NUREG/CR-2040
BNL-NUREG-51367

A Study of the Implications of Applying Quantitative Risk Criteria in the Licensing of Nuclear Power Plants in the United States

Prepared by S. Mitra, R. Hall/Brookhaven National Laboratory
A. Coppola/Grumman Aerospace Corporation

Brookhaven National Laboratory

Prepared for
U.S. Nuclear Regulatory Commission

Available from

GPO Sales Program
Division of Technical Information and Document Control
U. S. Nuclear Regulatory Commission
Washington, D. C. 20555

Printed copy price: $6.00

and

National Technical Information Service
Springfield, Virginia 22161

NUREG/CR-2040
BNL-NUREG-51367
RG, XA

A Study of the Implications of Applying Quantitative Risk Criteria in the Licensing of Nuclear Power Plants in the United States

Manuscript Completed: March 1981
Date Published: May 1981

Prepared by
S. Mitra, R. Hall/Brookhaven National Laboratory
A. Coppola/Grumman Aerospace Corporation

Brookhaven National Laboratory
Upton, NY 11973

Prepared for
Division of Systems and Reliability Research
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, D.C. 20555
NRC FIN A3222

CONTENTS

CONTENTS (Cont'd)

CONTENTS (Cont'd.)

CONTENTS (Cont'd.)

LIST OF FIGURES

LIST OF FIGURES (Cont'd.)

LIST OF TABLES

1. INTRODUCTION

This report describes the results of an investigation into the feasibility of developing and using a set of probabilistic risk criteria to help judge the safety of nuclear power plants. The principal aim of this report was to critically review and examine the implications and ramifications of the various proposals for a numerical risk criterion from a unified viewpoint. Brookhaven National Laboratory (BNL) performed the investigation for the Methodology and Data Branch of the U.S. Nuclear Regulatory Commission (NRC).

The desire for a set of numerical risk criteria may be recognized and appreciated when viewed against the background of the evolution of safety practices that has prevailed since the inception of the commercial nuclear power program in the U.S. The first formalized approach to nuclear safety has been the use of the Maximum Credible Accident (MCA). The MCA is defined as the postulated credible accident which poses a potential hazard greater than any other accident which is also considered to be credible. Engineering judgement relying on an intuitive informal estimate of probabilities was used to divide accidents into a 'credible' or an 'incredible' category. It was then necessary to show that a plant met the guidelines set forth in the Code of Federal Regulations, Title 10, Part 100 (10 CFR 100). Although 'incredible' accidents were dismissed from further consideration, an attempt was always made to incorporate a margin of safety into the safety system design in an attempt to protect against such accidents. The major weakness of the MCA approach lies in the intuitive basis for the classification of accidents. It might well be that accidents deemed to be incredible by virtue of their low frequency of occurrence may indeed have a relatively high frequency of occurrence when systematically analyzed using formal probabilistic approaches. Also, even though the "incredible" accidents may have low frequency, they may pose a greater risk because of their proportionately higher consequences. Another weakness of the MCA approach is the failure to quantify the safety built into a design.

In an attempt to improve the MCA approach, a method known as the Design Base Accident (DBA) approach evolved. The DBA approach principally consists of explicitly identifying low frequency high consequence accidents which must be designed against. Though subjective judgement on what is credible and what is incredible is avoided, subjectivity and intuitive judgements still prevail in the consideration of what accidents are to be included in the design base; there were still no systematic, formal probabilistic analyses to determine the frequencies and consequences of the accidents. In the DBA approach, safety is assumed when it is shown that plants are capable of withstanding the DBAs. In addition to its lack of systematic probabilistic analysis, the DBA concept has been criticized on the grounds that it does not adequately motivate vendors and utilities to improve reactor safety through design changes because of the fixed nature of the designs acceptable in the DBA approach.

The MCA and DBA are consequence oriented approaches. They concentrate on the consequence aspect of accidents; relatively little attention is paid towards the quantification of the frequency of an accident or the relative frequency of accidents in alternate designs. The weakness of the consequences oriented methods have led recently to the transition to a new approach known as Probabilistic Risk Assessment (PRA). The PRA technique attempts to quantify both the frequency as well as the resulting consequences of accidents. The Reactor Safety Study (WASH-1400) marks the first major step in the application of PRA techniques to nuclear power plant safety. The Reactor Safety Study, however, had weaknesses and was heavily criticized. The criticism was not, in general, directed toward the risk approaches used but the specific data and assumptions used.

The NRC's present statutory mandate in licensing nuclear power plants is basically inherited from the days of the U.S. Atomic Energy Commission (AEC) which preceeded the NRC. The mandate is very general, it calls for providing adequate protection to the health and safety of the public. An operating license is issued based on the finding that there is a reasonable assurance that the authorized activities can be conducted without endangering the health and safety to the public. The courts of law have provided considerable discretionary power to the NRC to define what the words 'adequate' and 'reasonable' mean. As yet, however, the NRC has not quantified in terms of probability and risk, the definition as to what constitutes 'reasonable' and 'adequate.'

In May 1979, the Advisory Committee on Reactor Safeguards (ACRS) recommended[1] that "consideration be given by the NRC to the establishment of quantitative safety goals of nuclear power reactors..." and that "Congress should be asked to express its views on the suitability of such goals and criteria in relation to other relevant aspects of our technological society." A report[2], directed by Rogovin, on the accident at TMI to the NRC's Commissioners and to the public, offered the following suggestion:

> "We do not suggest here that the existing safety review process be immediately supplanted by a more probabilistic review. What we are suggesting is that it be augmented, and that quantitative methods be used as the best available guide to which accidents are the important ones, and which approaches are best for reducing their probability or their consequences.
>
> We believe that the advantages of such an approach far outweigh the difficulties. We strongly urge that NRC begin, the long and perhaps painful process of, converting as much as is feasible of the present review process to a more accident-sequence-oriented approach. This conversion process may be difficult. It could easily take as much as a decade to accomplish. The time to begin is now."

Also the President's Commission on the accident[3] at TMI recommended that "continuing in-depth studies should be initiated on the probabilities and consequences (on-site and off-site) of nuclear power plant accidents, including the consequences of meltdown."

