define the time allowed for corrective action in that sequence. A time-dependent operator response model was developed that related the time available for correct or corrective action in an accident sequence to the probability of successful operator action. A time-dependent repair model was developed to account for the likelihood of recovery actions for a sequence, with these recovery actions being highly dependent on the system-failure modes. Data on human-error contributions were collected for each event and included in the fault or event trees both as common-mode fractions and as random system or component failure rates (Hannaman, 1981; Fleming et al., 1979).

In operating, testing, and maintaining equipment, human errors that cause component or system failures are treated explicitly in the system fault-tree analyses and implicitly in the method used to model the reliability characteristics of dependent failures in redundant systems (Fleming et al., 1978). The implicit treatment arises from the use of failure-rate and dependent-failure experience data that include contributions from human errors (Hannaman and Kelley, 1978).

The bases for the operator model are as follows:

1. Initially there is a probability of zero that the operator will respond instantaneously.

2. As time increases, the probability that an operator will take corrective actions increases.

3. If the operator discovers that his initial actions are insufficient for plant recovery, he will take further action until a stable condition is reached.

These factors indicate an increasing probability of operator success in time. The probability of success in this model increases until a time t_{max} is reached. The parameter t_{max} is the time available for operator action, determined from computer models that simulate the physical behavior of the system for the postulated accident and the transient response of key components. In a particular accident, the time available for operator action is determined by the transient thermal and structural response of the reactor core, vessel, structures, and containment. Usually a limiting component temperature or pressure defines the time available for operator action.

The likelihood that the operator will be able to take action to mitigate the consequences of an initiating event increases as the time available for such action increases. The time available to take such action is the time until the point at which such action will no longer significantly change the consequences of the event. The time within which 63 percent of trained operators will take successful action is the mean time to operator response (MTOR), the expected response time for an average, adequately trained operator. Data on the MTOR can be "obtained from measurement of operator response, estimates of knowledgeable experts, or development of a functional relationship for the most important variables contributing to the response time in the reactor control room environment" (Fleming et al., 1975). In the AIPA study, expert judgment was used to estimate MTOR, which was assumed to have a lognormal distribution. Confidence limits on the MTOR were determined by computing the standard deviation or by plotting the estimates and determining the variability graphically. To account for the effects of stress on operator performance, the estimates of MTOR were increased by 10 to 20 percent, in effect reducing the probability of correct operator action for a given time under stressful conditions.

The AIPA operator-response model is intended for HTGR conditions. For other situations, the probability distributions on time, MTOR, and their functional relationships should be investigated before applying this model (i.e., for short or long t_{max} other models may be useful).

The steps taken in applying the operator-response model were as follows (Fleming et al., 1975):

1. Determine the need for operator action in a branch-point fault tree.

2. Identify the operator's situation.

   a. Identify instrumentation that is operating, failed, etc., which may be dependent on the particular branch conditions.

   b. Identify the expected or trained-operator response, which may come from technical specifications or planned operator procedures (training).

3. Obtain data and analyze operator response.

   a. Utilize data sources (i.e., the Reactor Safety Study, abnormal occurrence reports, or expert opinion).

   b. Adjust data to include stress factors.

   c. Use the data range to determine uncertainty in the MTOR.

   d. Consider the interrelation of multiple operator actions within the same fault tree, which may require the use of a common-mode beta factor.

   e. Determine an upper limit (P_s), which is generally in the range of .99 to .9999.

4. Treat the resulting probability (P_{of}) and uncertainties as equipment-failure blocks in the fault-tree diagram (which may include the use of the sample computer code to determine the overall fault-tree uncertainty).

5. Use the time factor to help determine the range of consequences resulting from the two branches.

Although the consideration of human factors in the AIPA study was balanced between beneficial and detrimental actions in line with the objective of making realistic risk estimates, certain elements of the treatment may be viewed as conservative and still others as optimistic. Among the former are the use of maintenance data to quantify the timing of operator actions during accident situations and the omission from consideration of (1) human ingenuity to terminate the accident and (2) the mobilization of experts and technicians to supervise long-term external actions to mitigate the accident consequences. The most important class of actions whose omission can lead to underestimates of accident risk appears to be errors of commission that either initiate accidents or compound their consequences and those that cause the failure of multiple, otherwise independent, systems.

4.9.4 CONCLUSIONS

The methods outlined above have been applied in actual PRAs. There are, in fact, several other methods and models of human-reliability analysis in existence, but most of them have seen limited application or no application in PRAs as yet. The state of the art of human-reliability analysis is

changing rapidly at present. New methods are being developed, and older models are being revised and updated to accommodate the type of information needed for a PRA. The users of this guide are urged to investigate recent developments in human-reliability analysis that are or will shortly be available in the public literature. Limitations to these models should be observed carefully, and professionals with experience in human-performance techniques should be responsible for their use.

Of especial interest in current months are examples of "cognitive models," developed to provide estimates of errors made in diagnosing particular accident signatures and in deciding on corrective action. These are highly speculative and should be investigated with caution before application in a PRA. However, for such errors screening models are available, and they can be used more readily because of the extremely wide uncertainty bounds associated with them.

4.10 ASSURANCE OF TECHNICAL QUALITY

To ensure that the quality of any given human-reliability analysis is maintained and that the quality of the several analyses is constant, a program plan for the performance of these analyses should be developed. This plan should be developed by the director of the human-reliability analysis in conjunction with the PRA team leader.

To meet internal quality standards (those relating to any given human-reliability analysis), the plan should provide for scheduling the various stages of the analysis, integrating it into the entire PRA, and monitoring its progress. To this end, dates, places, personnel, and expected results should be identified. Working from the block diagram in Figure 4-2, for example, tables or charts should be set up itemizing each task; the elements necessary for its completion (including personnel); its relation to and/or interfaces with other PRA groups; the date, time, and place of its expected performance; the expected results; and the method of its documentation.

To meet external quality standards (those relating to human-reliability analyses performed for several plants), the plan should provide for cross-plant comparisons. This implies that the team leader for a new PRA should be familiar with the HRA program plan implemented in earlier PRAs, using this information to ensure that the control and documentation of the ongoing analysis are complete.

REFERENCES

Bell, B. J., and A. D. Swain, 1981. A Procedure for Conducting a Human Reliability Analysis for Nuclear Power Plants, NUREG/CR-2254, draft USNRC report for interim use and comments.

Dougherty, E. M., 1981. "The Human Element in a Probabilistic Risk Assessment," in Proceedings of the Myrtle Beach Workshop on Human Factors and Nuclear Safety, August 1981.

Dougherty, E. M., 1982. "Treating Human Interactions in Risk Assessment," in Proceedings of the ANS Topical Conference on PRA, April 1982, American Nuclear Society, Inc., La Grange Park, Ill.

Embrey, D. E., 1976. Human Reliability in Complex Systems: An Overview, NCSR.R10, National Centre of Systems Reliability, United Kingdom Atomic Energy Authority, Warrington, England.

Embrey, D. E., 1981. "The Use of Quantified Expert Judgment in the Assessment of Human Reliability in Nuclear Power Plant Operation," in Proceedings of the Human Factors Society 25th Annual Meeting, Human Factors Society, Santa Monica, Calif.

Fleming, K. N., et al., 1975. HTGR Accident Initiation and Progression Analysis Status Report, Vol. II, "AIPA Risk Assessment Methodology," U.S. Energy Research and Development Administration, Washington, D.C.

Fleming, K. N., et al., 1978. HTGR Accident Initiation and Progression Analysis Status Report: Phase II Risk Assessment, U.S. Department of Energy, Washington, D.C.

Fleming, K. N., F. A. Silady, and G. W. Hannaman, 1979. "Treatment of Operator Actions in the HTGR Risk Assessment Study," Transactions of the American Nuclear Society, Vol. 33

Hannaman, B., 1981. "Human Factor Considerations in the Accident Initiation and Progression Analysis," in Proceedings of the Myrtle Beach Workshop on Human Factors and Nuclear Safety, August 1981.

Hannaman, G. W., and A. P. Kelley, 1978. "Synthesis of Experience Data for Risk Assessment and Design Improvement of Gas-Cooled Reactors," in Proceedings of the ANS Topical Meeting on Probabilistic Safety, Los Angeles, May 8-10, 1978, American Nuclear Society, Inc., La Grange Park, Ill.

Meister, D., 1971. Comparative Analysis of Human Reliability Models, L0074-107, Bunker-Ramo Electronics Systems Division, Westlake Village, Calif.

Pew, R. W., S. Baron, C. E. Feehrer, and D. C. Miller, 1977. Critical Review and Analysis of Performance Models Applicable to Man-Machine Systems Evaluation, AFOSR-TR-77-0520, U.S. Air Force Office of Scientific Research, Bolling Air Force Base, Washington, D.C.

Raabe, P. H., et al., 1977. HTGR Accident Initiation and Progression Analysis Status Report, Vol. VIII, "Responses to Comments on AIPA Status Report," U.S. Energy Research and Development Administration, Washington, D.C.

Stillwell, W. G., D. A. Seaver, and J. P. Schwartz, 1982. Expert Estimation of Human Error Probabilities in Nuclear Power Plant Operations: A Review of Probability Assessment and Scaling, USNRC Report NUREG/CR-2255 (in press).

Swain, A. D., and H. E. Guttmann, 1980. Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications, NUREG/CR-1278, draft USNRC report for interim use and comment.

U.S. Department of Defense, 1981. Military Standard, Human Engineering Design Criteria for Military Systems, Equipment and Facilities, MIL STD-1472C, Washington, D.C.

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

USNRC (U.S. Nuclear Regulatory Commission), 1981a. Evaluation Criteria for Detailed Control Room Design Review, NUREG-0801.

USNRC (U.S. Nuclear Regulatory Commission), 1981b. Guidelines for Control Room Reviews, NUREG-0700, draft USNRC report for interim use and comment.

Wechsler, D., 1952. Range of Human Capacities, Williams & Wilkins, Baltimore, Md.

Wreathall, J., 1981. "Operator Reliability Model," in Proceedings of the Myrtle Beach Workshop on Human Factors and Nuclear Safety, August 1981.

Wreathall, J., 1982. Operator Action Trees--An Approach to Quantifying Operator Error Probability During Accident Sequences, NUS-4159, NUS Corporation, Gaithersburg, Md.

Chapter 5

Data-Base Development

5.1 INTRODUCTION

Two types of events identified during accident-sequence definition and system modeling must be quantified for the event and fault trees in order to estimate frequencies of occurrence for accident sequences: (1) initiating events (see Section 3.4.2) and (2) component failures, or primary events (see Section 3.5.3.1). This chapter describes how this quantification is performed.*

The quantification of initiating and primary events involves two separate activities. First the reliability model for each event must be established, and then the parameters of the model must be estimated. The quantification also involves various types of data analysis (e.g., a statistical analysis of raw information), the use of generic and specific data, and, in some cases, the collection and use of subjective data. The necessary data include component-failure rates, repair times, test frequencies and test downtimes, common-cause probabilities, and uncertainty characterizations. Also involved is the quantification of human errors, a subject not covered here because it is discussed in Chapter 4.

The objective of the task described in this chapter is to estimate the frequencies of the initiating events and the probability of the primary events identified in accident-sequence definition and system modeling (Chapter 3) and thus to develop a data base for accident-sequence quantification (Chapter 6). It is important to note that the output of this task must be consistent with the general approach chosen and the tools to be used in accident-sequence quantification. Before this task is performed, a decision will have been made as to whether the PRA will use a classical or a Bayesian framework for treating uncertainties. This decision will affect the way data are evaluated. In addition, the tools used in sequence quantification will also affect the data analysis, in that the data must be in a form compatible with the tools. For example, the data analysis may yield probability distributions for reliability models that cannot be exactly represented by any defined distribution (e.g., a gamma or a lognormal distribution), and yet the quantification tools require that all inputs be described by one of a set of predefined distributions. It will be the data analyst's job to make the data output fit this quantification requirement, by finding the "best" distribution to fit the actual result, and then to record any uncertainty (Chapter 12) that is thus introduced in the analysis. Hence, the task described in this chapter is closely linked with the tasks of Chapters 3, 6, and 12.

---

*The numerical quantities obtained by the procedures of this chapter are in a very strict sense estimates; that is, these quantities should be considered judgments of the values for the numerical quantities of interest.

5.2 OVERVIEW

The development of a data base for accident-sequence quantification is a multistep process involving the collection of data, the analysis of data, and the evaluation of appropriate reliability models. It produces tables that specify the quantity to be used for each event in the fault and event trees.

While the task of data-base development may seem to lie between the tasks of accident-sequence development and quantification (Chapters 3 and 6), it is most likely to be accomplished largely in parallel with accident-sequence development.

The steps that need to be addressed in developing a data base are outlined below, in the order the tasks would be accomplished. As in many engineering analyses, the order may be modified as the work progresses, or iteration may be required. It is also possible that time constraints, budget constraints, or study goals may allow, or even require, some steps to be shortened or bypassed. For example, instead of collecting and analyzing raw data, it may be sufficient to use data from a previous PRA study. This could save considerable time and cost, but it may diminish confidence in the results. Figure 5-1 indicates the flow of the steps outlined below.

Selection and Use of Event Models. The data analyst must select several types of models for event quantification: failure models, maintenance models, test models, and initiating-event models. The factors to be considered in these decisions are discussed in Section 5.3.

Data Gathering. Early in the PRA project, the gathering of all information that may be pertinent to events usually included in PRA studies should begin. At this point the development of accident sequences will not have been completed, and hence this early information gathering must rely on previous experience. The information should include published data reports, data from other PRA studies, and available information about the specific plant that is being analyzed. This task is described in Section 5.4.

Estimation of Model Parameters. After the models have been selected, their parameters must be evaluated. Two approaches to parameter estimation, the Bayesian approach and the classical approach, are described in Section 5.5.

Evaluation of Dependent Failures. It is generally recognized that dependent failures may make significant contributions to system unreliability. Section 5.6 addresses various methods available for estimating these contributions.

Uncertainties in Data. A major concern in a probabilistic risk assessment is the issue of uncertainty in the various evaluations. Section 5.7 discusses the factors in data-base development that contribute to uncertainty.

Documentation. The results and the process of data-base development must be documented. Guidelines for documenting the data base in a clear and consistent manner are presented in Section 5.8.

Assurance of Technical Quality. It is very important that the resultant data base be as accurate and as consistent as possible. Procedures for ensuring that the data base is of the best possible quality are presented in Section 5.9.

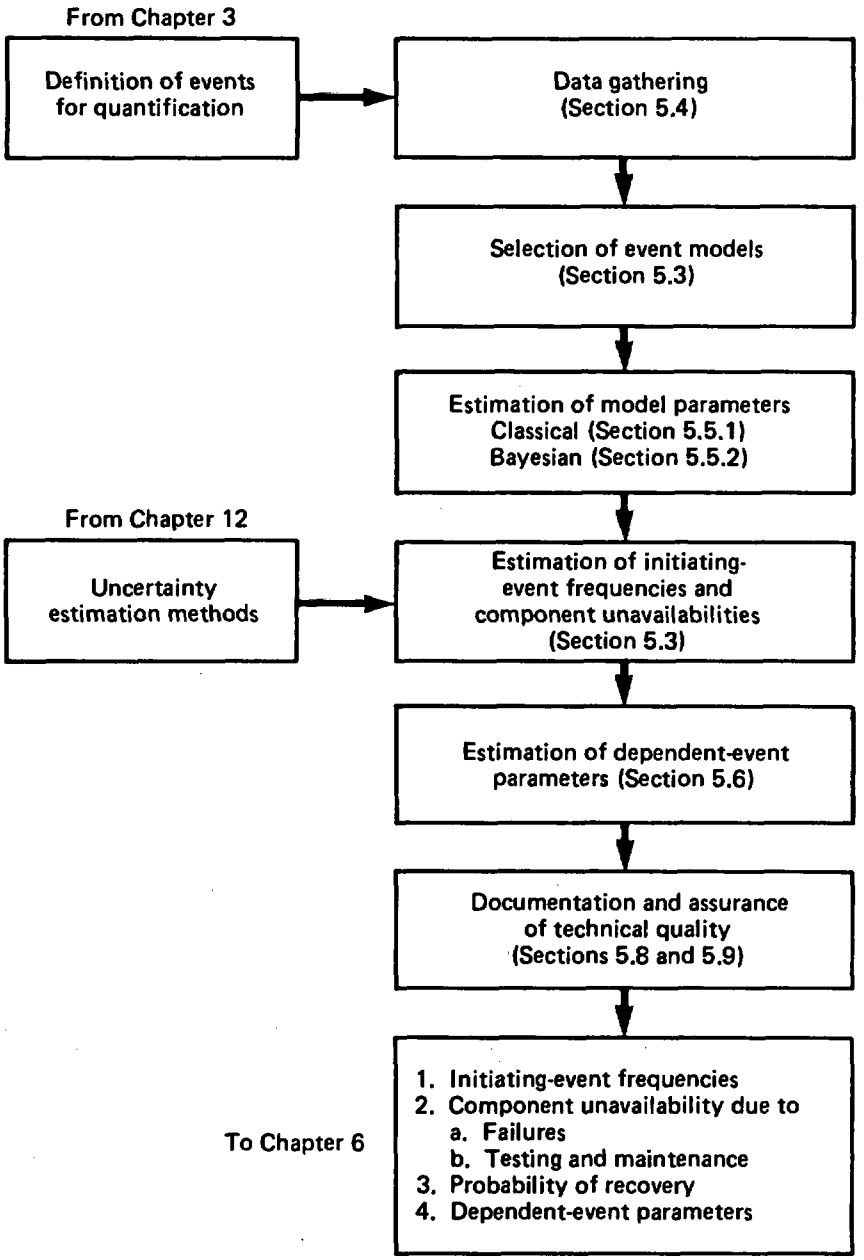

Figure 5-1. Inputs, outputs, and steps in data-base development.

5.3 EVENT MODELS AND THEIR USE

The primary events in the fault trees and event trees can be analyzed with four types of models: component-failure models, test-contribution models, maintenance-contribution models, and initiating-event models. The first three of these models provide estimates of the probability that a plant element cannot accomplish its design function because it has failed, is being tested, or is being maintained. The model for initiating events provides the estimated frequency of the specific event of interest.

5.3.1 COMPONENT-FAILURE MODELS

Component-failure models can be divided into two general types: time-related models and demand models. This section defines both types of models and explains their application.

5.3.1.1 Time-Related Models

5.3.1.1.1 Definition

Reliability as a function of time can be modeled by a number of probability distributions, the more common models being the exponential, the Weibull, the gamma, and the lognormal. Each represents a different type of failure process.

The exponential gives the distribution of time between independent events occurring at a constant rate. The Weibull gives the distribution of time between independent events occurring at a rate that varies in time. The gamma gives the distribution of time required for exactly k independent events to occur, assuming a constant rate of occurrence. An exponential distribution is a gamma with $k = 1$. The lognormal implies that the logarithms of lifetimes are normally distributed. There are also other models that provide for time-dependent failure rates, an example being the inverse Gaussian (Chhikara and Folks, 1977).

In most PRA studies, the exponential is the most commonly used time-to-failure distribution. It is used basically for two reasons: (1) many reliability studies have found the exponential justifiable on empirical grounds and (2) both the theory and the required calculations are simple. It is important to note that, even though the time to failure is not exponential over the entire life of the component, the in-use portion may be exponential. This assumes replacement by a component that is also in its exponential-behavior time period.

The validity of the assumptions underlying the choice of the exponential distribution can be examined by several methods. These methods are not discussed here because most PRAs have not found it necessary to justify their choices of reliability models. Should there be a need to examine the time-to-occurrence distribution, the graphical methods described by Hahn and

Shapiro (1967) and the analytical methods described by Mann et al. (1974) can be used.

In this chapter the exponential distribution will be used to model the time to component failure. The equation for the exponential distribution is

$$U(t) = 1 - e^{-\lambda t} \qquad (5\text{-}1)$$

which represents the cumulative probability that the event has occurred by time t. The parameter λ is the failure rate and is expressed in units of failures per unit time.

5.3.1.1.2 Use of Time-Related Models

Failure in Time: Standby

Many components in a nuclear plant are in a standby mode; that is, they are not used until needed or tested. Often such components are assumed to fail in time while in this standby mode.

Standby components are usually subjected to periodic testing, which occurs, for example, once a month or perhaps once a year. The time between tests is the length of time the component is exposed to failure without detection, and hence the term "fault-exposure time." This time is often designated by τ. The fault-exposure time τ is usually determined from plant procedures, but some caution should be used when examining a system for test intervals. As an example, consider the system in Figure 5-2. This system is tested in various pieces; that is, the logic is tested once a month, as are the spray pumps.

The sensors are calibrated once a year and are tested once a year through the logic. However, the entire system is never tested end to end. This results, in this example, in a specific contact never being tested during the life of the plant. Figure 5-3 focuses on this situation.

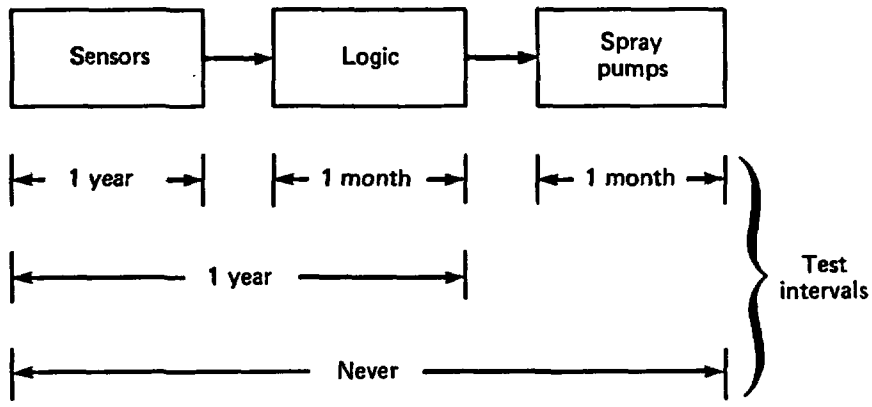

Figure 5-2. Test intervals for sample system.

The logic testing verifies that the coil is energized when the test contact closes and the light is illuminated. However, the contact for pump start is not tested. The analyst then must decide on a value of τ for this contact that is not directly tested during the life of the plant. Indeed, it may be deemed appropriate to assign a τ of 40 years. However, in this case a 40-year value for τ is inappropriate, because the contact is part of a relay that is tested in part and has an associated mean time to failure; thus, the relay will be periodically replaced and the untested contact will be renewed. It is therefore suggested that the τ for the untested element be the reciprocal of the mean time to failure of the tested elements in the relay combined through an OR operation.

In the present example, assume that the coil has a mean time to failure of 20 years and the tested contact has a mean time to failure of 5 years. These can be combined by adding the failure rate, defined to be the reciprocal of the mean time to failure, and then inverting the result; that is, $\tau = \left[(1/20) + (1/5)\right]^{-1} = 4$ years. Thus, it would be appropriate to use $\tau = 4$ years for the contact that is not directly tested.

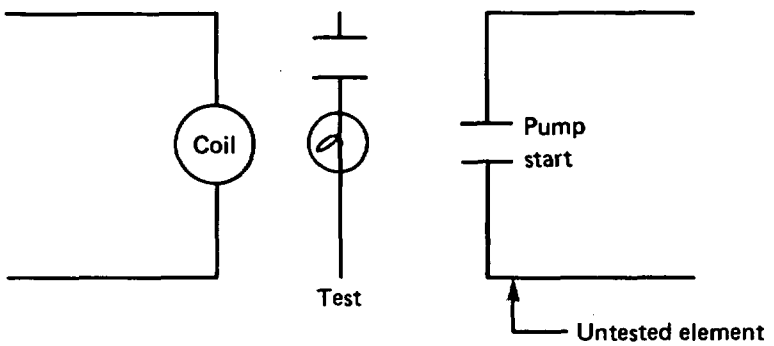

Figure 5-3. Interface schematic.

After determining an appropriate τ for each component that is modeled to fail in time during standby, it is necessary to define the unavailability due to each component's random-failure distribution in time. The expression for the availability of a component that fails in time over a period τ is given by the cumulative distribution function of the time-to-failure distribution for that component. For example, if a component is found to have an exponential failure density function (i.e., $f(t) = \lambda e^{-\lambda t}$), then the unavailability is given by

$$U(t) = 1 - e^{-\lambda t}$$

However, the demand on the safety systems and components occurs randomly in time. Thus, it is necessary to evaluate the unavailability function during the fault-exposure time τ. If it is assumed that the demand can occur with equal likelihood at any point in the τ interval, as it usually does, the

unavailability that should be used is the frequency-weighted unavailability* over the time period τ. Thus,

$$\bar{U} = \frac{1}{\tau} \int_0^{\tau} U(t)\, dt$$

or, for the exponential considered above,

$$U = \frac{1}{\tau} \int_0^{\tau} (1 - e^{-\lambda t})\, dt$$

$$= 1 + \frac{1}{\lambda\tau} (e^{-\lambda\tau} - 1)$$

$$= \frac{\lambda\tau}{2!} - \frac{(\lambda\tau)^2}{3!} + \frac{(\lambda\tau)^3}{4!} - \cdots$$

$$\approx \frac{\lambda\tau}{2}$$

Note that the often-used approximation for the frequency-weighted component unavailability assumes that (1) the failure density function is exponential and (2) higher-order terms of the exponential are negligible.

Failure in Time: Annunciated

For some components, failure is detected immediately (e.g., an annunciated failure). The probability that such a component is not available if needed is related to the frequency of failure and the average time needed to return the component to service. This unavailability is given by

$$U = \frac{\lambda T}{1 + \lambda T}$$

where λ is the failure rate and T is the average total time to respond to the failure, repair the component, and return it to service. Note that if λT is much smaller than unity, the unavailability may be approximated:

$$U \approx \lambda T$$

Failure in Time After Successful Start

It is often necessary to evaluate the probability of a component's starting successfully but failing in time before completing its mission.

---

*The term "frequency-weighted unavailability" is used here to distinguish between this quantity and a similar quantity, average (un)availability. See a reliability text, such as that by Barlow and Proschan (1975), for the definition and use of the term "average availability."

The mission time is here designated τ^*. The probability that a component fails before τ^* is given by the cumulative distribution function. For the exponential case,

$$R(\tau^*) = 1 - e^{-\lambda\tau^*}$$

$$\approx \lambda\tau^*$$

It should not be assumed that the failure rate λ in this case is the same as the failure rate in standby. Indeed, in estimating the rate for failures occurring after a successful start, the analyst must take into account any adverse environment as well as recognize differences between the rates of standby and operation failures.

Often, failure to start on demand and failure to run for some time τ^* are both included in the tree. It must be noted that failure to run is dependent on a successful start; that is, the probability of failure to run for τ^* hours must be modified by the probability of successful start. There are two possible approaches to modeling this combination in the fault trees: (1) as dependent events or (2) as one event.

If failure to start and failure to continue running after starting are separate events, they should be modeled as mutually exclusive events (see Figure 5-4).

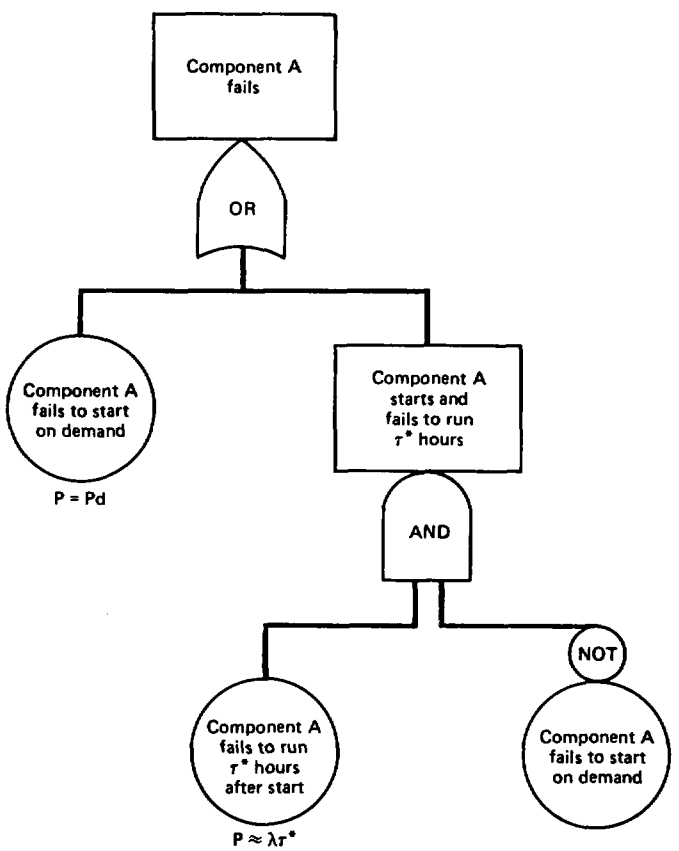

Figure 5-4. Modeling of mutually exclusive events.

If both modes are treated as one event, then

$$P_E = P_F + (1 - P_F)\ \lambda\tau^*$$

That is, the model accounts for the probability of failure to start on demand plus the probability of a successful start and failure to run for τ^* hours.

Recovery

It is possible that some events can be reversed in time to prevent core damage. There are data that provide recovery times for the loss of offsite power and emergency power. For accident sequences that are initiated by a loss of offsite power and the subsequent failure of all emergency diesels, recovery within a specified time can prevent core damage.

Such events can be broken into two parts: (1) frequency of loss or failure and (2) probability of recovery by time t, given loss or failure. This process is illustrated by the example given below, using point estimates. The data used in this example should not be taken for an actual assessment, though the results should be comparable with those of an actual assessment.

Example: Total Loss of AC Power (Station Blackout)

Loss of Offsite Power. The distribution for the duration of an offsite-power loss is given below. The data were collected from 46 sites where 45 losses occurred in 313.03 site-years, the rate of loss being .144 per site-year.

Duration (hours)	Percentage of events
<2	70
2 to 4	3
4 to 8	15
>8	12

Diesel Failure. Data from 36 plants were used to estimate the failure of diesel generators to start. If a configuration of three diesels is assumed and one diesel is needed for an adequate supply of power, the relevant probabilities for failure to start are as follows:

P(diesel 1 fails to start) = .0261

P(diesel 2 fails to start|diesel 1 has failed) = .234

P(diesel 3 fails to start|diesels 1 and 2 have failed) = .552

P(all three diesels fail to start) = .00337

The repair-time probabilities are

P(diesel not repaired within 2 hours) = .66

P(diesel not repaired within 4 hours) = .47

P(diesel not repaired within 8 hours) = .23

Probability of Station Blackout Given Duration. First we define the following:

D = duration of station blackout

L = duration of loss of station power

G = duration of diesel unavailability

S = event station blackout occurs in a year

Then for some period of time t,

$$P(D > t|S) = P(L > t \text{ AND } G > t|S)$$

$$= P(L > t|S)\ P(G > t|S) \quad \text{(assuming independence)}$$

If F_D is the failure of all diesels on demand and F_L is the loss of offsite power in a year, then assuming independence between diesel and offsite-power failures,

$$P(S) = P(F_D)\ P(F_L)$$

the probabilities being

$$P(F_L) = .144$$

$$P(F_D) = .0034$$

and

$$P(S) = 4.9 \times 10^{-4}\ yr^{-1}$$

Then

$$P(S \text{ and } D > t) = P(D > t|S)\ P(S)$$

For t = 2 hours:

$$P(S \text{ and } D > t) = (.30)\ (.66)\ (4.9 \times 10^{-4})$$

$$= 9.7 \times 10^{-5}\ yr^{-1}$$

For t = 4 hours:

$$P(S \text{ and } D > t) = (.27)\ (.47)\ (4.9 \times 10^{-4})$$

$$= 6.2 \times 10^{-5}\ yr^{-1}$$

For t = 8 hours:

$$P(S \text{ and } D > t) = (.12)\ (.23)\ (4.9 \times 10^{-4})$$

$$= 1.3 \times 10^{-5}\ yr^{-1}$$

5.3.1.2 Demand Model

Another type of model for describing component failures is the demand model. It is used to describe the failure of a component at the time of a demand for its use. The number of failures in n trials is described by the binomial distribution, and the demand model is appropriate for components that are in a dormant state until the moment of need, when they are switched on. The underlying assumption is that at each demand the probability of failure is independent of whether or not a failure occurred at any previous demand. The demand model is one that will be carried through this chapter and has been commonly used in PRAs.

The equation for the binomial distribution is as follows:

$$Pr(X \leq r) = \sum_{x=0}^{r} \binom{n}{x} p^x (1 - p)^{n-x} \qquad (5\text{-}2)$$

It gives the probability of r or fewer failures in n independent trials, given the probability of failure in a single trial is p. The parameter needed in this model is p, the probability of failure at each demand.

5.3.1.3 Demand Model vs. Time-to-Failure Model

Several very important factors should be taken into account when using the demand model. If the event being considered really could occur before the demand, then using the demand model "lumps" the failure rate into the instantaneous time of the demand. Thus, for different demand rates the probability of failure would actually be different, and if the demand model is used, a reasonable estimate is obtained only if the demand rates are similar. A component that behaves exactly as the demand model will have the same probability of failure on demand whether the demand occurs once per hour or once per decade.

The relationship between a failure-on-demand model and a failure-in-time model (assuming a constant failure rate) can easily be seen mathematically. The following assumptions are typical of this situation:

1. Component failures can be detected only at tests that occur every τ hours.

2. Components found failed are immediately repaired or replaced; components found operable are returned to service in working condition.

The data from such a situation yield x failures in N tests. The probability of failure on demand is P = x/N. Note that the results from successive tests are independent and that the exponential distribution allows a component to be considered as good as new after the test. Thus the number of tests failed has a binomial distribution with parameters N and $1 - e^{-\lambda\tau}$. The maximum-likelihood estimate (MLE) of $1 - e^{-\lambda\tau}$ is x/N, and thus the MLE of λ is

$$\hat{\lambda} = -\frac{1}{\tau} \ln(1 - P)$$

For small P, $\hat{\lambda} \approx P/T$, which is the usual estimate for $\hat{\lambda}$. However, this approximation is nonconservative. For example, if half the tests are failed,

$$\hat{\lambda} = \frac{\ln 2}{\tau} = \frac{0.69}{\tau}$$

where the approximation yields

$$\hat{\lambda} \approx 0.5/\tau$$

If it is necessary to obtain a new probability of failure on demand, P_1, for a new test period τ_1, the above relationships must be considered. The new demand probability is

$$\hat{P}_1 = 1 - \exp(-\hat{\lambda}\tau_1)$$

$$= 1 - \exp\left[-\frac{\tau_1}{\tau} \ln(1 - P)\right]$$

$$= 1 - (1 - P)^{\tau_1/\tau}$$

For example, if $P = 1 \times 10^{-2}$, τ = 720 hours (1 month), and τ_1 is 1 year, then $\tau_1/\tau = 12$, and

$$\hat{P} = 1 - \left[1 - (1 \times 10^{-2})\right]^{12} = 1.14 \times 10^{-1}$$

5.3.2 TEST CONTRIBUTIONS TO COMPONENT UNAVAILABILITY

Some test activities render a component or group of components unavailable to the system should a demand occur. Such an activity should appear on the appropriate tree as a separate event.

The probability that a component will be in testing when a demand occurs is simply the frequency of the test multiplied by the average

duration of the test, normalized by the time between the start of tests. For example,

$$P_T = \frac{(1 \text{ test/month})(L_T \text{ hr})}{730 \text{ hr/month}}$$

Here L_T is the average length of a test that occurs once every month.

The model often used in PRAs for the time to complete a test is the lognormal distribution. Although this assumption has not been extensively tested, several studies have found the lognormal distribution to provide a reasonable fit (Lapides, 1975; USNRC, 1975, Appendix III; McClymont and McLagan, 1982).