In response to these recommendations to expand risk analyses and to consider numerical risk criteria, the division of Systems Reliability Research in the NRC instituted research and development programs to more widely apply risk analyses and to evaluate implications of numerical risk criteria. As part of these programs, this work and this report specifically addresses the implications and ramifications of different proposed numerical criteria for various safety aspects of a nuclear power plant. These investigations particularly address the difficulties inherent in the formulation and implementation of practical quantitative risk criteria. The difficulties associated with formulating and implementing quantitative risk criteria stem from the fact that it is exceedingly difficult, if not perhaps impossible, to develop a workable, defensible set of quantitative risk criteria which allows for benefits, societal needs and equity for all parties concerned. The difficulty is compounded because numerical risk criteria are not only a technical matter but by their

very nature involve socio-economic, legal, and political concerns. This report, while mentioning the various other aspects when applicable, limits itself to technical aspects of risk criteria.

Despite the inherent difficulties, several risk criteria have been proposed and are examined in this report. They include the following, listed with their generic category and the chapter which discusses them:

- System unavailability criteria (Chapter 5), (eg. Canadian Reactor Siting Guide),
- Accident Probability Criteria (Chapter 6), (eg., Burns, Wall, Vesely),
- Release Criteria (Chapter 7), (eg. Farmer),
- Individual Risk Criteria (Chapter 8), (eg. Canadian Reactor Siting Guide),
- Societal Risk Criteria (Chapter 9), (eg. Canadian Reactor Siting Guide, Kinchin, Levine).

For each of the generic categories, the main features are described, their implications examined, and the inherent difficulties are discussed. The examples cited are used for illustrative purposes only, and the numerical values presented are not given any special significance. In order to relate the numerous criteria that can be formulated, a hierarchy was developed to provide a basic framework for this report. The report is divided into chapters which are described in the following paragraphs.

Chapter 2 discusses pros and cons of criteria that determine unacceptability as opposed to those that determine acceptability. The need for consistent implementation for any criteria is highlighted. Some approaches useful in formulating criteria are discussed including risk comparison, revealed preference, expressed preference, cost-benefit, and multi-attribute theory. Again, the advantages and disadvantages of these approaches are given.

Chapter 3 establishes a hierarchy of risk levels (1 to 5) for which various criteria can be formulated. At the first level is the risk number criteria which focus on individual and societal risks. Level 2 is the release criteria which concentrate on the frequencies and consequences of radioactive releases from nuclear power plants. Level 3 is the accident

probability criteria that attempt to control the frequencies of some defined accidents such as those which involve core melt. Level 4 is the system availability criteria which focus on the availabilities for process and safety systems. Level 5 is the component availability criteria which focus on hardware failures and human errors. This hierarchy provides the framework for the next five chapters, although in reverse order, since the top levels depend on the lower levels for their development.

Chapter 4, while brief, is nonetheless important and placed at the beginning of our technical discussion of all criteria. Since there are no examples given of criteria based on component availability, data and models required in the evaluation of component availability are also required in the evaluations of system availability and accident probability. Under this heading, human error or man-machine interface is also included, since the human is considered a component in the system.

Chapter 5 discusses criteria based on system availabiity and the problems with this approach. It describes both the U.S. and Canadian practices (defense in depth) in regards to system design requirements and describes some of the problems with system fault tree analysis and the data required. Human error and common mode failure contributions to system unavailability are also discussed.

Chapter 6 discusses criteria based on various accident categories. As an example, two types of core melt probability criteria are examined. One limits the frequency of core melt accidents for a given reactor in any given year and the other limits the frequency of these accidents from all reactors in any given year. Core melt probabilities are examined in detail and the implications of particular probability numbers or ranges are presented. Several recent proposed criteria are described as examples.

The criteria described in Chapter 7 deal with the frequencies of various amounts of different isotopes that may be released from nuclear power plants under accident conditions. The properties and implications of this kind of criteria are discussed. Three different forms of the criteria are described, and the Farmer proposal is presented as an example of one form which seeks to constrain both the amount and frequency of any release.

In Chapter 8, two forms of individual risk criteria are examined. One is the limit on radioactive dose to an individual per year and the other is a limit on the annual risk of death for an individual due to nuclear reactor accidents. The individual risk of death is also presented in terms of life shortening. Properties and implications of these criteria are discussed and an example is given (Canadian Safety Requirements).

In Chapter 9 two types of societal risk criteria are discussed. One attempts to control the site-specific societal risk that is associated with a plant and the other attempts to control the total societal risk from all reactors. They include a limit on expected consequences per unit time, limits on frequencies per unit time of events exceeding certain consequences, and limits of consequences and their associated magnitudes. Again, the Canadian Reactor-Siting Guide total population dose limits are given as an example of these criteria, and the proposals of Kinchin in the U.K. and Levine in the U.S. are also discussed.

Chapter 10 discusses the various elements of costs incurred as a result of nuclear accidents, and the present status of insurance and liability systems in use in the U.S. Comparison of health and property damage costs are discussed, as are cost-benefit analyses. Several forms of property damage criteria are described and their implications are presented.

Chapter 11 deals with the difficulties of risk analysis in general and with the information required for their evaluation. The information required is divided into two broad categories, one being the models required and the other is all the data and descriptive information required for accomplishing any valid risk assessment. Since this report is limited to a technological discussion, problems with socio-psychological aspects such as risk perception are not discussed. The chapter is quite detailed, however, on the uncertainties inherent in any risk assessment due to the models and data used.

Chapter 12 addresses such questions as: given risk criteria in any form, what are the regulatory procedures required for their implementation? How is the review of a risk assessment accomplished? It describes the existing problems in the area of probabilistic risk assessment and examines the

kind of standards required to induce a consistent practice within the profession. Several proposals for assisting in this effort are described including a Science Court for settling disputes, a system of certification of individual risk analysts to insure top-quality workmanship in analyses, and the certification of risk analysis studies, which may involve a review by a team of specialists derived from all interested parties (e.g., NRC, industry, and others).

NOTE: The reader should consult the particular chapter cited above for detailed references. Each chapter has a separate list of references.

Chapter 1 References

1. Carbon, M.W., ACRS Chairman, in a letter to J.M. Hendrie, NRC Chairman, May 16, 1979

2. Rogovin, M., et al, "Three Mile Island, A Report to the Commissioners and to the Public", Nuclear Regulatory Commission Special Inquiry Group, January 1980

3. Kemeny, J.G., et al, "Report of the President's Commission on the Accident at Three Mile Island", October, 1979

2. FACTORS AND APPROACHES IN THE FORMULATION OF NUMERICAL RISK CRITERIA

2.1 INTRODUCTION

In Chapter 1 we outlined the historical process which has led to a desire for numerical criteria for reactor safety. In truth, the previous criteria mentioned, including the MCA (maximum credible accident) and the DBA (design basis accident) had implicit numerical values associated with them. Since they are however limited in scope, the range of adequacy is questionable. Their biggest drawback is perhaps their lack of completeness; all possible accidents are not considered and core melt accidents are not addressed.

It is the need for completeness which makes new, all inclusive criteria desirable. It is also desirable that the new criteria be unambiguous and easily used by regulatory and decision-making authorities. This leads to the presently suggested numerical criteria under consideration by this report, with explicit values used to determine adequacy of any reactor. The criteria considered would also cover all aspects of reactor design, siting, construction and operation. Given a numerical criterion, a probabalistic risk assessment would be required to determine whether or not the criterion had been met. The numerical expression of the risk inferred by the assessment would be compared with the criterion to determine what action should be taken. The possible alternatives might include not only approval or disapproval but a conditional approval as well.

Having established a desire for numerical criteria, the form and content of these criteria must be established. The factors to be considered in attempting to devise numerical criteria include:

- Should the criteria be used to gauge acceptability or unacceptability
- What methods are required for consistency in implementing the criteria
- What are the available approaches or suggested approaches to establishing such a criteria

These factors will be considered in this section. Other factors, which are discussed in later sections include:

- How should criteria be formulated
- What method should be employed to include uncertainties
- Should criteria cover hardware, human error or both
- Should the criteria incorporate risk aversion
- How is compliance to be demonstrated
- What data models and calculations would be required to judge compliance
- In applying criteria, what actions and decisions should be taken and in what time period should they occur

2.2 ACCEPTABLE VS. UNACCEPTABLE

If a numerical criterion is established which, when met, requires no further proof of acceptability, then it is termed an acceptablity criterion If, on the other hand, a numerical criterion is established which requires additional proof of acceptability even though the criterion is met, then it is termed an unacceptability criterion. For an acceptable criterion, if the inferred risk is higher than the criterion, then action must be taken to reduce the risk; if the inferred risk is lower than the criterion, then the risk is deemed acceptable and no further action need be taken for approval. For an unacceptability criterion, if the inferred risk is higher than the criterion, then action must also be taken; however, if the inferred risk is lower, then additional tests or regulatory requirements must still be satisfied. Thus, it is necessary, but not sufficient to satisfy an unacceptability criterion.

The determination of whether a formulated criterion should be termed an acceptable or unacceptable limit depends to a great extent on the completeness it represents for meeting society's safety goals. To illustrate this concept of completeness, and acceptability versus unacceptability, let us take as a simple example from industry, a test performed in the manufacture of ball bearings. At the end of the manufacturing process, as a test, each ball is passed through sizing holes to choose those which meet the specified diameter. This sizing test is complete in itself and is termed a "no-go" or unacceptability test since any bearing that does not pass is rejected, but those that do pass

are subject to further tests, including roundness, hardness, finish, etc. Passing this unacceptability test is thus complying with an unacceptability criterion. Now, if a comprehensive test were devised which would incorporate tests for all the attributes of interest, the resultant test could then be termed an acceptance test.

In the case of a numerical risk criterion, it is theoretically possible to have an acceptability criteria if the models and data utilized in risk assessment are accurate enough and comprehensive enough to encompass all attributes of the risk involved. Oftentimes, however, models and data have large uncertainties. If data and models have large uncertainties, then a criterion can only be used as one tool in assessing the acceptability of a system, plant or action. This is true whether we are speaking of one criterion or a set of criteria. The preference for numerical criteria being unacceptability criteria is particularly true of complex technologies such as nuclear power which are also relatively new. Most of the existing reliability data used in assessment is based on similar components and equipment in other industries. The system and accident models (fault trees and event trees) are also, by their very nature, not exhaustive and are presently in a state of development as are human error and common cause failures (dependent failure modelling).

2.3 CONSISTENT IMPLEMENTATION

To consistently implement the criteria, approved modelling approaches, approved methods of quantification, and approved data sources must be specified. Thus, a criteria becomes not simply a statement of a safety goal, but also a specification of models, methods and data to be used in making a risk assessment. Without the specification of approved models, methods, and data, the statement of safety goals would be rather useless since the assessments could then be arbitrary and reviews would be almost impossible.

The depth of the specification of approved models, methods and data will depend on the intended uses of these safety goals and the amount of freedom that review committees have. If the numerical safety goals were used only as guidelines and are not strictly interpreted, or if a review group has a large amount of discretion and freedom in subjectively approving analysis, then the specifications can be quite loose and general. If numerical safety goals are strictly adhered to, and if the amount of subjectivity in a review

is to be kept at a minimum, then the specifications must be detailed and very specific. Consideration of models, methods and data specifications will be addressed in more detail in later sections when specific criteria are discussed.

2.4 APPROACHES TOWARDS DETERMINING ACCEPTABILITY OR UNACCEPTABILITY

Various approaches have been explored for assessing the acceptability or unacceptability of a given risk. These approaches can hypothetically be of use in formulating numerical risk criteria. Of the approaches explored, the following categories or listing will be considered:

1. Risk comparisons
2. Comparisons with natural hazards (background) (a special case of 1. above)
3. Information obtained from revealed preference
4. Information obtained from expressed preference
5. Cost-benefit or risk-benefit evaluations
6. Methods employing combinations of the above (multi-attribute)

All the approaches have certain advantages and disadvantages, and none yield results that are satisfactory to everyone. The approaches are discussed below.

2.4.1 Risk Comparison

One of the techniques which can be used to gauge a specific risk is to compare it (as measured in some units) to other risks (as measured in the same units) to which the individual and society are subjected. Comparison can give a relative ranking for any specific risk. The measure or numbers that we are using to characterize the risk can be assessed to be higher or lower than the equivalent measure of other risks. These other risks may be more familiar or acceptable. Whether these comparisons can determine whether or not the risk is acceptable is an open question. For example, people on the West Coast have a lower risk of death due to hurricanes than those living on the East Coast,

but a higher risk of death due to earthquakes and floods. Yet, very few people considered comparison of risks like these in determining preference for where to live. In the area of man-made hazards, some people in urban areas where mass transportation is easily available still choose to drive their own cars to work, even though there are statistics that show that mass transportation is generally less risky than the use of private automobiles. Again, these decisions on alternate modes of transportation do not seem to be made on the basis of risk comparison. Thus, in decisions made by individuals or alternatives, risk measures and comparisons are not able to represent all the important factors which affect the decision.