The equation for the lognormal distribution is

$$C(t) = \frac{1}{\sigma\sqrt{(2\pi)}} \int_{-\infty}^{\ln t} \exp\left[-\frac{(y-\mu)^2}{2\sigma_2}\right] dy \quad \sigma^2 \qquad (5\text{-}3)$$

This equation represents the cumulative probability that the event has been completed by time t. The parameters σ and μ can be expressed in other terms:

$$\mu = \ln M$$

$$\sigma = \frac{\ln(EF)}{1.64}$$

where the parameter M is the median time to completion and the error factor EF is the quantity that, when multiplied by the median, gives the time of completion that is equal to or longer than 95 percent of all times to complete the event.

Sections 5.5.1 and 5.5.2 show how to estimate the parameters of a lognormal time-to-completion distribution as either distributions or point estimates with confidence limits. Methods for propagating these uncertainty measures can be found in Chapter 12. These methods can be used to estimate the distribution or point estimate with confidence limits for P_T from the parameter distributions or point estimates and confidence limits. The quantity P_T is then the input required for the accident-sequence quantification discussed in Chapter 6.

5.3.3 MAINTENANCE CONTRIBUTIONS TO COMPONENT UNAVAILABILITY

A maintenance act is considered to be any unscheduled activity that causes a component or system to be taken out of service. It may be expected that repair takes place, but this repair may vary from the very simple to the very complex.

The evaluation of the maintenance contribution is similar to that of testing, except that maintenance acts occur randomly in time, whereas for

tests the time is fixed. The Reactor Safety Study (USNRC, 1975, Appendix III), for example, found that the time of maintenance for all components could be modeled by a lognormal distribution with 5th and 95th percentile points of 1 and 12 months, respectively. In most cases, it may be expected that the frequency of maintenance will exceed the frequency of failure for a component in the fault tree because the number of component failures requiring maintenance far exceeds the number of failures that completely negate a component's ability to function in its safety role. A good example is a motor-operated valve that must open to successfully perform its safety role. Failure to open occurs less frequently than valve-stem leaks, which require the valve to be taken out of service for repacking, but do not directly negate the safety role of the valve.

The probability that a component is in maintenance when a demand occurs is shown below as

$$P_M = \frac{f_M L_M}{1 + f_M L_M}$$

In this expression, f_M is the average frequency of required maintenance and L_M is the average length of the maintenance.

The lognormal distribution (see Equation 5-3) can be used for the time to complete maintenance, while the frequency of occurrence may be lognormal or exponential. Sections 5.5.1 and 5.5.2 show how to estimate the parameters of both the lognormal and the exponential distributions as either distributions or point estimates with confidence limits. Chapter 12 gives the methods for propagating the distribution or point estimate with confidence-limit parameters to the event P_M, which will then be a distribution or a point estimate with confidence limits. The quantity P_M, then, is the required input for accident-sequence quantification (Chapter 6).

5.3.4 INITIATING-EVENT MODELS

Initiating events are the occurrences that initiate an accident sequence. The desired measure for such events is frequency. A plant may experience tens of these events per year or only one in 10,000 years.

Initiating events are assumed to occur randomly in time, and they are usually assumed to occur at a constant rate. However, data on events that occur more frequently indicate that the rate of occurrence may be higher during the plant's first years than during subsequent years. There are insufficient data to predict whether or not the frequency of these initiators might increase in later life.

For purposes of this chapter it is assumed that the model for initiating events will be based on a constant rate of occurrence (the Poisson model).

It should be noted that in most PRAs initiating events are treated as single events. However, the initiating event can be quantified by

combining several events. This combination can be accomplished through a fault tree, an event tree, or a similar tool. While this may not affect the underlying event modeling and data analysis, it may require quantification tools that differ from those used to evaluate system/sequence frequency-weighted unavailability via fault trees, event trees, etc. That is, it may be necessary to quantify the synthesized initiating event as a frequency, rather than a probability.

5.4 DATA GATHERING

Before collecting and analyzing data, it is important to know what kind of data are needed. In a PRA the events of interest are modeled as events that occur randomly. In general, they occur either randomly in time or randomly at each challenge. Thus, for each classification of events, data will be either x events in time T or x events in n trials (or demands). In addition, if it is necessary to test the component-reliability models, the actual time history of the failures is needed. More specifically, if the failure of motor-operated valves to open when needed is a class of events to be evaluated, it will be necessary to search data sources to determine the number of occurrences for this event, either the number of demands or the time over which these events occurred, and when each failure to open occurred. It will also be useful to examine other data bases for information about the event of interest.

In general, for events involving components in safety systems, the quantity of interest is the probability that the component cannot perform its intended function when the initiating event occurs.

Thus, the objective of the data-gathering task is to obtain the raw information needed for estimating the event-model parameters identified in the preceding section: (1) the number of failures in time or the number of demands for reliability models; (2) the frequency and duration of tests for systems or components; (3) the frequency and duration of maintenance on components; and (4) the frequency of initiating events. The data may also be used to test the applicability of the event model; in this case, it is necessary to have the time of each failure. The sources of data may include plant records, existing data reports, and previous PRAs. This section describes various sources of available data and their attributes; it then discusses the process of data collection. It is strongly recommended that representative existing data sources be closely examined to establish clearly the type of data needed before beginning the collection of plant data.

5.4.1 EXISTING DATA SOURCES

As the data analyst proceeds to determine the appropriate reliability data, he finds a spectrum of available resources. In some cases a clearly appropriate source is available. In other instances, however, there are few sources of data whose content and format allow unambiguous selection. The

data analyst must decide on the appropriateness of the data he examines. The data source does not always specify what failure modes or mode is represented; whether, for example, the pump driver is included in all pump failures; what environment is applicable; or what the total population is. Often, additional research may be needed to discover the information not available in the reported data. Discussed below are the following sources that may be useful in building a data base for a PRA:

1. A report (EPRI, 1982a) on anticipated transients without scram.

2. A report (EPRI, 1982b) on the loss of offsite power at nuclear power plants.

3. A report (McClymont and McLagan, 1982) on diesel-generator reliability at nuclear power plants.

4. Data summaries of the licensee event reports submitted to the Nuclear Regulatory Commission.

5. The Reactor Safety Study (USNRC, 1975).

6. An IEEE data manual on electronic, electrical, and sensing components.

7. The Nuclear Plant Reliability Data System.

8. The National Electric Reliability Council.

A substantial number of other sources are summarized in Appendix C.

ATWS: A Reappraisal, Part III, "Frequency of Anticipated Transients," EPRI NP-2330. Published in 1982 by the Electric Power Research Institute (EPRI), this report contains information on the type and frequency of initiating events that lead to reactor scram. The information was gathered from about 60 percent of the nuclear power plants in the United States. Initiating events like pipe breaks are not included. The data are presented as incidents that resulted in a reactor scram and are sorted into categories. Since data analysis is minimal, the user must extract the information as needed and perform the necessary analysis.

Loss of Off-Site Power at Nuclear Power Plants: Data and Analysis, EPRI NP-2301. This 1982 report presents data on the frequency of loss and subsequent recovery of offsite power at nuclear power plants. The data were collected from the sites of 47 plants. Results are presented as events per site and by National Electric Reliability Council region. Data analysis includes point estimates for frequency with confidence limits, assuming a constant rate of occurrence. Recovery time is analyzed with a lognormal distribution for the time to recover. All raw data are reported to allow the user to perform his own analysis. This document is the most comprehensive source of data on the loss of offsite power for PRA usage.

Diesel Generator Reliability at Nuclear Power Plants: Data and Preliminary Analysis, EPRI NP-2433 (McClymont and McLagan, 1982). This report presents data related to the reliability of emergency diesel generators.

The sources include plant records, utility records, and licensee event reports submitted to the Nuclear Regulatory Commission. The data include both raw information and estimates of event-model parameters. The report details failure to start, failure to continue running, and repair times.

Data Summaries of Licensee Event Reports at U.S. Nuclear Power Plants. Published by the Nuclear Regulatory Commission, these data summaries are available as six separate reports:

1. Diesel Generators (NUREG/CR-1362; EG&G-EA-5092).

2. Pumps (NUREG/CR-1205; EG&G-EA-5044).

3. Valves (NUREG/CR-1363; EG&G-EA-5125).

4. Selected Instrumentation and Control Components (NUREG/CR-1740; EG&G-EA-5388).

5. Primary Containment Penetrations (NUREG/CR-1730; EG&G-EA-5188).

6. Control Rods and Drive Mechanisms (NUREG/CR-1331; EG&G-EA-5079).

They describe the results of analyses of component failures reported to the Nuclear Regulatory Commission in licensee event reports. Component failures are reported for individual plants, by reactor vendor, by failure mode, and for all plants considered together. Included are failure rates, failures on demand, and some information on repair times. The estimates of event-model parameters, however, are based on estimates of population, demands, and exposure time. Hence, the statistical analysis includes estimated information together with actual plant data.

Reactor Safety Study, WASH-1400, Nuclear Regulatory Commission, 1975. Appendix III of this report, "Failure Data," contains the failure data used in the study, including raw data from 1972, notes on test time, notes on maintenance time and frequency, the results of a human-reliability analysis, aircraft-crash probabilities, estimates of the frequency of initiating events, and some information on common-cause failures. From the assembled information, this appendix also defines the "assessed range" for each failure rate. The authors state, however, that "this data may not be sufficiently detailed, general, or accurate enough for use in other quantitative reliability models or in applications involving greater specificity."

IEEE Project 500 Data Manual, Institute of Electrical and Electronics Engineers, Inc. This document contains data for electronic, electrical, and sensing components. The reported values are mainly synthesized from the opinions of some 200 experts. Each expert has submitted a low, a recommended, and a high value for the failure rate under normal conditions and a maximum value that would be applicable under all conditions (including abnormal ones). The pooling of estimates was done by geometric averaging, a method judged to be a better representation of expert estimates, which are often given as negative powers of 10. While some estimates include hard

data, the reader is not made aware of which estimates are based only on opinion, on hard data, or a combination of both.

Nuclear Plant Reliability Data System (NPRDS), Southwest Research Institute. The NPRDS collects failure data on safety-related systems and components. At present, 61 plants are reporting data. The data are compiled and disseminated in periodic reports to the participants of the program and other potential users. In addition, special searches of the data base may be requested by the participants and others, or the users can access the data through their computer terminals. Typical information that NPRDS provides includes the following:

1. The plant operating mode (i.e., operating, standby, and shutdown).
2. The calculated in-service hours of the system.
3. Outage times.
4. Number of failures per million in-service hours.
5. Number of applicable tests.
6. Number of actuations for standby equipment.
7. Component failure modes and effects.

The main disadvantage is the dependence of the NPRDS on regular participant reporting. If no report is received from a participant in a reporting period, it is assumed that no failures have occurred. In the near future, data from plants with irregular reporting will be filtered from the data base to avoid this disadvantage.

National Electric Reliability Council (NERC). On January 1, 1979, the Edison Electric Institute (EEI) transferred to NERC the responsibility for operating its equipment-availability data system--the prime utility-industry source for the collection, processing, analysis, and reporting of information on power-plant outages and overall performance. The Unit Year Summary computer program produces a report for each individual unit, including statistics for the latest year and cumulative statistics for the life of the unit. In addition, the Equipment Availability Task Force produces annually a report on equipment availability for a 10-year period. Finally, the EEI has established a procedure for processing special requests for the analysis of reliability data.

5.4.2 COMPONENT-DATA COLLECTION FROM NUCLEAR POWER PLANTS

At present, no nuclear plant keeps records of component reliability for the specific purpose of using them as data for risk assessments. The PRAs that have been conducted to date have had to depend on other sources for plant-specific data. These sources include many plant records and procedures that may be available to the PRA analysts. The usefulness of a particular source depends on the reliability models chosen to represent components in system fault trees. On the other hand, the availability (or the absence) of various data sources may affect the choice of models by a system analyst. Table 5-1 lists the most common parameters used to represent components, the data required to derive estimates of the parameters, and the potential sources of such data at plants. How these sources can be used to extract needed information is briefly explained below.

Table 5-1. Sources of plant data

Parameter	Data requirements	Potential sources
1. Probability of failure on demand	a. Number of failures	Periodic test reports, maintenance reports, control-room log
	b. Number of demands	Periodic test reports, periodic test procedures, operating procedures, control-room log
2. Standby failure rate[a]	a. Number of failures	See 1a above
	b. Time in standby	Control-room log
3. Operating failure rate[a]	a. Number of failures	See 1a above
	b. Time in operation	Control-room log, periodic test reports, periodic test procedures
4. Repair-time distribution parameters	Repair times	Maintenance reports, control-room log
5. Unavailability due to maintenance and testing	Frequency and length of test and maintenance	Maintenance reports, control-room log, periodic test procedures
6. Recovery	Length of time to recover	Maintenance reports, control-room log
7. Human errors[b]	a. Number of errors b. Opportunities	Maintenance reports, control-room log, periodic test procedures, operating procedures

[a]See Section 5.3.1.1.

[b]While this chapter does not deal with the evaluation of human errors, it is likely that a search for plant-specific data would find human-error data to supplement the analysis methods described in Chapter 4.

5.4.2.1 Periodic Test Reports and Procedures

Periodic test reports and procedures are a potential source of data on failures, demands, and operating time for components that are tested periodically. Test reports for key components or systems typically contain a description of the test procedure and a checklist to be filled out by the

tester as the steps are performed. For example, in an operating test of an emergency diesel generator, the procedure may call for starting the diesel and running it for an hour. The record of a specific test would report whether or not the diesel started and whether it ran successfully for the entire hour. Another example is a test of emergency system performance, in which the procedure calls for the tester to give an emergency signal that should open certain flow paths by moving some motor-operated valves and starting one or more pumps. The position of the valves and the operation of the pump are then verified, giving records of whether the valves and pumps responded successfully to the demands. As shown by these examples, records of periodic tests provide a self-contained tally of demands on some components, as well as the failure (and success) of the component given these demands.

When failures are reported in periodic tests, however, the failure mode should be examined carefully, if possible, before the failure is included in a failure-parameter estimate to be used in system fault trees. In the diesel-generator example, the report may note that the result of the test was unsatisfactory because the diesel tripped on a signal of low oil pressure, high oil temperature, or the like. Since many of these trips are disabled by a LOCA signal, such an event should not be counted in deriving a failure-parameter estimate for a fault tree that is part of a LOCA sequence, even though the test report indicated an unsatisfactory performance by the diesel generator. If, on the other hand, the diesel would have failed if the trip was bypassed, it must be counted as a failure. Similarly, a test report on diesel-generator operability may log an unsatisfactory result due to an air-compressor failure. Such a failure would cause a diesel-generator failure to start only if it occurred in conjunction with a leak in the diesel air tank. In this instance, the test report indicates a failure even though no actual demand was placed on the diesel.

If the records of actual periodic tests are not readily available, the test procedures can be used to estimate the number of testing demands or the operating time during tests for a component over a period of time. To do this, the number of demands or the operating time of a single test can be multiplied by the frequency of the test and the pertinent calendar time. Of course, this approach is valid only if the tests are conducted at the prescribed frequency. Some tests may in fact be conducted at more frequent intervals than those stated in the procedures. Plant personnel should be interviewed to determine what adjustments are necessary.

If this approach is used, a count of failures must be obtained from different sources (e.g., maintenance reports). Since these sources may not indicate clearly which failures occurred during the periodic tests considered, the failure-parameter estimates derived by this approach are probably conservative. In order to correctly match failures with demands or operating time for a component, the number of demands or the duration of operating time occurring outside periodic tests must be obtained. Such information is usually much more difficult to extract from typically available data sources.

5.4.2.2 Maintenance Reports

Reports of maintenance on components are potential sources of data on failures, repair times after failure, and other unavailability due to maintenance. These reports typically include the following:

1. A plant identification number for the component undergoing maintenance and a description of the component.

2. A description of the reason for maintenance.

3. A description of the work performed.

4. An indication of the time required for the work or the duration of the component's unavailability.

The report may indicate that maintenance was needed because the component failed to operate adequately or was completely inoperable. Such an event may then be added to the count of component failures. The maintenance report often gives information about the failure mode and mechanism as well as the amount of time spent on repair after the failure was discovered. Such information must be interpreted carefully, because the actual repair time may cover only a fraction of the time the component was unavailable between the detection of the failure and the completion of repairs. In addition, the repair time is often given in terms of man-hours, which means that the actual time spent on repair could be shorter, depending on the size of the work crew; the use of recorded man-hours would therefore lead to a conservative estimate of repair time. The complete out-of-service time for the component can, however, be derived, because the maintenance record often states the date on which the failure was discovered and the date on which the component was made available after repair.

Maintenance reports that record preventive maintenance can be used to estimate the contributions of these actions to component unavailability. Again, the report may show that a component was taken out of service on a certain date and restored some time later, giving a sample of the duration of maintenance. The frequency of these events can be derived from the number of preventive-maintenance reports in the calendar time considered.

Unfortunately, not all maintenance reports present all of the information listed above. Often, the descriptions of a component's unavailability or the work performed are unclear (or missing altogether), requiring guesswork as to whether an unfailed component was made unavailable by maintenance or whether the maintenance was the result of component failure. An additional problem that has already been mentioned is the difficulty in matching up the failures recorded in maintenance reports with the demands or operating times reported in other documents.

5.4.2.3 Operating Procedures

Operating procedures can be used to estimate the number of demands on certain components in addition to demands occurring during periodic tests.

This estimate is obtained by multiplying the number of demands imposed on a component during a procedure by the number of times the procedure was carried out during the calendar time of interest. Unfortunately, the latter number is not always easily obtained. For procedures followed during plant startup or shutdown, the number of times the procedure was performed should be readily obtainable, but for procedures followed during operation, this information will be available only from the control-room log.