Even when risk is the dominant factor in decision making, commonly used measures of risk ("risk numbers") do not often explain all the factors associated with the risk. In the case of electric power plants, for example, risk comparisons might show that the additional risk incurred by the addition of the plant to the community is much smaller than most of the existing risks due to natural and man-made hazards in the same community. The public and the decision makers oftentimes do not believe the analysis because of perceived omissions. Even having accepted the siting of an electric generating plant in a particular location, the choice between options for the plant (fossil vs. nuclear), is not necessarily made on the basis of comparison of inferred risk numbers. Estimates of risks to the population for both fossil and nuclear plants, even though they may be low, are oftentimes perceived as having large uncertainties and not being demonstrable.

There are, however, instances where options were inferred to reduce risk (by comparison) and decisions are made apparently based on risk evaluations, even though all the individuals involved may not perceive the same risk. An example of an action apparently taken largely on the basis of risk comparison (in this case risk reduction) is the inclusion of seat belts in all cars sold in the U.S., even though individual passengers often choose not to make use of these safety devices.

It has been suggested by some authors that the examination of existing hazards can be useful in determining what is an acceptable risk.[1, 2] Farmer[3] for instance established a risk curve for nuclear power plants that was lower than the risk of early fatalities due to naturally occuring

thyroid cancer or leukemia for the population at risk, implying that since the latter is "accepted" by society, any risk which is mathematically lower should be "acceptable". In order to avoid controversy over definition of terms, we could substitute "less objectionable" for "acceptable".

Comparison of risk measures gives only part of the solution. By their very nature, as previously stated, the commonly used measures of risk (probability of fatalities, expected consequences, probability vs. consequence curves, etc.) cannot incorporate all the factors associated with the risk. Also, the risk measures do not account for improvements in present risks. Figures 2.1, 2.2 and 2.3 are examples of risk comparisons between nuclear power plant hazards and other common hazards in the U.S. Figures 2.1 and 2.2 were taken from WASH-1400[4] using data from 1900 to 1973. Figure 2.3 was taken from BNL NUREG 51338 a Brookhaven National Laboratory report by A. Coppola and R.E. Hall,[5] and uses data from 1938 to 1977. The Coppola/Hall data shows a decrease in the high consequence area for the common hazards in the U.S. for the later period.

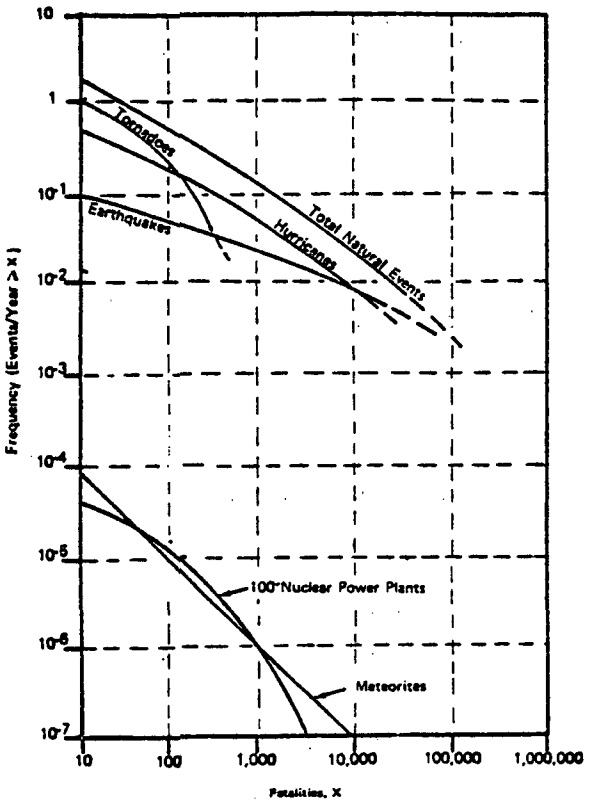

Fig. 2.1 Frequency of Natural Events Involving Fatalities (Ref. 4)

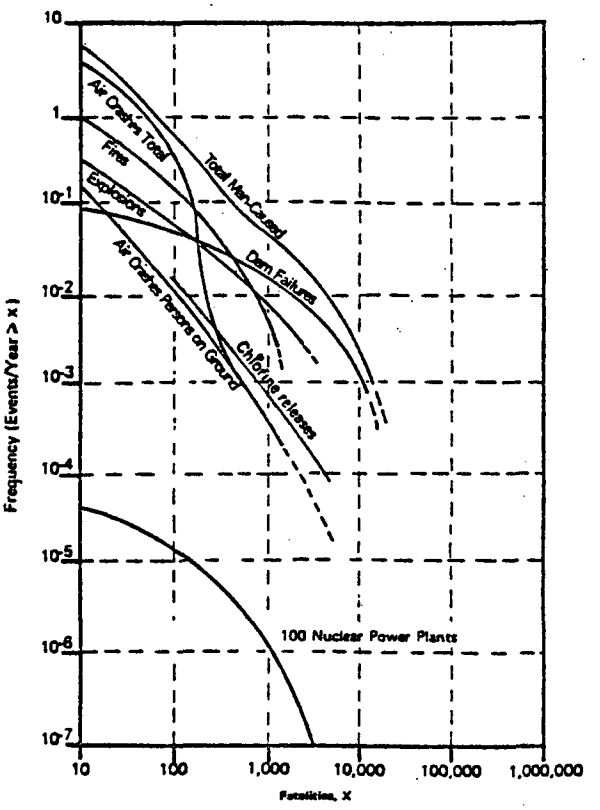

Fig. 2.2 Frequency of Man-Caused Events Involving Fatalities (Ref. 4)

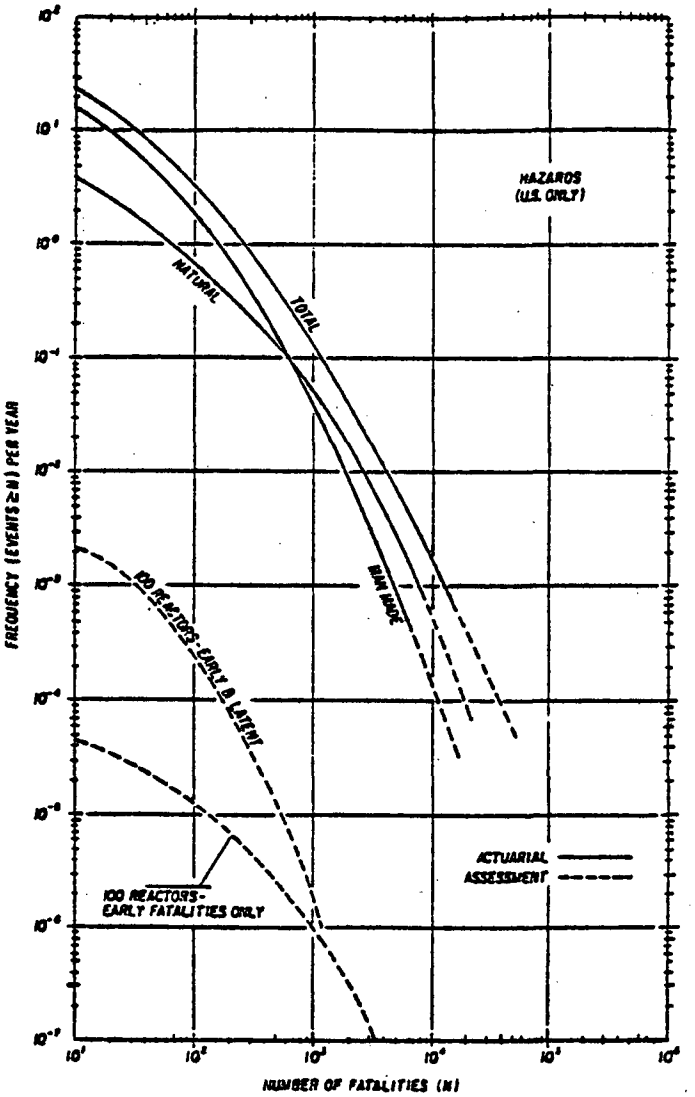

Fig. 2.3 Hazards (U.S. Only)(Ref. 5)

The data, even if the risk were constant with time, do not provide an obvious guide or suggestion as to what an acceptability criterion should be for nuclear power plant risk.

Comparisons may be used to indicate what is unacceptable. For example, one might argue that the curve of probability versus consequences for nuclear power should not be higher than the highest probability-consequence curve for man-made or natural events. We might even argue that it should not be higher than the lowest curve for natural or man-made events. These arguments, however, do not say where the criterion should be, only where it should not be

(i.e. giving criteria on what is unacceptable, but not what is acceptable). Furthermore, these arguments do not consider benefits or limitations due to uncertainties in methodology and data.

2.4.2 Natural Hazard (Background)

A special or limited form of the risk comparison technique is one that uses only risk measures of certain natural hazards which are considered background risks. The selection of an appropriate background hazard for comparison is not always possible. For instance, in the case of fire, an individual is either close enough to be burned or is not. Repeated exposure to non-harmful fires does not present any risk to an individual from burns. it is therefore difficult to think of any kind of fire as a "background risk." In the case of nuclear power plants however, all individuals are exposed to some form of ionizing radiation due to other causes, which does constitute a risk from an equivalent hazard. We call this background or natural radiation. This background radiation risk comes from many sources including the following:

1) Cosmic (solar and other extra-terrestrial sources)

2) Terrestrial (naturally radioactive materials such as rock, brick, metal deposits, etc.)

3) Man-made (medical X-rays, weapons tests, etc.)

If we include man-made sources, the better term is background radiation and not natural radiation, but the following comments apply equally to the case where only cosmic and terrestrial sources are included.

If we take an average dose per individual in the United States from these background causes (in the order of 100 mrem per year) and divide this number by a suitable "protection" factor, say 10, which is purported to give a dose that will result in no additional detectable fatality or morbidity from ionizing radiation, we end up with a background dose (of about 10 mrem per year) which can be argued to be "harmless". We can use this "harmless" dose as a target for technologies such as nuclear power plants. This target might then be considered an acceptable risk criterion which might be achievable in the U.S. for both normal operation and for accidents. The problem is that

this criterion is an average value and like any statistical average value does not take into account the extra concern associated with high consequence, low probability events.

If we are concerned with distribution of risk to specific individuals, and if we try to set an individual risk criterion using this approach of comparing with the background, the formulation of criteria becomes much more complex and tenuous. Figures 2.4 and 2.5 taken from Reference 6 show the variation of individual dose rates for cosmic and terrestrial sources. Using this data we can estimate a low value approaching 20 to 30 mrem/yr for individuals living on the east and gulf coast areas to a high value of over 150 mrem/yr for individuals living in the Colorado plateau area. The large variation in individual background dose rates is a major complicating factor and makes it difficult to select any one value as a criterion, and shows that the use of an average dose as discussed above is somewhat artificial.

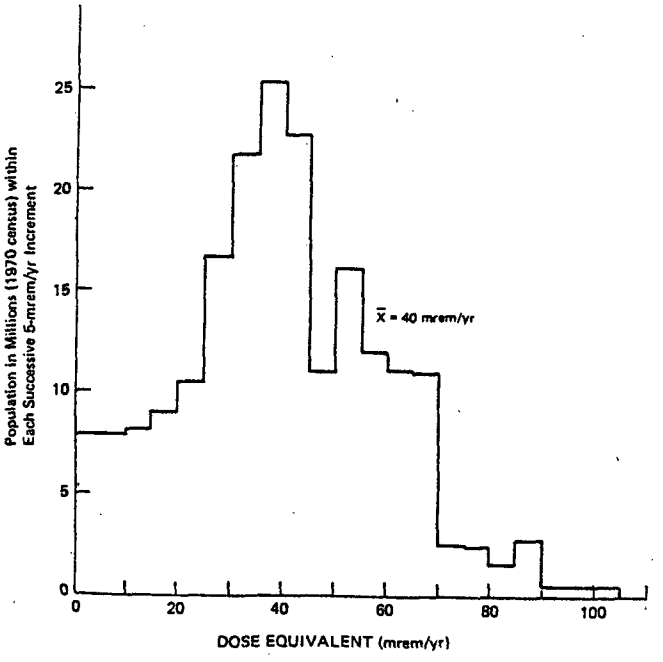

Fig. 2.4 Population distribution vs. dose-equivalent rate of radiation from terrestrial sources. (Ref. 6).