5.4.2.4 Control-Room Log

Many of the gaps in a component-reliability data base compiled from test and maintenance records can be filled by examining the control-room log, which is a chronological record of important events at the plant. For example, the log has records of demands made (e.g., pumps and diesel generators) at times other than periodic tests. It notes the starting and stopping times for these components, thus supplying operating-time data. The log also notes the initiation of various operating procedures, thus adding to the information about demand. Furthermore, it records periods when certain components and systems are out of service, and in this the log is often more accurate than the maintenance reports.

There is, however, a problem with using the control-room log as a source of component data: all events in the log are listed chronologically, without being separated by system, type of event, or any other category. The analyst must therefore search through many irrelevant entries to find those needed for the data base. The additional accuracy that is supplied to the estimates of component-failure parameters by data from the log may not be worth the effort needed to search through several years of the plant history recorded in the log.

5.5 ESTIMATION OF MODEL PARAMETERS

After model selection, the parameters of the models can be estimated. Two methods of estimation are described in this chapter and are complemented by the relevant methods in Chapters 6 and 12: (1) classical methods and (2) Bayesian methods.

A Bayesian analysis allows the augmentation of available data by quantified personal opinion. The analyst quantifies his belief about the parameters (unknown constants) in the model, exclusive of the information in the data, by a probability distribution; that is, he not only models the occurrence of accidents probabilistically but also develops a probability model for his beliefs about such occurrences. The data analyst should be aware that this may be difficult to do, and it will be even more difficult to convince the community at large to adopt his degree of belief as their own.

In a classical analysis, knowledge and expertise also play a role, but less formally, in general serving only as aids in choosing probability models and relevant data. For example, data obtained under normal operating conditions may or may not be applicable to accident conditions. An understanding of the situation is needed to resolve this question. Once such questions are resolved, a classical analysis lets the data "speak for themselves." The users of a classical analysis must be aware that limited data can lead to imprecise estimates. Though the introduction of a quantified degree of belief can improve the apparent precision of risk estimates, it may be useful and informative to do both a Bayesian and a classical analysis, thus allowing the reader of a PRA to separate the data and the belief components of the results.

5.5.1 CLASSICAL ESTIMATION

5.5.1.1 Point Estimation

Reliability and availability models involve a variety of parameters, such as component-failure rates and expected repair times, that need to be estimated in order to estimate the probability of specific accident sequences. Choosing a point estimate can involve a variety of considerations, depending on the information available. If data are available and it is desired to obtain estimates that are strictly functions of the data, then, for the models commonly used in risk analysis, point estimators are well established. The point estimators generally used for the binomial, Poisson, and lognormal models, and appropriate data, are given below.

Binomial Distribution. The data, parameter, and estimate for binomial models are as follows:

Data: f failures in n demands. The number of demands is known; the outcomes, success or failure, are statistically independent; and the failure probability is constant across these demands.

Parameter: p, the probability of failure on demand (dimensionless).

Estimate:

$$p^* = f/n$$

Poisson Distribution. For Poisson models, the data, parameter, and estimate are the following:

Data: f failures (or occurrences of an initiating event) in T time units. The quantity T is known; failures occur independently and at a constant rate in time and across different items, which may be combined to obtain the data.

Parameter: λ, the failure rate (number of failures per unit time).

Estimate:

$$\lambda^* = f/T$$

Lognormal Distribution. The data, parameters, and estimates for lognormal models are as follows:

Data: n independent positive observations, $X_1, X_2, \ldots, X_n$, such as repair times, whose logarithms are modeled as being normally distributed.

Parameters: μ, the expected value of $t = \log_e(X)$ and σ^2, the variance of t.

Estimates:

$$\mu^* = \sum_{i=1}^{n} \frac{t_i}{n} = \bar{t} \quad \text{for the sample mean}$$

$$\sigma^{2*} = \frac{\sum(t_i - \bar{t})^2}{n - 1} = s_t^2 \quad \text{for the sample variance}$$

All the estimates given here are unbiased, which means that, on the average, they equal the parameter being estimated. Moreover, all but σ^{2*} are maximum-likelihood estimators. Additional details pertaining to these estimates are available in a text by Mann et al. (1974), which also provides statistical estimators for other models, such as the Weibull and gamma distributions, and other situations, such as a fixed number of failures/random operating-time estimates of the failure rate λ.

Classical point estimates are attempts to identify single parameter values indicated by the data. As such, they are data summaries, and information is necessarily lost in the summarization. The loss is serious in the case of point estimation because the amount of data going into the estimates is lost. For example, one failure in 10,000 hours yields the same point estimate of a failure rate as do ten failures in 100,000 hours, but clearly more information is present in the latter case. If this information is ignored or not communicated, an incomplete analysis results. Two classical methods by which the amount of information pertaining to parameters of interest can be conveyed are standard errors and statistical confidence intervals.

5.5.1.2 Standard Errors

If the data-yielding process described above is repeated, the parameter estimates will vary; that is, in another n demands or T time units, the number of failures will vary (in a manner described by the probability models used to analyze those data). Furthermore, the n repair times collected in the future would differ from those observed at present. The variance over such repetitions of the estimators described above provides a measure of the information contained in the point estimates obtained. The larger the variance, the less reliable the point estimate. In general, the variance of an estimator is not known, but it can be estimated in these cases. The square

root of the estimated variance of an estimator is termed the "standard error of the estimate." For the parameters considered in the preceding section, the standard errors (s.e.) are as follows:

Binomial:

$$\text{s.e. } (p^*) = \left[\frac{p^*(1 - p^*)}{n}\right]^{1/2}$$

Poisson:

$$\text{s.e. } (\lambda^*) = \left(\frac{\lambda^*}{T}\right)^{1/2}$$

Lognormal:

$$\text{s.e. } (\mu^*) = \frac{\sigma^*}{n^{1/2}}$$

$$\text{s.e. } (\sigma^{2*}) = \sigma^{2*}\left(\frac{2}{n - 1}\right)^{1/2}$$

(The information contained in an estimated variance is usually conveyed by reporting the degrees of freedom, n - 1 in the case considered here, rather than a standard error.)

One way in which standard errors are used is to obtain approximate classical confidence limits on the parameter of interest. For example, the point estimate plus or minus twice its standard error provides a crude 95-percent confidence interval on the parameter. Thus, a large standard error, relative to the point estimate, indicates that the data do not provide a very clear indication of the parameter. If only a point estimate is given, this information about the data is lost, and an unwarranted and misleading aura of precision may result. Without standard errors, any comparison of point estimates, say for the purpose of ranking accident sequences, may be misleading.

5.5.1.3 Interval Estimation

A given set of data, say f failures in T hours, can occur in sampling from a variety of Poisson distributions. That is, many other values of λ besides $\lambda^* = f/T$ can give rise to this particular outcome. Some values of λ, however, are more consonant with the data than others. This realization is the basis for classical confidence intervals, whose purpose is to identify ranges of parameter values that are consonant with the data to some specified extent. For example, suppose an upper 95-percent limit on λ is found to be $\lambda_{95} = 10^{-4}$ failures per hour. This means that, for λ values greater than 10^{-4}, the observed data are in the extreme 5 percent of possible outcomes; such λ values are not very consistent with the data. Values

of λ less than 10^{-4} are less unconsonant with the data. Both upper and lower confidence limits, at any specified confidence level, can be obtained, and the interval between these limits is termed a "classical confidence interval." Classical confidence intervals have the property that, in repeated sampling, the probability that the confidence interval will contain the parameter of interest is at least at the specified confidence level.

As indicated above, approximate confidence intervals on a parameter can be obtained from a point estimate and its standard error. For the three distributions considered here, though, exact confidence limits or better approximations can be readily obtained.

Binomial Distribution

The upper $100(1 - \alpha)\%$ confidence limit on p is obtained by solving

$$\alpha = \sum_{x=0}^{f} \binom{n}{x} p^x (1 - p)^{n-x}$$

for p. The lower $100(1 - \alpha)\%$ confidence limit on p is obtained by solving

$$\alpha = \sum_{x=f}^{n} \binom{n}{x} p^x (1 - p)^{n-x}$$

for p. Tables, slide rules, and computer programs are available for solving these equations (Green and Bourne, 1972; Hald, 1952). A useful approximation for small f, large n is

$$P_U(1 - \alpha) = \frac{\chi^2(2f + 2; 1 - \alpha)}{2n}$$

$$P_L(1 - \alpha) = \frac{\chi^2(2f; \alpha)}{2n}$$

where $P_U(1 - \alpha)$ and $P_L(1 - \alpha)$ are the upper and the lower $100(1 - \alpha)\%$ confidence limits, respectively, and $\chi^2(m, \gamma)$ denotes the 100 γ-percentile of the chi-squared distribution with m degrees of freedom. The interval between $P_L(\alpha)$ and $P_U(\alpha)$ constitutes a $100(1 - 2\alpha)\%$ confidence interval.

Poisson Distribution

The upper and the lower $100(1 - \alpha)\%$ confidence limits on λ are obtained by solving the following equations:

$$\lambda_U(1 - \alpha) = \frac{\chi^2(2f + 2; 1 - \alpha)}{2T}$$

$$\lambda_L(1 - \alpha) = \frac{\chi^2(2f; \alpha)}{2T}$$

Note that, mathematically, confidence limits on a failure rate λ are similar to those on a failure probability p, with time units replacing the number of demands.

Lognormal Distribution

The upper and the lower $100(1 - \alpha)\%$ confidence limits on μ are obtained from

$$\bar{t} \pm t(n - 1, 1 - \alpha)(\sigma^*/n^{1/2})$$

where $t(f,\gamma)$ denotes the γ-percentile of the Student's t distribution with f degrees of freedom.

For the upper and the lower $100(1 - \alpha)\%$ confidence limits on σ^2, the following equations are used:

$$\sigma_U^2(1 - \alpha) = \frac{(n - 1)\ \sigma^{2*}}{\chi^2(n - 1, \alpha)}$$

$$\sigma_L^2(1 - \alpha) = \frac{(n - 1)\ \sigma^{2*}}{\chi^2(n - 1, 1 - \alpha)}$$

As already discussed, classical confidence intervals supplement point estimates as a summary of the data-based information about the parameters of a probability model. They also serve to provide guidance on the parameter ranges that should be covered in a sensitivity analysis (see Chapter 12). That is, if one is interested in the change in an accident-sequence probability that results from a change in a component parameter, confidence intervals provide a plausible range over which the component parameter should be varied.

Occasionally, in probabilistic risk assessments classical confidence limits are misinterpreted as percentiles on a probability distribution of the parameter. Because confidence limits are derived under the assumption that these parameters are constants, not random variables, such an interpretation is unwarranted, except perhaps as a Bayesian degree-of-belief distribution, given a uniform prior distribution. One reason confidence limits are given a distributional interpretation is to provide input to probabilistic uncertainty analyses (Chapter 12). One could view such an analysis as a mathematical device for obtaining approximate classical confidence limits on an accident-sequence probability, given data pertaining to the parameters in the accident model, but better methods are available (Chapters 6 and 12). One particular treatment of confidence limits that should be avoided is the fitting of distributions to classical confidence limits on failure rates or probabilities.

An example of the application of classical techniques is included in Section 5.5.2.5, where the result can be compared with Bayesian treatments of the same data.

5.5.2 BAYESIAN ESTIMATION

The Bayesian approach is similar to the classical approach in that it yields "best" point estimates and interval estimates, the intervals representing ranges in which, we are confident, the parameter really lies. It differs in both practical and philosophical aspects, though. The practical distinction is in the incorporation of belief and information beyond that contained in the observed data; the philosophical distinction lies in assigning a distribution that describes the analyst's belief about the values of the parameter. This is the so-called prior distribution.

The prior distribution may reflect a purely subjective notion of probability, as in the case of a Bayesian degree-of-belief distribution, or any physically caused random variability in the parameter, or some combination of both. Physically caused random variations in a parameter like a failure rate may stem from plant and/or system effects, operational differences, maintenance effects, environmental differences, and the like. The distribution that describes this physically caused random variation in the parameter is sometimes referred to as the "population variability" distribution (Apostolakis et al., 1980) and can be represented by a Bayesian prior distribution. However, such random variation in the parameter can also be modeled by classical methods, using compound distributions in which the population-variability distribution becomes the mixing distribution. On the other hand, if the prior distribution embodies subjective probability notions regarding the analyst's degree of belief about the parameter, the Bayesian method is the appropriate framework for making parameter estimates. A comparative discussion of both interpretations of the notion of probability, the subjective and the relative-frequency notions, is given by Parry and Winter (1981).

Whether the analyst does or does not have objective relative-frequency data, he will often have other information based on engineering designs, related experience in similar situations, or the subjective judgment of experienced personnel. These more or less subjective factors will also be incorporated into the prior distribution--that is, into the description of his prior knowledge (or opinions) about the parameter.

The Bayesian method takes its name from the use of Bayes' theorem and the philosophical approach embodied in the 18th-century work of the Rev. Thomas Bayes (modern reproduction, 1958). Bayes' theorem (see Section 5.5.2.1.1) is used to update the prior distribution with directly relevant data. Here the term "generic data" will be used to refer to parameter-related information that is nonspecific to any particular plant or application, being an aggregation over more than one use condition. A prior distribution is often based on such generic data sources (Apostolakis et al., 1980). A PRA for a particular plant, of course, requires not generic data but rather estimates that are specific to the plant or application. Bayes' theorem then updates the prior distribution with plant-specific evidence and has the effect of "specializing" the prior to the specific plant. The updated, or specialized, prior is called the "posterior distribution" because it can be derived only after the plant-specific evidence is incorporated. The prior reflects the analyst's degree of belief about the parameter before such evidence; the posterior represents the

degree of belief after incorporating the evidence. Plant-specific estimates are then obtained from the posterior distribution as described in Sections 5.5.2.3 and 5.5.2.4.

5.5.2.1 Essential Elements of the Bayesian Approach

This section considers the essential elements of the Bayesian approach to data reduction. It presents a brief discussion of Bayes' theorem, the basic notions of Bayesian point and interval estimation, and a step-by-step outline of the procedures for obtaining Bayesian estimates.

The main benefit in using the Bayesian approach to data reduction is that it provides a formal way of explicitly organizing and introducing into the analysis assumptions about prior knowledge. This knowledge may be based on past generic industry-wide data and experience, engineering judgment, expert opinion, and so forth, with varying degrees of subjectivity. The parameter estimates will then reflect this knowledge. A noteworthy feature of the nuclear industry is that such prior information is often available to the extent that it may contribute more to knowledge about the parameter than does the more directly applicable (but sparse) plant-specific information.

5.5.2.1.1 Bayes' Theorem

The fundamental tool for use in updating the generic prior distribution to obtain plant- or application-specific parameter estimates is Bayes' theorem. If the parameter of interest is a failure rate λ (number of failures per unit time), Bayes' theorem states that

$$f(\lambda|E) = \frac{f(\lambda)\ L(E|\lambda)}{\int_0^\infty f(\lambda)\ L(E|\lambda)\ d\lambda} \qquad (5\text{-}4)$$

where $f(\lambda|E)$ is the posterior distribution, the probability density function of λ, conditional on the specific evidence E; $f(\lambda)$ is the prior distribution, the probability density function of λ based on generic information but incorporating no specific evidence E; and $L(E|\lambda)$ is the likelihood function, the probability distribution of the specific evidence E for a given value of λ.

If the parameter of interest is the probability of failure on demand, p, rather than a failure rate λ per unit time, then λ is simply replaced by p in Equation 5-4. However, the likelihood function will differ for the different cases, as shown in Sections 5.5.2.3.1 and 5.5.2.4.

In certain special cases, the integral on the right-hand side of Equation 5-4 can be done analytically to give a closed-form expression for the posterior distribution. The term "conjugate prior" is used to describe the prior-distribution form that conveniently simplifies the integration. For example, if the likelihood function is the Poisson distribution (see

Section 5.5.2.4), then the gamma family represents the conjugate prior: the posterior distribution will be expressible in closed form as another gamma distribution. Section 5.5.2.2.3 will discuss this in more detail. In general, a closed-form integration will not be possible, and numerical techniques must be used; alternatively, the continuous prior distribution can be approximated by a discrete approximation and the integral replaced by a sum. An example of the latter approach has been given by Apostolakis et al. (1980).

Numerical integration or a discrete approximation is often needed when the generic data include a precise description of a prior distribution, so that the analyst lacks the flexibility to choose a mathematically tractable form for it. For example, if a lognormal prior distribution is specified for λ and the likelihood is the Poisson distribution, then the posterior distribution cannot be obtained analytically in closed form. On the other hand, if we have incomplete information, this choice can be made from the conjugate family of distribution (see Section 5.5.2.2.3), which yields the mathematical convenience and resultant simplicity of a closed-form expression for the posterior distribution. Sensitivity studies can then be used to examine the effects of this choice.

The discrete form of Bayes' theorem is

$$f(\lambda|E) = \frac{f(\lambda_i)\ L(E|\lambda_i)}{\sum_{i=1}^{m} f(\lambda_i)\ L(E|\lambda_i)} \qquad (5\text{-}5)$$

where λ_i $(i = 1,2,\ldots,m)$ is a discrete set of failure-rate values. The prior and posterior distributions are approximated by the discrete functions $f(\lambda_i)$ and $f(\lambda_i|E)$, respectively.

The discrete form of Bayes' theorem is mathematically convenient and is sometimes used as an approximation to the continuous form given by Equation 5-4 when the denominator in Equation 5-4 cannot be evaluated in closed form. In such cases, the range of the parameter is carved into a set of intervals and the probability content of each interval is then associated with a single point inside the interval.

There are two important issues that should be raised in conjunction with the discrete-prior approach. First, it sometimes happens that the use of a discretized approximation to a continuous prior does not produce a meaningful well-spread posterior distribution (see Apostolakis et al., 1980, Examples 2 and 3). In such cases, the prior distribution must be finely spread in the appropriate region after the initial posterior distribution has been obtained. Thus, the method may require more than one iteration to produce a meaningful posterior, and such recursive procedures may be unacceptable. Second, if continuous priors of a specified form (e.g., a lognormal distribution) are discretized, the results may be interpreted as a crude approximation to the integration in Equation 5-4. A better approximation is to use Equation 5-4 in conjunction with an appropriate numerical integration method, such as the Gauss quadrature, thus maintaining in effect a continuous prior distribution. This is the approach used by Ahmed et al. (1981).