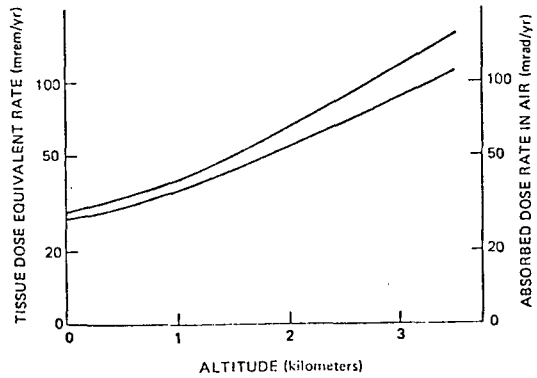

Fig. 2.5 Long-term average dose rates from cosmic radiation. The charged-particle absorbed dose rate in air or tissue is shown in the lower curve and the total DE rate (charged particles plus neutrons) is shown in the upper curve for a depth of 5 cm in a 30 cm thick slab of tissue. A quality factor of 2-10 was assumed for the range of energies within the neutron component. (Ref. 6).

A comparison with background risk thus gives the same problems as those associated with risk comparisons in general. The comparison gives information on the relative ordering of risks but gives no information on acceptability. Also, what is unacceptable is not uniquely defined.

2.4.3 Revealed Preference

Chauncey Starr[7] first proposed the method of revealed preferences for examining acceptability. The method essentially examines past activities which have been accepted by society to infer acceptable risk standards for present and future activities. He analyzed historical data on risks to infer a relationship between risk and benefit. The analysis of historical data to infer risk acceptability is the essential feature of the revealed preference method. By defining risk in terms of expected fatalities per unit of exposure, and defining benefit in terms of dollars per person (i.e. hourly wage received), Starr inferred accepted risk versus benefit curves for voluntary and involuntary risks (including different degrees of voluntary), and showed

increasing risks accepted for increasing benefits received (a cubic power law was indicated). The problem with the revealed preference approach is that it assumes that what was accepted in the past is acceptable now and in the future. Society in the past may have tolerated certain risks or may have even been unaware of their size. What has been accepted or implicitly tolerated in the past may not be accepted or be tolerated at present or in the future. Conversely, what was not tolerated in the past may be tolerated in the future because of additional considerations such as increased benefits received. Therefore, the approach is not a reliable method for inferring the unacceptability or acceptability for new risks or new technology such as nuclear power plants.

One characteristic of the revealed preference method is that both benefits and risks were considered in the analysis. In this sense it was one of the first approaches to use multi-attribute considerations for risk acceptability. The difficulty in these considerations is the definition and quantification of risks and benefits on commensurable terms so that they may be compared. It is not clear, for instance, that hunting, skiing, and smoking provide similar benefits, nor that railroads provide less benefits than general aviation as indicated in some revealed preference analyses,[7, 8] and as inferred in Fig. 2.6.

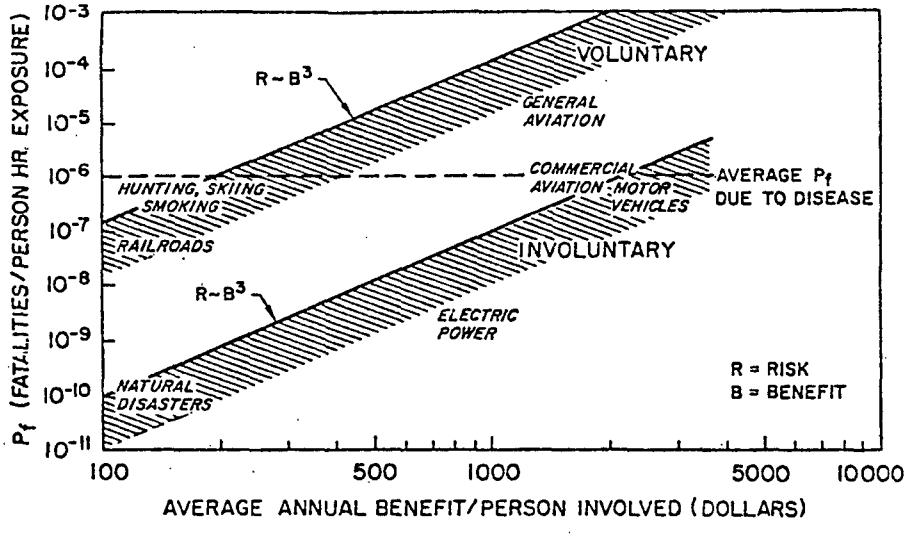

Fig. 2.6 Revealed Risk-Benefit Relationships (Ref. 8).

2.4.4 Expressed Preferences

The method of expressed preferences described by Otway,[9] and Slovic et al,[10] uses the stated preference of individuals to infer acceptability or unacceptability for any given risk. The individuals' preferences of what risk they would accept or would tolerate are solicited by surveys of various forms. This method can be viewed as an extension of Starr's approach and considers additional factors involved in risk acceptability, the most important ones being risk aversion and risk perception of individuals measured by direct polling of groups in the population. Comparing the survey results of risk rankings with other rankings such as those by experts or actuarial data is what in essence is used to give an indication of the public's "perception" of the risks. Using the expressed preference technique, Slovic et al,[10] inferred that perception was apparently dependent on several factors in addition to the degree of voluntary or involuntary involvement. The subjects of the poll were asked to rate 30 risks in nine separate factors as follows:

- voluntary - involuntary
- effect immediate - effect delayed
- chronic - catastrophic
- common - dread
- certain not fatal - certainly fatal
- known to exposed - not known to exposed
- known to science - not known to science
- controllable - not controllable
- new - old

It was found that there was great correlation between the responses for some of these attributes, and concluded that only about half were required to obtain the desired "perception factors". In particular, the nine characteristics could be collapsed into two dimensions, each representing a specific combination of the original nine characteristics. The vertical dimension approximates a level of technological sophistication, and the horizontal dimension primarily reflects the liklihood of a mishap being fatal. Fig. 2.7, taken from Ref. 11 illustrates this relationship.

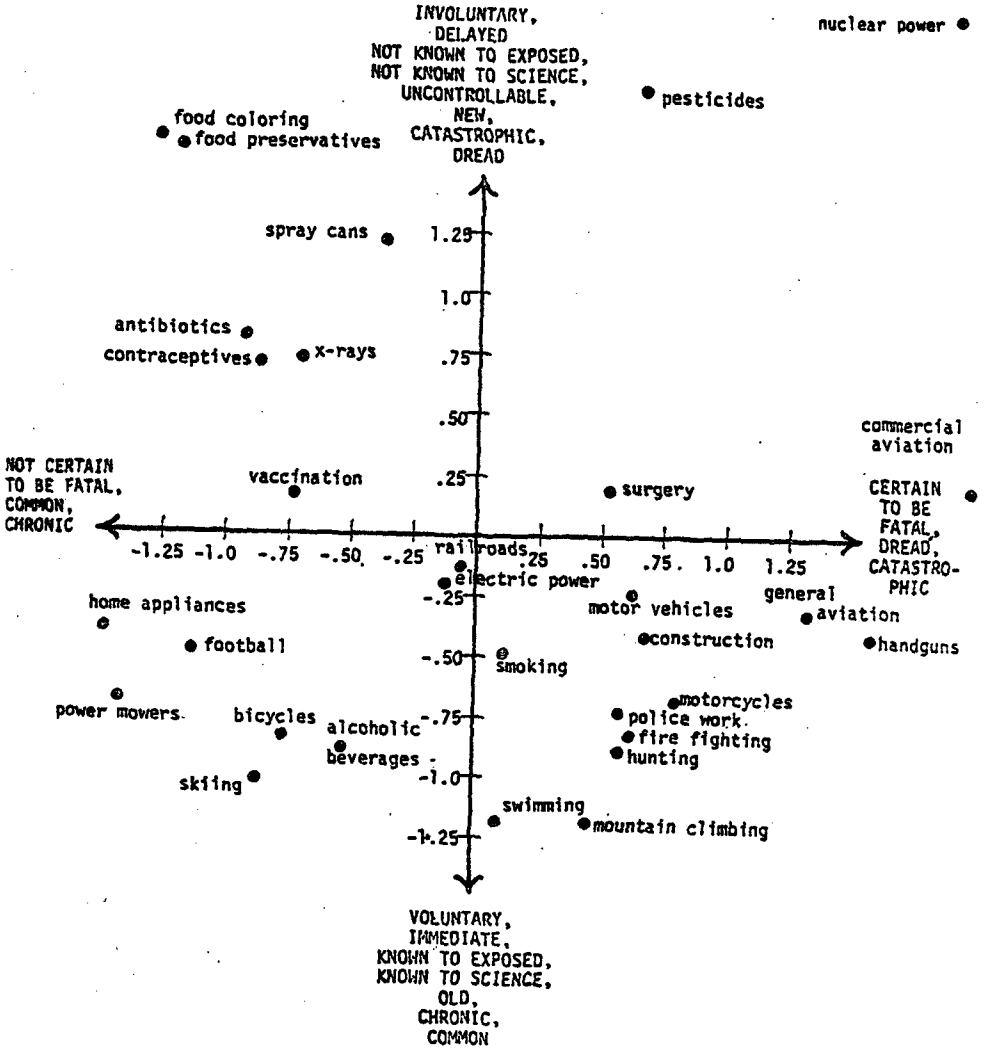

Fig. 2.7 Location of Risk Items within the Two-Factor Space (Ref. 11).

It should be noted that nuclear power, one of the risks included in the survey, scored at or near the extreme on all of the characteristics associated with high risk. According to the authors,(10) "its risks were seen as involuntary, delayed, unknown, uncontrolled, unfamiliar, potentially catastrophic, dreaded, and severe (certainly fatal)."

The use of this technique does not result in any recommended criteria for acceptability, but it does illustrate the public's attitude toward certain risks. The expressed preference technique does indicate that the public's perception of and aversion to risks should theoretically be included in a criterion which defines an acceptable risk level. This is because the perceived risk may be quite different from the calculated risk. Risk perception or risk aversion at present have not been incorporated in any dependable manner in any proposed criteria. (Some notable attempts have been made in this area, see for example Greismayer Simpson and Okrent, "The Use of Risk Aversion in Risk Acceptance Criteria?", UCLA-ENG-7970, June 1980).

Use of this method reduces the dependence on past history as a predictor of social preferences, but adds some variability to any prediction made. Since factors such as the public's knowledge or awareness of a certain risk affect the perception of that risk, we would expect this perception to be a changing phenomenon for new risks.

2.4.5 Risk Benefit or Cost Benefit

Risk-benefit or cost-benefit methods are most often practiced by business and government when confronted with a small number of alternative actions when benefits and costs (or risks) can be enumerated and put in commensurate terms. In fact, the underlying principle for these approaches is that both risk and costs can be expressed in the same terms (dollars most often used) and thereby the choice between alternatives can be made by finding the lowest net cost or highest net gain. Risk-benefit and cost-benefit methods are not conducive to establishing an "acceptable" level of risk since varying levels of acceptability are established, depending upon benefit and value considerations. When the alternatives and all the benefits and costs (risks) can be enumerated and quantified in commensurate terms, then these approaches are valid, since what is acceptable depends on the benefits and particular alternatives available.

When considering nuclear plants, however, at present all the benefits and risk measures which must be considered and their appropriate quantification have not been identified or agreed upon. Also, the alternatives depend upon case by case considerations, and general guidelines for nuclear risks are not easily obtained from these methods, even if benefits and risk measures were enumerated.

2.4.6 Combinations and Multi-Attribute Theory

There have been numerous attempts at combining various aspects of the methods and techniques listed above in order to produce a method for establishing acceptability without the inherent disadvantages of each of the techniques listed above. For the most part, these have led to complex systems requiring value measures or utility functions to make decisions. These approaches are again designed to select one of several stated alternatives. The input required to use these approaches, particularly as applied to nuclear power plants, are in the formulative stages and generally not usable at present. When these formal decision approaches are extended and modified to address the acceptability of nuclear power plant risk and benefits, and when input data become available and are understood, then these approaches will be usable tools for decision making. In the meantime, much development work needs to be done in this area.