The denominator of either Equation 5-4 or Equation 5-5 can be thought of simply as a normalizing factor that makes the posterior distribution integrate or sum to unity. Thus, Bayes' theorem can be stated verbally as simply saying that the posterior distribution is proportional to the product of the prior distribution and the likelihood function.

5.5.2.1.2 Bayesian Point and Interval Estimation

The prior distribution summarizes the uncertainty in a parameter as reflected by prior judgment and/or the generic data sources on which the prior is based. Similarly, the posterior distribution summarizes the uncertainties in the plant-specific value of the parameter as reflected by the combined influence of both the prior distribution and the likelihood function. In either case, it is frequently desired to obtain either a point or an interval estimate of the underlying parameter.

A Bayesian point estimate is a single value that, in some precisely defined sense, best estimates or represents the unknown parameter. Two commonly used point estimates are the mean and the median (50th percentile) of the prior or the posterior distribution. The mean of a distribution is the Bayesian estimate that minimizes the average squared error of estimation (averaged over the entire population of interest), while the median is the one that minimizes the average absolute error. Thus, either the mean or the median of the prior distribution can be used as a point estimate of the unknown generic parameter; likewise, the mean or the median of the posterior distribution can be used as a point estimate of the unknown plant- or application-specific parameter. The properties of the two estimators are discussed by Martz and Waller (1982). The mean or the median would be found by conventional statistical procedures: using the prior distribution, the mean of a failure rate λ is given by

$$\mu_\lambda = \int_0^\infty \lambda f(\lambda)\, d\lambda$$

while the median is the solution to

$$F(\lambda) = \int_0^\lambda f(t)\, dt = .5$$

$F(\lambda)$ denoting the cumulative distribution function. Using the posterior distribution, the prior $f(\lambda)$ would be replaced by the posterior $f(\lambda|E)$ in Equations 5-6 and 5-7.

Now consider the problem of obtaining an interval estimate for λ, using either the prior or the posterior distribution, depending on whether one is concerned with a generic or a specific failure rate. Suppose we want a probability of $(1 - \gamma)$ that the interval estimate really includes the unknown failure rate. (For example, $\gamma = .05$ for .95 probability.) We can

obtain a 100(1 - γ)% two-sided Bayes probability interval estimate of λ by solving the two equations

$$\int_0^{\lambda_L} f(\lambda)\, d\lambda = \frac{\gamma}{2} \tag{5-8}$$

and

$$\int_{\lambda_U}^{\infty} f(\lambda)\, d\lambda = \frac{\gamma}{2} \tag{5-9}$$

for the lower end point λ_L and the upper end point λ_U. It follows immediately that $P(\lambda_L < \lambda < \lambda_U) = 1 - \gamma$. Such an interval is often called a "Bayesian confidence interval"; we avoid that term here because it is not a confidence interval in the classical sense. The coefficient (1 - γ) is the subjectively defined probability that the interval estimate (λ_L, λ_U) contains λ.

For a Bayesian interval estimate of an unknown plant-specific failure rate, the posterior distribution $f(\lambda|E)$ would replace the prior distribution $f(\lambda)$ in Equations 5-8 and 5-9. The interval estimate (λ_L, λ_U) would then be such that $P(\lambda_L < \lambda < \lambda_U\ E) = 1 - \gamma$.

Analogous results hold when the parameter of interest is a failure-on-demand probability p rather than a failure rate λ.

5.5.2.1.3 Step-by-Step Procedure for Bayesian Estimation

The PRA analyst goes through several steps in Bayesian data reduction. For estimating a parameter like a component-failure rate or a failure-on-demand probability, the steps are as follows:

1. Identify the sources and forms of generic information to be used in selecting an appropriate prior distribution for the parameter (see Section 5.5.2.2.1).

2. Select a prior-distribution family if none has been specified as part of the generic information (see Sections 5.5.2.2.2 and 5.5.2.2.3).

3. Choose a particular prior distribution by reducing and/or combining the generic data from step 1 (see Sections 5.5.2.2.4 through 5.5.2.2.8).

4. Plot the prior and summarize it by determining its mean, variance, and selected summary percentiles.

5. If generic estimates are required, determine them from the prior as in Section 5.5.2.1.2.

6. If plant- or application-specific estimates are required, then--

   a. Obtain data representing operating experience with the specific component.

   b. Identify an appropriate form for the likelihood function (see Sections 5.5.2.3.1 and 5.5.2.4.1).

   c. Use Bayes' theorem to get the posterior distribution (see Section 5.4.2.1.1).

   d. Plot the posterior distribution on the same page with the prior and summarize the posterior in the same manner as in step 4.

   e. Compare the prior and the posterior distributions to see the effect of the specific data.

   f. Obtain the desired estimates from the posterior distribution.

7. Investigate the sensitivity of the results to the prior distribution.

5.5.2.2 Determining Prior Distributions

A fundamental part of any Bayesian estimation procedure is the selection and fitting of a prior distribution. This section considers "generic" data that can be used to determine a prior distribution, including sample sources of such data, and then discusses some methods for reducing or combining such data in fitting a prior. Subsequently, several classes of priors that have been found useful in reactor applications will be introduced. Particular emphasis is given to the class of noninformative prior distributions, useful when there are few or no prior generic data. Lognormal, gamma, and beta prior distributions are presented for possible use when prior generic data are available.

5.5.2.2.1 Sources of Data for Use in Bayesian Estimation

Three types of information about the reliability parameter of interest are often available: (1) engineering knowledge about the design, construction, and performance of the component; (2) the past performance of similar components in similar environments; and (3) the past performance of the specific component in question. The first two types constitute the "generic" information (or data) and may include varying degrees of subjective judgment. The third type, constituted of objective data, is the "plant- or application-specific" information (or data).

Generic Data

Generic data may be available in many forms. The analyst may have raw (unreduced) failure data or reduced failure-rate data in the form of point or interval estimates, percentiles, and so forth.

Two sources of failure-rate data that have been previously used (Apostolakis et al., 1980) in nuclear plant PRAs are the Reactor Safety Study (RSS) and the IEEE Std-500 Data Manual. The RSS data have been updated in a recent report (Murphy, 1980) that summarizes the generic (and some specific) component-failure-rate data that are currently available for nuclear plant PRAs. The use of both of these sources is described by Apostolakis et al. (1980).

Another method of using raw generic data for determining a prior distribution is described by Kaplan (1981a); it uses Bayes' theorem to determine the prior distribution.

Plant- or Application-Specific Data

There are several sources of plant- or application-specific data that can be used via Bayes' theorem to determine posterior distributions suitable for application-specific estimates. Reliability data bases like the Nuclear Plant Reliability Data System (NPRDS), the In-Plant Reliability Data System (IPRDS), and the NRC licensee event reports (LERs), all of which report on component populations and failure events, are good sources of plant-specific data. Such data are also often available in summary form in secondary reports derived from these basic sources.

5.5.2.2.2 Noninformative Prior Distributions

"Noninformative" prior distributions are a class of priors that loosely minimize the relative importance of the prior (compared with the data) in generating a posterior estimate. There are many ways of precisely quantifying this basic notion and hence a variety of classes of noninformative priors and corresponding methods for their attainment in practice. The notion adopted here for the noninformative prior is that of Martz and Waller (1982), in which, roughly speaking, a prior is said to be noninformative if the plant-specific data serve only to change the location of the corresponding likelihood and not its shape. This and other notions have also been discussed by Jeffreys (1961), and a summary of the relevant literature on this subject has been presented by Parry and Winter (1981).

Noninformative priors are useful when little or no generic prior information is available; they should not be used when there is such information, because they deliberately downgrade its role in the estimation process. Frequently, Bayesian estimates from noninformative priors are identical with, or very close to, the classical estimates, a fact illustrating the versatility of the Bayesian method. However, interval estimates generated by their use are probability intervals, not classical confidence intervals. Section 5.5.2.3.2 presents the noninformative prior

for failure-on-demand probabilities, and Section 5.5.2.4.2 does so for failure rates. Since noninformative priors contain no generic information, it may be preferable to avoid their use when even minimal generic prior data are available.

5.5.2.2.3 Natural Conjugate Prior Distributions

Natural conjugate prior distributions have the property that, for a given likelihood function, the posterior and prior distributions are members of the same family of distributions. In such cases, the posterior distribution has a closed-form analytical representation (at least to the extent that the prior does), and accordingly the expressions for computing the Bayesian point and interval estimates can usually be represented in terms of well-defined probabilities. This will be seen in Sections 5.5.2.3.3 and 5.5.2.4.3. The parameters of such priors are often especially easy to interpret, playing the role of prior failure data entirely analogous to the specific data used in the likelihood function. This will also be illustrated in Sections 5.5.2.3.3 and 5.5.2.4.3. Such families of priors are often rich enough and flexible enough to permit the analyst to model reasonably a wide range of prior data that may be encountered (Martz and Waller, 1982). Finally, there are well-developed methods for fitting natural conjugate priors to generic prior data. Some of these will be discussed in Sections 5.5.2.2.6 and 5.5.2.2.7.

For these reasons, natural conjugate priors have found application in nuclear plant PRAs (see, for example, Apostolakis and Mosleh, 1979). Their use is recommended (see, for example, Ahmed et al., 1981) whenever the exact form of the prior has not been specified as part of the generic prior data, but the data are sufficient to determine a reasonable member of the natural conjugate family. If incomplete information exists on the prior, as often happens, the analyst will have the flexibility to select the form of the distribution, and the conjugate prior is often the natural selection. However, a sensitivity analysis should be performed to confirm this choice.

5.5.2.2.4 Using Generic Data Sources

The generic prior data must be reduced to a form that permits the selection of a specific prior distribution from a suitable family. For example, if a lognormal family has been selected, the two lognormal parameters must be determined from the generic data. If there are multiple sets of generic prior data, these must likewise be reduced to a common consensus prior.

A Single Source

For convenience consider the case of failure-rate (per unit time) estimation. If a two-parameter prior distribution is to be fitted, such as a lognormal or a gamma distribution, the generic data must contain at least two independent pieces of information. For example, the generic data may consist of upper and lower limits on the failure rate. Each of these limits

is then equated to its theoretical counterpart derived from the prior family considered. Since each theoretical expression will be a function of the two prior parameters, the two equations can be solved simultaneously for the values of the two parameters.

Example 1. Given that a diesel generator starts successfully, its subsequent hourly failure rate is given in the Reactor Safety Study as a lognormal distribution with 5th percentile $\lambda_L = 3 \times 10^{-4}$ and 95th percentile $\lambda_U = 3 \times 10^{-2}$. For the lognormal distribution we have the pair of equations given by

$$\Phi\left[\frac{\ln(3 \times 10^{-4}) - \xi}{\sigma}\right] = 0.05$$

and

$$\Phi\left[\frac{\ln(3 \times 10^{-2}) - \xi}{\sigma}\right] = 0.95$$

where ξ and σ are parameters of the lognormal family (Section 5.5.2.4.4) and $\Phi(\cdot)$ is the standard normal cumulative distribution function. Since $\Phi(-1.645) = 0.05$ and $\Phi(1.645) = 0.95$, we have

$$\ln(3 \times 10^{-4}) - \xi = -1.645\sigma$$

and

$$\ln(3 \times 10^{-2}) - \xi = 1.645\sigma$$

from which $\xi = -5.81$ and $\sigma = 1.40$. Thus, the fitted lognormal prior based on the RSS data becomes

$$f(\lambda) = \frac{1}{1.40\lambda\sqrt{(2\pi)}} \exp\left[-\frac{1}{2(1.40)^2}(\ln\lambda + 5.81)^2\right] \quad (0 < \lambda < \infty)$$

An alternative technique is considered in Section 5.5.2.2.8.

Similar techniques can be used for generic data like means or medians. However, if only a "best" point estimate is given (as in some of the IEEE Std-500 cases), there will usually be a need for some additional specification by the analyst. First, he must decide whether to use the mean, median, or mode of the distribution as the suitable central value representing the "best" estimate. Second, the analyst may have to introduce a second parameter value in order to define a distribution without ambiguity. For example, suppose one is to fit a gamma prior for a failure rate when the only available datum is the mean of the generic rate. Since the mean does not uniquely determine a gamma distribution, the variance could also be introduced and treated as an unspecified parameter.

Often the prior data from a single generic source are inconsistent in the sense that no common prior distribution can be fitted to the data.

There is no universally accepted method of rectifying such inconsistencies, but any of several approaches could be taken. One would be to take the set of all priors implied by the generic data and define some "most conservative" criterion to select a single prior from the set. Another would be to consider the entire set of priors as representing multiple sources of generic data and employ the procedures suggested in the discussion that follows.

Combining Multiple Sources

Often, multiple sources of generic prior data must be reduced to a single prior distribution that satisfactorily reflects and incorporates the views of each source. The multiple sources might be generic data from two or more studies (e.g., the RSS or IEEE Std-500) that report on the same generic component; they may consist of the opinions of several experts about the same component; or, as noted above, the multiple "sources" may consist of the set of unrectified priors obtained from a single inconsistent source.

Three procedures are suggested for forming a consensus prior distribution, although several methods are described in the literature (see for example, Eisenberg and Gale, 1959; Brown and Helmer, 1964; Winkler, 1968; Stone, 1961; Winkler and Cummings, 1972; De Groot, 1974; and Morris, 1974, 1977). For convenience, consider a failure-rate estimation as before. If each source provides both a point and an interval estimate, the first method is to pool (combine) the estimates by means of simple geometric averaging techniques:

$$\hat{\lambda} = \left(\prod_{i=1}^{n} \hat{\lambda}_i \right)^{1/n} \qquad (5\text{-}10)$$

This is equivalent in effect to forming the usual arithmetic average of failure rates described by their logarithms. This estimate implicitly assumes that the underlying sources are statistically independent and of equal importance. If the sources are unequal in their contribution to the consensus prior, a weighted geometric mean could be used with weights chosen to reflect the importance of each source.

Martz and Bryson (1982) have developed a classical statistical model for combining multiple sources of data. The resultant maximum-likelihood consensus point estimator is a weighted geometric mean of the individual estimates in which the weights are simple functions of the uncertainty bounds supplied by each data source. A corresponding consensus confidence-interval estimator is also provided. The maximum-likelihood point estimator of Martz and Bryson (1982) reduces to Equation 5-10 under two conditions: if each data source reports the exact same range of uncertainty, and if there is no location bias in the individual estimates.

The above pooling method was used to synthesize the opinions of some 200 experts in developing the IEEE Std-500 data base. Martz and Waller (1978) examined the effectiveness of this approach in a simulation and concluded that the method produced good point estimates; however, the combined interval estimates generally tended to be too narrow and thus had less than the desired assurance.

The second method yields a consensus prior that is generally more diffuse (spread out) than that obtained by the method just described. This method, discussed by Winkler (1968) and Stone (1961), is often referred to as the "mixture method." It involves fitting a suitable prior to each generic source and then combining the individual prior distributions by forming a mixture,

$$f(\lambda) = \sum_{i=1}^{n} w_i f_i(\lambda) \qquad (5\text{-}11)$$

The coefficients w_i are positive weights that sum to 1. Winkler (1968) suggests several methods for determining the weights. In the absence of any reason for preferring one source over another, the selection $w_i = n^{-1}$ is an obvious possibility. An interesting feature of this method is that it may yield a non-unimodal prior distribution. If such a mixture is used as a prior distribution, the corresponding posterior distribution from Equation 5-4 will also be a mixture of the individual (component) posterior distributions, namely,

$$f(\lambda|E) = \sum_{i=1}^{n} w_i f_i(\lambda\ E) \qquad (5\text{-}12)$$

where the new (updated) weights are

$$w_i' = \frac{w_i \int_0^\infty f_i(\lambda)\ L(E|\lambda)\ d\lambda}{\sum_{i=1}^{n} w_i \int_0^\infty f_i(\lambda)\ L(E|\lambda)\ d\lambda} \qquad (i = 1,2,\ldots,n) \qquad (5\text{-}13)$$

Since this method generally yields a more diffuse consensus prior than does geometric averaging, it provides more-conservative interval estimates. For this reason it is often preferred. However, it should be pointed out that the mixture method is computationally more difficult; numerical methods are frequently required for determining such quantities as the prior moments and percentiles.

A third method has been described by Kaplan (1981b) and earlier by Guttman (1970). This method, called a "two-stage" Bayesian procedure by Kaplan, uses a Bayesian procedure for forming the prior (stage 1) before combining the prior with the likelihood function (stage 2).

To describe the two-stage method, assume that the problem to be solved is to estimate the failure rate of machine S and express the degree of confidence in this failure rate given the following relevant information:

E_1: engineering knowledge of the design and construction of the machine

E_2: past performance of similar machines in similar applications

E_3: past performance of the specific machine in question

The information E_3 is of the format

$$E_3 = \langle h_s, T_s \rangle$$

that is, a doublet stating that machine S has failed h_s times in T_s years. This information is used in Bayes' theorem:

$$f(\lambda|E_1,E_2,E_3) = \frac{f(\lambda|E_1,E_2)L(E_3|\lambda_1,E_1,E_2)}{\int_0^\infty f(\lambda|E_1,E_2)L(E_3|\lambda_1,E_1,E_2)}$$

where $f(\lambda|E_1,E_2,E_3)$ is the posterior probability distribution for λ_S. This distribution expresses the final state of knowledge about λ_S in light of all the evidence E_1,E_2, and E_3. On the right, $f(\lambda|E_1,E_2)$ is the "prior" distribution representing the state of knowledge without information E_3 but including E_1 and E_2.

This use of Bayes' theorem to incorporate the specific evidence E_3 is a conventional application of Bayes' theorem and is the second stage of the two-stage approach. The first stage of the two-stage approach is aimed at determining the prior $f(\lambda|E_1,E_2)$, from the information E_2, which is of the form

$$E_2 = \{\langle h_1,T_1\rangle, \langle h_2,T_2\rangle,\ldots,\langle h_M,T_M\rangle\}$$

E_2 then is the set of doublets giving the operating experience of a set of M components deemed similar to that being analyzed.

To use E_2, this set of M components is thought of as a sample from an infinite population Q of similar components. Considering the whole of Q, there is a frequency distribution $\Phi(\lambda)$, where λ is the failure rate of a member of Q, such that $\Phi(\lambda)\,d\lambda$ is the fraction of the population with failure rates in the interval $d\lambda$. Kaplan denotes $\Phi(\lambda)$ as the "population variability curve" for the population Q.

If the population variability curve was known, it could be used as a prior, that is,

$$f(\lambda|E_1,E_2) = \Phi(\lambda)$$

Since $\Phi(\lambda)$ is now known, it is necessary to express what is known or can be inferred about $\Phi(\lambda)$ from the evidence E_2. For this purpose, consider the function $\Phi(\lambda)$ as being imbedded in a space of functions $\Phi(\lambda)$. Then a probability distribution, call it $f(\Phi|E_1,E_2)$ over this space F of functions exists, expressing knowledge of where, in F, Φ is located. For this purpose, Kaplan writes the "first-stage" application of Bayes' theorem in the form

$$f(\Phi|E_1,E_2) = \frac{f(\Phi|E_1)L(E_2|\Phi,E_1)}{\int_0^\infty f(\Phi|E_1)L(E_2|\Phi,E_1)}$$

Thus $f(\Phi|E_1,E_2)$ is the state of knowledge about Φ "posterior" to having the information E_3.

Once $f(\Phi|E_1,E_2)$ is known, then the desired prior $f(\lambda|E_1,E_2)$ to the second stage of the process is calculated from

$$f(\lambda|E_1,E_2) = \int_F f(\Phi|E_1,E_2)\, d\Phi$$

Kaplan (1981b) uses discretization techniques to find the population-variability curve. This can be illustrated by choosing a two-parameter family of lognormal* curves as follows:

$$\Phi_{ij}(\lambda) = \frac{1}{\sqrt{(2\pi)}\ \lambda\sigma_j} \exp\left\{-\frac{[\ln(\lambda/\mu_i)]^2}{2(\sigma_j)^2}\right\}$$

where the two parameters μ_i, σ_j range over a discrete "grid." Thus,

$$p(\Phi_{ij}|E_1,E_2) = \frac{p(\Phi_{ij}|E_1)\ p(E_2|\Phi_{ij},E_1)}{\sum_{i=1}^{I}\sum_{j=1}^{J} p(\Phi_{ij}|E_1)\ p(E_2|\Phi_{ij},E_1)}$$

and

$$p(E_2|\Phi_{ij},E_1) = \prod_{m=1}^{M}\left[\int_0^\infty \Phi_{ij}(\lambda)\ \frac{(\lambda T)_m^{K_m}}{K_m!} \exp(-\lambda T_m)\, d\lambda\right]$$

where M is the number of components with data K_m failures in T_m hours.

The prior $p(\Phi_{ij}|E_1)$ is the information that describes the grid of the parameters μ_i and σ_j. This is determined from experience, or it could be a noninformative prior.

A further simplification can be made by finding a "best estimate" for Φ, or the mean value for the distribution $p(\Phi_{ij}|E_1,E_2)$; that is,

$$\bar{\Phi}(\lambda) = \sum_{ij} \Phi_{ij}(\lambda)\ p(\Phi_{ij}|E_1,E_2)$$

This could then become the final prior for combining with the likelihood function from E_3.

---

*The choice of this family of lognormal curves should be regarded as illustrative. Any desired family of curves could be used, subject only to the requirement that somewhere in the family there would be at least one good approximation to the true variability curve Φ.

5.5.2.2.5 Using Expert Opinion

Expert opinion is often used for a prior probability distribution when other information is inadequate. If neither physical nor theoretical models are available and relative frequency is unavailable as well, subjective assessment is the only alternative for obtaining a probability. The practical feasibility of this alternative is supported not only by theoretical foundations that show judgments about uncertain events can be expressed as probabilities but also by practical assessment procedures. Holloway (1979) reviews the basis for these procedures and gives examples for several assessment approaches. The following summary of assessment procedures draws on his book. After this summary, well-known cautions and guidelines for interpreting and reviewing expert opinions are presented to highlight the care and caveats that must accompany the quantitative assessment.

However, the user of this guide should be cautioned against the indiscrete use of the methods described in this section. These techniques and results are not necessarily applicable to PRAs, which often treat extremely small probabilities of various events. More research is needed to determine the direct applicability of these methods and findings to PRAs. The user should be aware that the subjective estimates frequently used in PRAs can have large biases and errors.

Assessment Lotteries

An assessment lottery is a physical example of a random process. The uncertainty represented by the lottery must be easily recognized by the expert and have definite, objective probabilities. Such a lottery is the reference scale that measures an expert's degree of belief about the uncertain event. The operational definition for subjective probability, then, is the fraction of this reference uncertainty scale that makes an expert just indifferent between the assessment lottery and the feeling of uncertainty toward the event being assessed.

One example of assessment lotteries is an urn containing balls of different colors, some fraction being one color and the rest the other color. Drawing a ball at random from the urn is supposed to provide a visualization of an objective probability. Spetzler and Stael von Holstein (1975) developed and clinically tested another procedure that uses the spinning of a reference wheel as the assessment lottery. Their experience has shown that these probability wheels provide a strong visual image of an uncertain process.

Assessment Procedures

Two approaches to subjective probability assessment are in practical use, either the direct approach or the indirect approach. With the direct approach, the expert is asked to declare the probability number associated with the feeling of uncertainty for the occurrence of an event. With the indirect approach, an expert is asked to choose between a reference assessment lottery and the uncertain feeling (the degree of belief) in an opinion or judgment. Until an expert has shown an ability both to form a knowledgeable opinion and to assess, unaided, a probability for the degree of belief associated with that opinion, the indirect approach is preferred. The

well-known difficulties in obtaining useful subjective probability assessments are summarized below in the section entitled "Validity of Expert Opinion." These difficulties are magnified by inexperienced, unaided direct assessments. The references in that section give some experience comparing the two approaches.

The direct approach has the expert state a number that represents the assessment of the probability. Some studies have shown it possible for people to become better at assessing their own feelings of uncertainty as probabilities (see for example, Stael von Holstein, 1970; Lichtenstein et al. 1977). This improvement in direct assessment comes from specific training and guided practiced discipline rather than by trial and error. A good direct assessment comes from one who is both an experienced expert in what is known about a technical area (as well as how much is not known) and an experienced expert on how to express that judgment with little cognitive bias. This is an uncommon combination of expertise.

Assessment lotteries are used in the indirect approach to disclose the subjective probability. This external reference is used as a scale to measure the internal degree of belief an expert holds toward an opinion. Dividing between the expert and the assessors the responsibility to provide both a well-founded, knowledgeable judgment and an accurate representation of that judgment as a probability allows the use of expert opinion in PRAs. Most technical experts are not practiced, good probability assessors of themselves. Using the indirect approach improves the quality of expert opinion over that obtained by unaided, inexperienced direct assessment. Fischhoff et al. (1981) have shown that people qualified as technical experts are by no means qualified as probability assessors of that expertise.

Assessment Models

The representation used to model the uncertain event, either intuitively or formally, is a significant part of obtaining a good assessment. How the expert thinks about the problem of giving a judgment on the event likelihood should be recorded (see the discussion on "Recording Expert Opinion," page 5-44). It is this representation that fashions the eventual probability that is assessed. If disputes or questions arise in reviewing the quality of the expert opinion, a brief description of the thought model can focus the issue to a particular facet of that judgment.

Often, the expert is better able to provide a judgment by refining the event description into underlying events or factors. This formal assessment model can be subdivided until the expert finds it easy to examine each part, provide an opinion conditioned on each one, and review the formally computed probability of the original event for completeness and accuracy. This aid to assessment relieves an expert from making logical, or procedural, errors in combining the underlying knowledge. Reducing this source of error with the use of assessment models allows the assessor to focus on revealing a more subtle bias in the judgment.

Validity of Expert Opinion

The validity of a subjective assessment comes from two distinct parts: the knowledge content provided by the expert and the procedural process

provided by the assessor. If the expert is playing both of these roles, the distinction blurs, but it is still useful to describe the source of inaccuracies.

The content factor is evaluated from the credentials provided by the expert. Identifying who knows what and how much is a routine task for a professional community. Even for a recognized expert, a peer review can use the assessment model to judge whether or not all the significant factors were included in the expert's opinion. Inaccuracies, disputes, omissions, and limits to knowledge can then be examined to improve the accuracy of the substantive, or content, portion of the probability assessment.

The procedural process is more difficult to evaluate. The judgmental processes used by the expert, the effect the assessor has on expanding or limiting the formation of the expert's opinion, the effect of misunderstandings, and the natural cognitive limits on human information processes are all hidden factors in a practical assessment. Clinical studies, however, have examined these process factors that affect expert opinion. These studies provide a catalog of possible sources of inaccuracy due to bias and the extent of their effect.

It is well known that various biases may accompany the subjectively quantified assessments of an expert. For example, Alpert and Raiffa (unpublished work, 1969) found that experts often overestimate the degree of certainty of their estimates and claim too high a level of assurance. They observed that interval estimates for which 98-percent assurance was claimed tended in reality to have an assurance of about 70 percent (i.e., to include the correct value 70 percent of the time). Alternatively stated, interval estimates are often too narrow for the assurance level that is claimed. Tversky and Kahneman (1974) attribute such bias in part to the phenomenon of "anchoring": the expert tends to focus, or "anchor," on an initial guess and is reluctant to deviate too far from that guess in accounting for possible misjudgment. The results of such studies suggest that the assurance associated with expert-supplied interval estimates should be reduced from that claimed. For example, if a 90-percent interval estimate is solicited, then the interval could perhaps be considered to be an actual 70-percent interval in fitting a prior.

It is also well known that the manner chosen to encode (solicit) the subjective probabilities held by the expert is crucial and may significantly affect the quality of the information (see, for example, Du Charme and Donnell, 1973; Winkler, 1967; and Seaver et al., 1978). Spetzler and Stael von Holstein (1975) describe and recommend a structured-interview procedure and suggest a number of techniques for reducing biases in the quantification of judgment.

Holloway (1979) finds two findings from these studies encouraging. First, persons who are procedural experts in obtaining probability distributions are able, by using a variety of assessment techniques, to elicit consistent, well-founded judgments from substantive experts. Second, the substantive experts who are knowledgeable about the event being assessed are able to learn quickly about the significant procedural factors of probability assessment.

Recording Expert Opinion

The procedure used for assessing expert opinion and the assessment model used by the expert to construct the judgment should be described in a record of the expert opinion.

A subjective probability is an evaluation. The important procedural and substantive factors in that evaluation should be recorded, like any other engineering analysis, to permit a peer review to determine the quality of that result.

This record does not have a standard format; however, with time and experience, one may evolve. Nevertheless, the probability number can be meaningless without a description of how it was obtained and what its principal foundations were.

5.5.2.2.6 Beta Prior Distributions

The beta family of prior distributions is the conjugate family when failure-on-demand probabilities are estimated with a binomial likelihood function (Section 5.5.2.3). To fit a beta prior, values of the two prior beta parameters must be selected.

Martz and Waller (1982) present a table-lookup procedure, along with two sets of tables, that can be directly used to determine the beta-parameter values. Two situations are considered: (1) when the prior mean and 5th percentile of the prior distribution of failure-on-demand probabilities are specified and (2) when the prior mean and 95th percentile are specified. The procedure then yields directly the two beta parameters, as described by Martz and Waller with examples.

Mosleh and Apostolakis (1982) also describe a procedure for determining the beta-parameter values corresponding to various combinations of 5th, 50th, and 95th percentiles as well as the mean. Their procedure is to approximate the beta distribution as a gamma distribution and use corresponding techniques for determining the gamma parameters. Ahmed et al. (1981) have developed a computer code, called BURD, that finds the beta-parameter values corresponding to specified 5th and 95th percentile values.

5.5.2.2.7 Gamma Prior Distributions

The gamma family of prior distributions is the conjugate family when failure rates are estimated with a Poisson likelihood function (Section 5.5.2.4). The gamma family is a two-parameter family, and both parameter values must be identified by specifying some two conditions.

Martz and Waller (1982) present a simple procedure for determining the values of both parameters when two percentiles are given, corresponding to tail areas of 0.5, 1, 2.5, 5, 10, 25, 50, 75, 90, 95, 97.5, 99, or 99.5

percent. Mosleh and Apostolakis (1982) also present a procedure for determining the two gamma-parameter values for specified pairs of values--the (5th, 95th), (5th, 50th), (50th, 95th), (mean, 5th), or (mean, 95th). Ahmed et al. (1981) describe the use of the BURD code to determine the gamma-parameter values for specified 5th and 95th percentile values.

5.5.2.2.8 Lognormal Prior Distributions

The lognormal distribution is frequently used as a prior distribution for failure rates, especially when the failure rates typically encountered are so low (say, 10^{-6} per demand or per unit time) as to make a logarithmic transformation attractive. Apostolakis et al. (1980) make use of lognormal priors, as did the Reactor Safety Study. We consider here a simple procedure for determining the lognormal parameters ξ and σ (see Section 5.5.2.4).

Suppose that two symmetrically located percentiles are specified for the lognormal, denoted by λ_γ and $\lambda_{1-\gamma}$, where $0 < \gamma < 0.5$. Thus,

$$P(\lambda < \lambda_\gamma) = P(\lambda > \lambda_{1-\gamma}) = \gamma$$

The geometric mean of the percentiles is defined as

$$M = (\lambda_\gamma \lambda_{1-\gamma})^{1/2}$$

and a generalized error factor is

$$EF = (\lambda_{1-\gamma}/\lambda_\gamma)^{1/2}$$

Then the desired parameter values are

$$\xi = \ln M \quad \text{and} \quad \sigma = \ln EF/z_{1-\gamma} \tag{5-14}$$

where $z_{1-\gamma}$ is the $100(1 - \gamma)$th percentile of a standard normal distribution. In this case the mean, the variance, the mode, and the median of the fitted lognormal distribution can be found from the parameters as follows:

Mean: $\exp(\xi + \sigma^2/2)$

Mode: $\exp(\xi - \sigma^2)$

Median: $\exp(\xi) = M$

Variance: $[\exp(2\xi + \sigma^2)][\exp(\sigma^2) - 1]$

It is further observed that M is the median of the lognormal distribution and that the two percentiles are $\lambda_{1-\gamma} = (EF)(M)$ and $\lambda_{\gamma} = M/(EF)$, in accord with the notion of an error factor.

Example 2. On reconsidering Example 1, where $\lambda_{0.05} = 3 \times 10^{-4}$ and $\lambda_{0.095} = 3 \times 10^{-2}$, we find immediately that $M = 3 \times 10^{-3}$ and $EF = 10$. These are then substituted into Equation 5-14 to obtain $\xi = -5.81$ and $\sigma = 1.40$, for the latter making use of the fact that $z_{0.95} = 1.645$. Equations 5-15 give for the mean, mode, median, and variance the values 8×10^{-3}, 4×10^{-4}, 3×10^{-3}, and 4×10^{-4}, respectively.

Apostolakis et al. (1980) present a similar method for fitting a lognormal prior when, in addition to the two symmetric percentiles λ_{γ} and $\lambda_{1-\gamma}$, the median is also specified. Their method requires resolution of the evident inconsistency when the geometric mean of the upper and lower percentiles is not equal to the specified median.

5.5.2.3 Estimating Failure-on-Demand Probabilities

5.5.2.3.1 Binomial Likelihood Function

The binomial distribution is the distribution of the number of failures, r, out of n independent demands, on each of which the component has a constant failure-on-demand probability p. Given this statistical framework, the likelihood in Equation 5-4 is the binomial distribution, given by

$$L(E|p) = \frac{n!}{r!\ (n - r)!}\ p^{r}(1 - p)^{n-r} \qquad (5\text{-}16)$$

for $r = 0,1,2,\ldots,n$ and the parameter p between 0 and 1. If the parameter p is small (as usually happens in a PRA) and n is sufficiently large, then Equation 5-16 will usually be most conveniently approximated by the Poisson distribution, to be discussed in a slightly different context in Section 5.5.2.4:

$$L(E|p) = (np)^{r} \exp(-np)/r! \qquad (5\text{-}17)$$

where, because the number of demands is so large in comparison with the number of failures, r is treated as being able to assume any nonnegative integral value. The large values of r thus contribute negligibly to the probability distribution.

In the Bayesian approach, the parameter p is regarded as a random variable with a specified prior distribution. Returning now to the general binomial context, we consider three methods of generating a prior: (1) a noninformative prior; (2) a natural conjugate beta prior; and (3) a lognormal prior. The next three sections consider three priors, presenting in the interests of conciseness only the major results and formulas required to compute appropriate moments and estimates. Details can be found in the text by Martz and Waller (1982).

5.5.2.3.2 Noninformative Prior Distribution

One prior density is calculated from

$$[p(1 - p)]^{-0.5}/\pi \quad (0 \leq p \leq 1)$$

The prior mean, median, and variance are as follows:

Prior mean: 0.5

Prior median: 0.5

Prior variance: 0.125

and the prior 100(1 - γ)% symmetric probability interval is obtained from

$$\frac{0.5}{0.5 + 0.5F_{1-\gamma/2}(1,1)}; \frac{0.5F_{1-\gamma/2}(1,1)}{0.5 + 0.5F_{1-\gamma/2}(1,1)}$$

where $F_{1-\gamma}(a,b)$ is the 100(1 - γ)th percentile of an F-distribution with a and b degrees of freedom.

The posterior density, after r failures in n demands, is obtained from

$$\frac{\Gamma(n + 1)}{\Gamma(r + 0.5)\ \Gamma(n - r + 0.5)} p^{r-0.5}(1 - p)^{n-r-0.5} \quad (0 \leq p \leq 1)$$

and the formulas for calculating the posterior mean, median, and density are as follows:

Posterior mean: $(r + 0.5)/(n + 1)$

Posterior median: $$\frac{r + 0.5}{r + 0.5 + (n - r + 0.5)\ F_{0.5}(2n - 2r + 1,\ 2r + 1)}$$

Posterior variance: $$\frac{(r + 0.5)(n - r + 0.5)}{[(n + 1)^2\ (n + 2)]}$$

and the posterior 100(1 - γ)% symmetric probability interval is obtained from

$$\frac{r + 0.5}{r + 0.5 + (n - r + 0.5)\ F_{1-\gamma/2}(2n - 2r + 1,\ 2r + 1)};$$

$$\frac{(r + 0.5)\ F_{1-\gamma/2}(2r + 1,\ 2n - 2r + 1)}{n - r + 0.5 + (r + 0.5)\ F_{1-\gamma/2}(2r + 1,\ 2n - 2r + 1)}$$

5.5.2.3.3 Beta Prior Distribution

For the beta prior distribution, the prior density is obtained from

$$\frac{\Gamma(n_0)}{\Gamma(r_0)\,\Gamma(n_0 - r_0)}\, p^{r\hat{0}-1}(1 - p)^{n\hat{0}-r\hat{0}-1} \qquad (0 \leq p \leq 1)$$

where the positive values n_0 and r_0 are parameters of the beta distribution but may be interpreted as the numbers of demands and failures, respectively, in the prior data. The prior mean, median, and variance are calculated as follows:

Prior mean: r_0/n_0

Prior median: $$\frac{r_0}{r_0 + (n_0 - r_0)\, F_{0.5}(2n_0 - 2r_0,\ 2r_0)}$$

Prior variance: $$\frac{r_0(n_0 - r_0)}{n_0^2(n_0 + 1)}$$

and the formula for the prior $100(1 - \gamma)\%$ symmetric probability interval is

$$\frac{r_0}{r_0 + (n_0 - r_0)\, F_{1-\gamma/2}(2n_0 - 2r_0,\ 2r_0)};$$

$$\frac{r_0\, F_{1-\gamma/2}(2r_0,\ 2n_0 - 2r_0)}{n_0 - (r_0 + r_0)\, F_{1-\gamma/2}(2r_0,\ 2n_0 - 2r_0)}$$

The posterior density is given by

$$\frac{\Gamma(n + n_0)}{\Gamma(r + r_0)\,\Gamma(n - r + n_0 - r_0)}\, p^{r+r\hat{0}-1}(1 - p)^{n-r+n\hat{0}-r\hat{0}-1} \qquad (0 \leq p \leq 1)$$

and the other formulas are as follows:

Posterior mean:

$$(r + r_0)/(n + n_0)$$

Posterior median:

$$\frac{r + r_0}{r + r_0 + (n - r + n_0 - r_0)\, F_{0.5}(2n - 2r + 2n_0 - 2r_0,\ 2r_0 + 2r_0)}$$

Posterior variance:

$$\frac{(r + r_0)\,(n - r + n_0 - r_0)}{(n + n_0)^2\,(n + n_0 + 1)}$$

Posterior $100(1 - \gamma)\%$ symmetric probability interval:

$$\frac{r + r_0}{r + r_0 + (n - r + n_0 - r_0)\,F_{1-\gamma/2}(2n - 2r + 2n_0 - 2r_0,\ 2r + 2r_0)};$$

$$\frac{(r + r_0)\,F_{1-\gamma/2}(2r + 2r_0,\ 2n - 2r + 2n_0 - 2r_0)}{n - r + n_0 - r_0 + (r + r_0)\,F_{1-\gamma/2}(2r + 2r_0,\ 2n - 2r + 2n_0 - 2r_0)}$$

5.5.2.3.4 Lognormal Prior Distribution

The lognormal distribution is often used as a prior distribution on p, but its parameters must be so chosen that the probability density outside the actual range of p--that is, above the value p = 1--is sufficiently small to be ignored or effectively truncated. Apostolakis and Kaplan (1981) discuss the effect of such a truncation. As noted earlier, the lognormal was used as a prior in the Reactor Safety Study (USNRC, 1975) and in Apostolakis et al. (1980) as well as in other PRAs.

The prior density is obtained from the formula

$$\frac{1}{\sigma p\,\sqrt{(2\pi)}}\,\exp\left[-\frac{1}{2\sigma}\,(\ln p - \xi)^2\right] \quad (p > 0)$$

The prior moments, etc., are given in Section 5.5.2.2.8, and the prior $100(1 - \gamma)\%$ symmetric probability interval is calculated by using the following:

$$\left[\exp(\xi - z_{1-\gamma/2}\sigma);\ \exp(\xi + z_{1-\gamma/2}\sigma)\right]$$

The posterior distribution cannot be obtained in closed form. However, the approximation given in Equation 5-5 can be used to approximate the posterior distribution where $f(p_i)$ denotes the area under the lognormal prior over an interval represented by $p = p_i$ and $L(E\ p_i)$ denotes either Equation 5-16 or 5-17 evaluated at $p = p_i$ for the selected set of discrete values p_i $(i = 1,2,\ldots,m)$.

5.5.2.4 Estimating Constant Failure Rates

5.5.2.4.1 Poisson Likelihood Function

A common assumption in reliability models is that failure times are independent, with a common exponential (constant failure rate) distribution. It follows that the distribution of the number of failures r in a fixed total operating time T has a Poisson distribution. In this case the likelihood function that is defined in Equation 5-4 is the Poisson density given by the following:

$$L(E|\lambda) = (\lambda T)^r \exp(-\lambda T)/r! \quad (r = 0,1,2,\ldots)$$

where λ denotes the constant failure rate.

We consider three cases: (1) one noninformative prior distribution; (2) a natural conjugate gamma prior distribution; and (3) a lognormal prior distribution on λ.

5.5.2.4.2 Noninformative Prior Distribution

The various formulas for the noninformative prior distribution are as follows:

Prior density: $\lambda^{-0.5}$ (an improper distribution) $(\lambda > 0)$

Posterior density: $\dfrac{T^{r+0.5}}{\Gamma(r+0.5)} \lambda^{r-0.5} \exp(-\lambda T)$ $(\lambda > 0)$

Posterior mean: $(2r + 1)/(2T)$

Posterior median: $\chi^2_{0.5}(2r + 1)/(2T)$

where $\chi^2_{1-\gamma}(n)$ is the $100(1 - \gamma)$th percentile of a chi-square distribution with n degrees of freedom.

Posterior variance: $(2r + 1)/(2T^2)$

Posterior $100(1 - \gamma)\%$ symmetric probability interval:

$$\left[\chi^2_{\gamma/2}(2r + 1)/(2T);\ \chi^2_{1-\gamma/2}(2r + 1)/(2T)\right]$$

5.5.2.4.3 Gamma Prior Distribution

The prior density is obtained from

$$\frac{\beta_0^{\alpha_0}}{\Gamma(\alpha_0)} \lambda^{\alpha_0-1} \exp(-\beta_0\lambda) \quad (\lambda > 0)$$

where the positive shape parameter α_0 can be interpreted as the prior number of failures in β_0 prior total operating time. (β_0, also positive, is the scale parameter.)

The other formulas are as follows:

Prior mean: α_0/β_0

Prior median: $\chi^2_{0.5}(2\alpha_0)/(2\beta_0)$

Prior variance: α_0/β_0^2

Prior $100(1 - \gamma)\%$ symmetric probability interval:

$$[\chi^2_{\gamma/2}(2\alpha_0)/(2\beta_0);\ \chi^2_{1-\gamma/2}(2\alpha_0)/(2\beta_0)]$$

Posterior density:

$$\frac{(\beta_0 + T)^{\alpha_0+r}}{\Gamma(\alpha_0 + r)} \lambda^{\alpha_0+r-1} \exp[-(\beta_0 + T)\lambda] \quad (\lambda > 0)$$

Posterior mean: $(\alpha_0 + r)/(\beta_0 + T)$

Posterior median: $\chi^2_{0.5}(2\alpha_0 + 2r)/(2\beta_0 + 2T)$

Posterior variance: $(\alpha_0 + r)/(\beta_0 + T)^2$

Posterior $100(1 - \gamma)\%$ symmetric probability interval:

$$[\chi^2_{\gamma/2}(2\alpha_0 + 2r)/(2\beta_0 + 2T);\ \chi^2_{1-\gamma/2}(2\alpha_0 + 2r)/(2\beta_0 + 2T)]$$

5.5.2.4.4 Lognormal Prior Distribution

The prior density is obtained from

$$\frac{1}{\sigma\lambda\sqrt{(2\pi)}} \exp[-(\ln \lambda - \xi)^2/2\sigma)^2] \quad (\lambda > 0)$$

The prior moments, etc., are given in Section 5.5.2.2.8, and the prior $100(1 - \gamma)\%$ symmetric probability interval is calculated as follows:

$$\exp(\xi - z_{1-\gamma/2}\sigma);\ \exp(\xi + z_{1-\gamma/2}\sigma)$$

The posterior distribution cannot be obtained in closed form. However, the discrete approximation in Equation 5-5 can be used to approximate the posterior distribution, or numerical integration can be used in conjunction with Equation 5-4. There $f(\lambda_i)$ denotes the area under the lognormal prior in the vicinity of λ_i and $L(E|\lambda_i)$ denotes the likelihood (density function) above evaluated at the chosen discrete set of values λ_i $(i = 1,2,\ldots,m)$.

5.5.2.5 Example: Failure of Diesel Generators To Start

Presented below is an example from Apostolakis et al. (1980). The frequency with which diesel generators fail to start (measured in terms of the failure rate per demand) was assumed in the Reactor Safety Study to have alognormal distribution with 5th and 95th percentiles of 10^{-2} and 10^{-1}, respectively. Thus, using the procedure outlined in Section 5.5.2.2.8, we find that $\xi = 3.45$ and $\sigma = 0.70$ are the two lognormal parameter values. The prior mean, mode, median, and variance are then found to be 0.04, 1.9×10^{-2}, 3.2×10^{-2}, and 1×10^{-3}, respectively.

Suppose now that the evidence E from a certain plant consists of $r = 5$ failures in $n = 227$ test demands (see Section 5.5.2.3). Table 5-2 shows the discretized lognormal prior and calculations required to compute the corresponding posterior distribution by means of Equation 5-5; values smaller than 10^{-4} have been treated as equal to zero.

Figure 5-5 shows a plot of the discretized prior and posterior distributions and gives a graphic illustration of the change in the generic prior brought about by the influence of the plant-specific evidence. The posterior mean and variance are computed to be 0.025 and 8.2×10^{-5}, respectively. The effects of the plant-specific evidence are, first, to shift the distribution of the failure-to-start probability toward lower values and, second, to reduce the dispersion.

Another alternative Bayesian procedure is to approximate the binomial likelihood with a Poisson distribution (see Section 5.5.2.3.1) and to assign a conjugate gamma prior distribution to the corresponding failure rate. Taking the 5th and 95th percentiles to be 10^{-2} and 10^{-1}, respectively, and using the procedure of Martz and Waller (1982) (see Section 5.5.2.2.7) yields a gamma prior distribution with the shape parameter $\alpha_0 = 2.4$ and

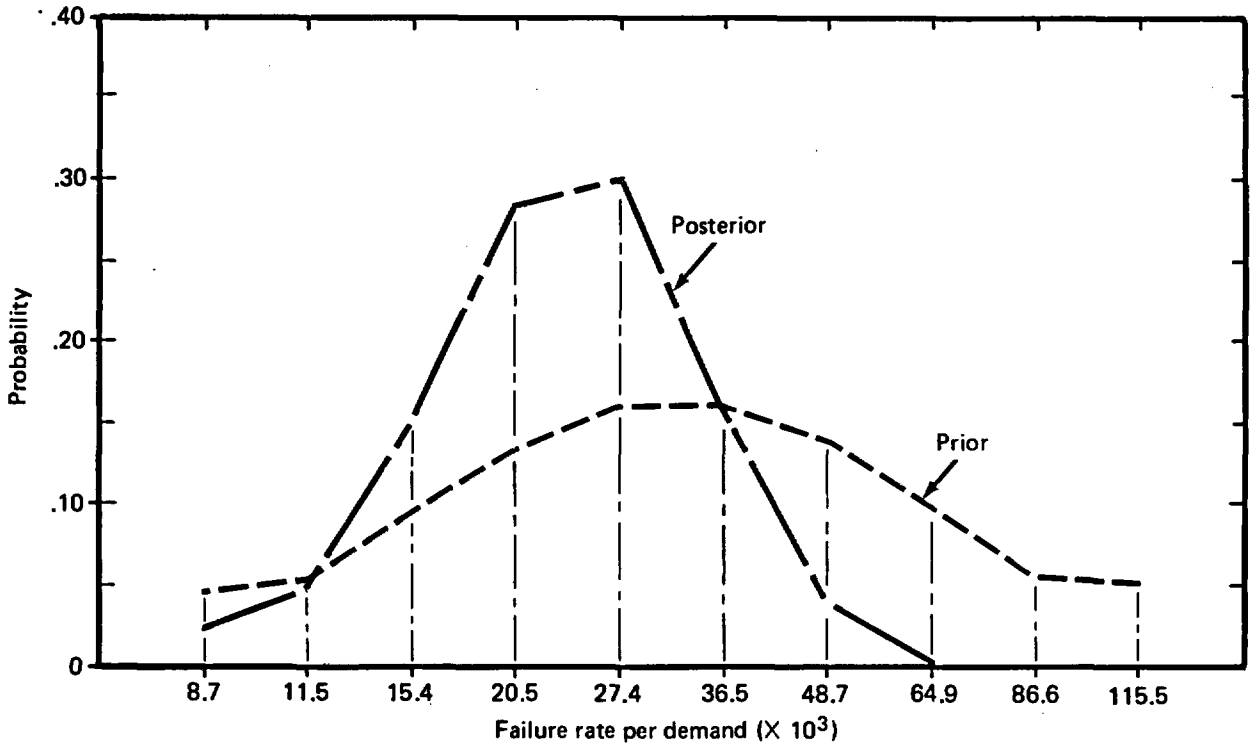

Figure 5-5. Prior and posterior histograms for diesel-generator failure to start. From Apostolakis et al. (1980).

the scale parameter $\beta_0 = 52.68$. Using the results in Section 5.5.2.4.3, the posterior distribution is another gamma distribution with the shape parameter 7.4 and the scale parameter 279.68. The corresponding posterior mean and variance are computed to be 0.026 and 9.5×10^{-5}, respectively. The posterior 5th, 50th, and 95th percentiles are also easily computed to be 0.013, 0.038, and 0.045, respectively.

Consider now the estimation of the probability of diesel-generator failure to start by the classical methods of Section 5.5.1. The data, f/n = 5/227, lead to a maximum-likelihood estimate of $p^* = .022$, which has a standard error of .0097. Note that the square of this standard error is 9.5×10^{-5}, which is slightly larger than the Apostolakis posterior variance. The difference in precision reflects the effect of the selected prior distribution.

Table 5-3 gives lower and upper classical confidence limits on the failure-to-start probability for a variety of confidence levels. It presents both the exact solutions to the expressions given in Section 5.5.1.3 and the chi-squared approximations. Both sets of confidence limits are shown to four decimals only to illustrate the close agreement between the exact and the approximate bounds for these data.

Because of the discretizing that is used, it is difficult to compare the Bayesian results in Table 5-2 with the classical results in Table 5-3. Qualitatively, however, both analyses suggest strongly that the failure probability of interest is between .01 and .05. As one method of

comparison, note that data of 7.5 failures in 300 demands would yield a maximum-likelihood estimate and a squared standard error essentially equal to Apostolakis' posterior mean and variance; thus, his prior effectively contributed additional data of 2.5/73 to his results.

In general, all three analyses of these data agree quite closely, even though the interpretation is quite different. The main reason for this agreement is the rather large quantity of plant-specific data, which results in a likelihood that dominates the prior distribution in the Bayesian analysis.

Table 5-2. Estimation of diesel-generator failure to start by the Bayesian method[a]

Failure rate (failure to start)	Prior probability	Likelihood	(Prior) x (likelihood)	Posterior probability
.0087	.0500	.0343	.0017	.0206
.0115	.0587	.0750	.0044	.0529
.0154	.0967	.1320	.0128	.1535
.0205	.1350	.1734	.0234	.2815
.0274	.1596	.1544	.0246	.2963
.0365	.1596	.0820	.0131	.1572
.0487	.1350	.0218	.0029	.0353
.0649	.0967	.0023	.0002	.0027
.0866	.0587	.0001	.0000	.0000
.1155	.0500	.0000	.0000	.0000
Sum	1.0000		.0831	1.0000

[a]From Apostolakis et al. (1980).

Table 5-3. Classical confidence limits on the probability of diesel-generator failure to start (Five failures in 227 attempts)

Confidence level (%)	Exact solution		Chi-squared approximation	
	Lower	Upper	Lower	Upper
50	.0205	.0249	.0206	.0249
75	.0149	.0325	.0148	.0327
90	.0108	.0405	.0107	.0407
95	.0087	.0458	.0087	.0463
97.5	.0072	.0507	.0072	.0513
99	.0057	.0567	.0056	.0577

5.6 EVALUATION OF DEPENDENT FAILURES

To support the analysis of dependent failures, which are discussed in detail in Section 3.7, appropriate data must be gathered. In gathering these data, it is necessary to establish what events will be classified as dependent and whether the beta-factor method or the binomial failure-rate (BFR) model will be used. An alternative approach is to use the various data reports by Atwood. These reports (Atwood, 1980a, 1982a,b; Atwood and Steverson, 1982a,b) include point estimates and confidence levels for the BFR model for a number of components at nuclear plants. Furthermore, the binomial failure rates can be used to estimate a beta factor, if desired. In addition, a computer code, BFR (Atwood and Suitt, 1982) is available to assist in the evaluation of data.

5.6.1 CLASSIFICATION OF EVENTS

A number of definitions have been used for the classification of events as dependent failures. Indeed, EPRI began a program in 1982 to refine the definition of such failures and thereby establish clearly which events involve dependences. The definition used here is consistent with Atwood's reports, but the data analyst may find it necessary to revise this definition for a particular study. For example, the analyst may wish to treat all multiple failures as if they were attributable to common causes, regardless of the mechanisms that caused the failure.

For this discussion, then, events that are simultaneous because of some external shock to the events are dependent. Two events occurring in the same time frame without such a shock are not considered to be dependent.

The data reports mentioned above (Atwood, 1982a,b; Atwood and Steverson, 1982a,b) give several examples of the classification of events, and these documents should be examined before the classification of specific data is begun.

5.6.2 CALCULATION OF PARAMETERS

The method presented here for the calculation of dependent-event parameters is that of Atwood and Steverson. Again, their documents should be consulted for additional detail and examples.

The quantities of interest are the following:

p	probability that a specific component fails, given that a shock occurs
m	number of components simultaneously susceptible to a shock
λ	failure rate for an individual component, not counting failures due to a common-cause shock

μ	rate of common-cause shocks
$\lambda_+ = \mu(1 - q^m)$	rate of shocks that cause at least one component failure--that is, rate of "visible" shocks (here $q = 1 - p$)
$r_1 = \lambda + \mu p$	rate at which a specific component fails, either because of individual failure or because of a common-cause shock
$r_k = \mu p^k$, $k \geq 2$	rate at which a specific set of k components fails simultaneously (because of a common-cause shock)
r_k/r_1	probability, given that a certain component has failed, that specific k components will also fail at the same time.

The quantities $r_1, r_2, \ldots$ are the relevant rates for fault-tree analysis. If a cut set of a fault tree involves k pumps, $k \geq 1$, then the relevant rate is r_k. The beta factor for any cut set can be estimated from the ratio r_k/r_1, where there are k elements in the cut set.*

The data set for any dependence must then be broken down such that the analyst is comfortable with including all the events as a single kind of shock. While this seems undesirable, the alternative requires obtaining multiple parameters for each shock from a data set that is probably small. Uncertainty methods should be used to allow for the variability in the parameters.

Basically, it is necessary to estimate the parameters p, λ, and μ. The analyst should refer to a report by Atwood (1980b) to estimate these parameters. The other parameters can be evaluated from p, λ, and μ.

5.7 UNCERTAINTIES

The data-development process, as presented herein, includes both classical and Bayesian viewpoints of uncertainty in parameter estimation. While these techniques treat, to some extent, the uncertainty that is related to the amount of data and the variability due to differences between data sources, there are other uncertainties that are not treated at all. This section briefly describes the potential sources of uncertainty and methods of judging their effects. In addition, Chapter 12 should be consulted for an overview of the treatment of uncertainty.

---

*Note that in Section 3.7 the beta factor is defined somewhat differently. For $k = 2$, these definitions are identical. When $k > 2$, the beta factor defined in Section 3.7 is a compromise among the various quantities r_k/r_1.

5.7.1 SOURCES OF UNCERTAINTY

Before discussing sources of uncertainty, it is important to remember what one may be uncertain about. This chapter has so far presented methods for estimating the following:

1. The failure rate of components.

2. The probability that components (or systems) fail on demand.

3. The probability that components (or systems) are unavailable because of testing or maintenance.

This estimation process involves the use of various models and estimates of the parameters in these models. Thus, there may be uncertainty in the models and/or the parameters.

Since the analyst first chooses a model for the data items, there is obviously some uncertainty in that selection, as no physical occurrence exactly fits a mathematical model. Next, there is uncertainty in the parameter of that model, even given that the model is correct. The sources for parameter uncertainty include (1) the amount of data, (2) the diversity of data sources, and (3) the accuracy of data sources.

5.7.2 PROCEDURES FOR TREATING MODELING UNCERTAINTIES

The first source of uncertainty mentioned above is that of model choice. The best way to determine the effect of this choice is to try another model--that is, perform a sensitivity assessment. The difference in the point estimate and confidence interval can then be reported. It is not expected that this will be an important contribution to uncertainty, and hence these extra evaluations need be done only for dominant events where the model does not seem to fit well.

5.7.3 PROCEDURES FOR TREATING PARAMETER UNCERTAINTIES

Uncertainty in the data parameters is already treated explicitly in the data process for certain sources by including uncertainty due to the amount of data. In addition, the data process can include differences between sources of data--that is, variability of an event's rate (or probability) of occurrence from one facility to another. In addition, the data process can be used to incorporate inaccuracies in the data sources. Of course, judgment is likely to enter into the process at this point. For example, in using data from licensee event reports, the number of demands is often estimated. Instead of treating this estimate as constant, the Bayesian approach could treat it as a random variate, while the classical approach could treat this value as a point estimate with error bounds.

5.8 DOCUMENTATION OF THE DATA BASE

An important aspect of developing the data for accident-sequence evaluation is to document the various steps of the process. This includes not only the final numbers but also the various assumptions and sources of information. The reader should be able to trace each data item from the fault tree or event tree back to the source, with each assumption and calculation apparent.

Documentation should include the output of the data process (i.e., the numbers used in quantification) and the general data base used in the PRA. These two types of documentation are discussed below.

5.8.1 DOCUMENTATION OF THE GENERAL DATA BASE

The general data base for the PRA includes all work from the source of data through the numerical results for the general types of events evaluated.

5.8.2 DOCUMENTATION OF DATA APPLIED TO EACH MODEL

The basic inputs to the task of accident-sequence quantification, and the outputs of the data process, are the numerical representations of each event. Forms like those shown in Figures 5-6 and 5-7 should be used to tie the specific events to the general data base.

Figure 5-6 is an example of a data table for hardware events. The first two columns, event name and description, come from the fault tree or the event tree. They give the alphanumeric code for an event and a brief description. The third column, the failure rate or probability of failure on demand, gives the data from the general data base for the type of event modeled. Note that the type of distribution and the parameters are included. The fault exposure time or mission time applies to events that occur as a function of time (either failure in time after a successful start or failure in time during standby). This time, then, is the length of time the component must survive to ensure success or the time between tests.

An example of tabular format for documenting test or maintenance acts is shown in Figure 5-7. The first column gives the event name as it appears in the fault tree or event tree. The second column is a brief description of the event. The third and fourth columns list the model used for act frequency and the model for the duration of the act. Note that these values could be average values, distributions, or point estimates with error factors. The fifth column contains a list of all the components included in the one act. For a test, this is often several components. This list helps to indicate the level in the tree where the act is modeled. Also included is a column for indicating the source of the information used to develop the act models.

BASIC EVENTS: HARDWARE						
Event name	Description	Failure rate or failure-on-demand probability	Fault exposure time or mission time (τ)	Data source	Quantification model	Comments
EVLV12	Valve fails to open	Lognormal 1×10^{-3} per demand Error factor = 3	NA	Reactor Safety Study	Distribution: lognormal 1×10^{-3} (3) mean: 1.3×10^{-3}	
EPM12F	Pump fails to start	Lognormal 1×10^{-3} per demand Error factor = 3	NA	Reactor Safety Study	Distribution: lognormal 1×10^{-3} (3) mean: 1.3×10^{-3}	
EPM12D	Pump discontinues running after start	Lognormal 3×10^{-5} per hour Error factor = 10	24 hr	Reactor Safety Study	Distribution: lognormal 7.2×10^{-4} (10) mean: 1.9×10^{-3}	
ECL12D	Clutch fails during mission	Lognormal 1×10^{-6} per hour Error factor = 20	24 hr	Reactor Safety Study	Distribution: lognormal 2.4×10^{-5} (20) mean: 1.3×10^{-4}	

Figure 5-6. Example of data table for hardware.

BASIC EVENTS: TEST AND MAINTENANCE ACTS							
Event name	Description	Frequency-of-act model	Duration-of-act model	Components in act block	Data source	Quantification model	Comments
EHPIMA	Maintenance of HPI leg A	1/3 month	Lognormal 4 hr Error factor = 1.5	Manual valve 11, MOV-12, pump	Plant data	Distribution: lognormal 1.8×10^{-3} (1.5) Point estimate: 1.9×10^{-3}	

Figure 5-7. Example of data table for test or maintenance acts.

The most important column in the tables is the quantification model. This column is the output of the data section and the input to sequence quantification. It includes the distribution and mean (or point estimate and interval estimates) for each specific event. Note that for time-dependent events it is a function of τ and the failure rate (see Section 5.5).

5.9 ASSURANCE OF TECHNICAL QUALITY

The term "assurance of technical quality," as used here, refers only to the quality of the data base that results from the procedures given in this chapter. Many factors affect the quality of the data base, including the overall programming, planning, and scheduling, as well as budget limitations; such items are discussed in Chapter 2, Section 2.3.3. The objective of this section is to address the items that will enhance the data quality within the program constraints.

The most beneficial activities to maximize quality are reviews and checks. As each data quantity is produced, it should be checked against other data bases. Major discrepancies should be justified. Other staff members should review the event quantifications for their models and cross-compare with others with the same type of events. Finally, the team leader should review the data, using his experience to look for unusual results. Of course, outside peer review is an important part of the review process, though feedback for revision via this path usually takes longer than does feedback within the study.

Documentation is the key to the quality of the data base. The data analyst should keep a notebook to document his decisions and assumptions. This notebook will make final documentation easier and make the data traceable from event results back to the source. It is also important to carefully document computer runs so that, if necessary, the runs producing particular results can be found. Often a keypunch error can result in an incorrect result.

REFERENCES

Ahmed, S., D. R. Metcalf, R. E. Clark, and J. A. Jacobsen, 1981. BURD--A Computer Program for Bayesian Updating of Reliability Data, NPGD-TM-582, Babcock & Wilcox, Lynchburg, Va.

Apostolakis, G., and S. Kaplan, 1981. "Pitfalls in Risk Calculations," Reliability Engineering, Vol. 2, pp. 135-145.

Apostolakis, G., S. Kaplan, B. J. Garrick, and R. J. Duphily, 1980. "Data Specialization for Plant-Specific Risk Studies," Nuclear Engineering and Design, Vol. 56, pp. 321-329.

Apostolakis, G., and A. Mosleh, 1979. "Expert Opinion and Statistical Evidence: An Application to Reactor Core Melt Frequency," Nuclear Science and Engineering, Vol. 70, pp. 135-149.

Atwood, C. L., 1980a. Common Cause and Individual Failure and Fault Rates for Licensee Event Reports of Pumps at U.S. Commercial Nuclear Power Plants, draft, EGG-EA-5289, EG&G Idaho, Inc., Idaho Falls, Idaho.

Atwood, C. L., 1980b. Estimators for the Binomial Failure Rate Common Cause Model, USNRC Report NUREG/CR-1401.

Atwood, C. L., 1982a. Common Cause Fault Rates for Pumps: Estimates Based on Licensee Event Reports at U.S. Commercial Nuclear Power Plants, January 1972-September 1980, USNRC Report NUREG/CR-2098.

Atwood, C. L., 1982b. Common Cause Fault Rates for Instrumentation and Control Assemblies: Estimates Based on Licensee Event Reports at U.S. Commercial Nuclear Power Plants, 1976-1978, USNRC Report NUREG/CR-2771.

Atwood, C. L., and J. A. Steverson, 1982a. Common Cause Fault Rates for Diesel Generators: Estimates Based on Licensee Event Reports at U.S. Nuclear Power Plants, 1976-1978, USNRC Report NUREG/CR-2099.

Atwood C. L., and J. A. Steverson, 1982b. Common Cause Fault Rates for Valves: Estimates Based on Licensee Event Reports at U.S. Commercial Nuclear Power Plants, 1976-1980, USNRC Report NUREG/CR-2770.

Atwood, C. L., and W. J. Suitt, 1982. User's Guide to BFR, A Computer Code Based on the Binomial Failure Rate Common Cause Model, USNRC Report NUREG/CR-2729.

Barlow, R. E., and F. Proschan, 1975. Statistical Theory of Reliability and Life Testing, Holt, Rinehart and Winston, Inc., New York.

Bayes, T., 1958. "Essay Toward Solving a Problem in the Doctrine of Chances" (reprinted), Biometrika, Vol. 45, pp. 293-315.

Brown, B., and O. Helmer, 1964. Improving the Reliability of Estimates Obtained from a Consensus of Experts, P-2986, the Rand Corporation, Santa Monica, Calif.

Chhikara, R. S., and J. L. Folks, 1977. "The Inverse Gaussian Distribution as a Lifetime Model," Technometrics, Vol. 19, pp. 461-468.

De Groot, M. H., 1974. "Reaching a Consensus," Journal of the American Statistical Association, Vol. 69, pp. 118-121.

Du Charme, W. M., and M. L. Donnell, 1973. "Intrasubject Comparison of Four Response Modes for 'Subjective Probability' Assessment," Organizational Behavior and Human Performance, Vol. 10, pp. 108-117.

Eisenberg, E., and D. Gale, 1959. "Consensus of Subjective Probabilities: the Pari-Mutuel Method," Annals of Mathematical Statistics, Vol. 30, pp. 165-168.

Fischhoff, B., P. Slovic, and S. Lichtenstein, 1981. "Lay Foibles and Expert Fables in Judgments About Risks," in R. O'Riordan and R. K. Turner (eds.), Progress in Resource Management and Environmental Planning, Vol. 3, John Wiley & Sons, Chichester, England.

Green, A. E., and A. J. Bourne, 1972. Reliability Technology, Wiley-Interscience, New York.

Guttman, I., 1970. "Tolerance Regions: Classical and Bayesian," Griffin Statistical Monographs, Griffin, London, England.

Hahn, G. J., and S. S. Shapiro, 1967. Statistical Models in Engineering, John Wiley & Sons, Inc., New York, Chapter 8.

Hald, A., 1952. Statistical Theory with Engineering Applications, John Wiley & Sons, Inc., New York.

Holloway, C. A., 1979. Decision Making Under Uncertainty: Models and Choices, Prentice Hall, Englewood Cliffs, N.J.

Jeffreys, H., 1961. Theory of Probability, 3rd ed., Clarendon Press, Oxford, England.

Kaplan, S., 1981a. "On a Two-Stage Bayesian Procedure for Determining Failure Rates from Experiential Data," IEEE Transactions on Power Apparatus and Systems (preprint).

Kaplan, S., 1981b. "On the Method of Discrete Probability Distributions in Risk and Reliability Calculations--Applications to Seismic Risk Assessment," Risk Analysis, Vol. 1, No. 3.

Lapides, M. E., and E. L. Zebroski, 1975. Use of Nuclear Plant Operating Experience To Guide Productivity Improvement Programs, EPRI SR-26-R, Electric Power Research Institute, Palo Alto, Calif.

Lichtenstein, S., B. Fischhoff, and L. D. Phillips, 1977. "Calibration of Probabilities: The State of the Art," in H. Jungermann and G. DeZeeuw (eds.), Decision Making and Chance in Human Affairs, D. Reidel, Amsterdam, the Netherlands.

Mann, N. R., R. E. Shafer, N. D. Singpurwalla, 1974. Methods for Statistical Analysis of Reliability and Life Data, John Wiley & Sons, Inc., New York.

Martz, H. F., and M. Bryson, 1982. "On Combining Data for Estimating the Frequency of Low-Probability Events with Application to Sodium Valve Failure Rates," to be published in Nuclear Science and Engineering.

Martz, H. F., and R. Waller, 1978. An Exploratory Comparison of Methods for Combining Failure-Rate Data from Different Data Sources, LA-7556-MS, Los Alamos National Laboratory, Los Alamos, N.M.

Martz, H. F., and R. Waller, 1982. Bayesian Reliability Analysis, John Wiley & Sons, New York.

McClymont, A., and G. McLagan, 1982. Diesel Generator Reliability at Nuclear Power Plants: Data and Preliminary Analysis, EPRI NP-2433, Electric Power Research Institute, Palo Alto, Calif.

Mosleh, A., and G. Apostolakis, 1982. "Some Properties of Distributions Useful in the Study of Rare Events," to be published in IEEE Transactions on Reliability.

Morris, P. A., 1974. "Decision Analysis Expert Use," Management Science, Vol. 20, pp. 1233-1241.

Morris, P. A., 1977. "Combining Expert Judgment: A Bayesian Approach," Management Science, Vol. 23, pp. 671-693.

Murphy, J., 1980. Component Failure Rates for Nuclear Plant Safety System Reliability Analysis, draft report, U.S. Nuclear Regulatory Commission, Washington, D.C.

Parry, T. W., and P. W. Winter, 1981. "Characterization and Evaluation of Uncertainty in Probabilistic Risk Analysis," Nuclear Safety, Vol. 22, pp. 28-42.

Seaver, D. A., D. V. Winterfeldt, and W. Edwards, 1978. "Eliciting Subjective Probability Distributions on Continuous Variables," Journal of Organizational Behavior and Human Performance, Vol. 21, pp. 379-391.

Spetzler, C. S., and C. A. S. Stael von Holstein, 1975. "Probability Encoding in Decision Analysis," Management Science, Vol. 22, pp. 340-358.

Stael von Holstein, C. A. S., 1970. Assessment and Evaluation of Subjective Probability Distributions, the Economic Research Institute, Stockholm School of Economics, Stockholm, Sweden.

Stone, M., 1961. "The Opinion Pool," Annals of Mathematical Statistics, Vol. 32, pp. 1339-1342.

Tversky, A., and D. Kahneman, 1974. "Judgment Under Uncertainty: Heuristics and Biases," Science, Vol. 185, pp. 1124-1131.

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

Winkler, R. L., 1967. "The Assessment of Prior Distributions in Bayesian Analysis," Journal of the American Statistical Association, Vol. 62, pp. 776-800.

Winkler, R. L., 1968. "The Consensus of Subjective Probability Distributions," Management Science, Vol. 15, pp. B61-B75.

Winkler, R. L., and L. L. Cummings, 1972. "On the Choice of a Consensus Distribution in Bayesian Analysis," Organizational Behavior and Human Performance, Vol. 7, pp. 63-76.