2.5 PRESENT USE OF THE AVAILABLE METHODS AND APPROACHES

None of the approaches discussed above will yield acceptability criteria that are satisfactory to everyone or even a wide majority. They can, however, if consistently applied, be used to establish unacceptability criteria. For example, if an activity produced a risk, measured by a probability versus consequence curve, which was higher than the risk from any previous or existing man-made activity, then that new activity would be supect and prone to be termed unacceptable. nly after careful consideration of its other attributes and benefits would the new risk be entertained at all.

The approaches previously discussed can thus be used to check consistency and implication of any formulated criterion but by themselves cannot be used to actually formulate criteria for what is acceptable.

Furthermore, particularly with respect to risk assessments of nuclear power plants, there are large uncertainties in risk models, methods of quantification and data. Because of these large uncertainties, when a calculated risk falls below some criterion, we cannot generally have reasonable confidence that the calculated value is really correct, and that the risk is acceptable. On the other hand, if the calculated risk is higher than some criterion and if extreme conservatism has not been used, then we can be reasonably confident that the risk is unacceptable.

Once a criterion has been proposed, we can use the various approaches mentioned above to examine the implications of the postulated initial criterion. From these implications, and from critiques by various concerned groups and individuals, modifications to the criterion will be suggested. A new iteration of this review process will be started. Hopefully, starting with criteria on unacceptability, as the iterations proceed and new knowledge gained, we will get closer to the formulation of acceptability criteria.

The following chapters will attempt to show a part of this iterative review process, especially with regard to implications of certain postulated criterion.

Chapter 2 References

1. Levine, S., "TMI and the Future of Reactor Safety" presented at the AIF International Public Affairs Workshop, Stockholm, June 1980

2. Kinchin, G.H., "Design Criteria, Concepts and Features Important to Safety and Licensing", UKAEA, 1979

3. Farmer, F.R., "Siting Criteria - A New Approach" presented at IAEA Symposium on Containment and Siting of Nuclear Power Reactors - Vienna, 1967

4. Rasmussen, N. et al, "Reactor Safety Study", WASH-1400, (NUREG 75/014), U.S. Nuclear Regulatory Commission, 1975.*

5. Coppola, A. and Hall, R.E., "A Risk Comparison" NUREG/CR-1916, BNL NUREG-51338, 1981.**

6. National Academy of Sciences report, "The Effect on Populations of Exposure to Low Levels of Ionizing Radiation", 1980

7. Starr, C., "Social Benefit vs. Technological Risk", Science, Volume 165, 1969, pp. 1232-1238

8. Starr, C., "Benefit-Cost Studies in Socio-Technical Systems" Committee on Public Engineering Policy Report, National Academy of Engineering, Washington, D.C., April 1971, pp. 17-42

9. Otway, H.J., "Risk Assessment and the Social Response to Nuclear Power" Journal of British Nuclear Energy Society, Vol. 16, No. 4, pp. 327-33, October 1977

10. Slovic, P., Fischoff, B., and Lichtenstein, S., "Rating the Risks", Environment, Volume 21, No. 3, April 1979, pp. 14-39

11. Fischoff, B., et al, "How Safe is Safe Enogh? A Psychometric Study of Attitudes Towards Technological Risks and Benefits", UCLA Eng. 7717, January 1977

*Available free upon written request to the Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

**Available for purchase from the NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission, Washington, DC 20555, and/or the National Technical Information Service, Springfield, VA 22161.

3. TYPES OF CRITERIA WHICH CAN BE FORMULATED

3.1 HIERARCHICAL STRUCTURE FOR RISK CRITERIA

This section is concerned specifically with nuclear power plants and the development of a conceptual framework for considering different criteria, each focusing on a particular aspect of nuclear power plant risks. The framework allows investigation of the specific coverage of a particular criterion and comparison of different criteria with regard to the areas to which they apply. This framework is called a hierarchical structure for risk criteria. Risk criteria are placed in our hierarchical structure according to their coverage. The most general criteria addressing unacceptable risks to society or to the individual are placed at the top of the structure. These criteria are the "most general" in the sense that they are not directly concerned with how reliable nuclear safety systems should be, or how nuclear plants should be designed, or how population siting should be established. These top level criteria are concerned only with the final risk to society or to an individual as measured by some particular risk number, and are called risk number criteria.

The criteria immediately below the risk number criteria are those which address the unacceptable amounts of radioactivity that could be released to the environment from accidents. These release criteria again are not concerned with how reliable safety system should be, or how a plant is specifically designed or operated, as long as amounts of radioactive releases and their associated frequencies are not above some unacceptable criterion level.

Below the release criteria are criteria which address unacceptable probabilities for different kinds of accidents; these are called accident probability criteria. Below these are other criteria which address unacceptable levels of availabilities for systems in nuclear power plants, termed system availability criteria. The lowest level criteria are those which address unacceptable levels for component availabilities and human error rate probabilities, which are called component availability criteria with "component,"

used to mean both hardware and personnel. As evidenced by their very definitions, the criteria address more specific sources of risk as they progress to lower levels.

Figure 3.1 is a diagram of the various kinds of criteria and their locations in the hierarchial structure. As shown, the risk number criteria are divided into those concerned with measures of individual risk (such as the probability of an individual dying) and those concerned with measures of societal risk (such as the expected number of fatalities from accidents). The accident probability criteria are divided into those addressing total core melt probability and those addressing specific accident scenarios (sequences). The system availability criteria are divided into those addressing process systems (non-safety systems) and safety systems. Finally, the component availability criteria address constraints on hardware and human failures.

In practice, any one criterion, or some set, or all the criteria at the different levels may be used. Which criterion or which set is used depends upon, among other things, the specific utilization, the specific purpose for using the criteria, and the availability of manpower, data, and methods required in calculating results to compare against the specific criterion level. Figure 3.2 is a simple diagram of the type of information (e.g., data, models, and assumptions) required in a risk evaluation study which is intended to show compliance with a specific risk criterion. These various considerations, and the specific ways the criteria may be expressed are discussed in broad terms in the rest of this chapter and in greater depth in subsequent chapters. The essential features of the different levels of criteria in our hierarchy are described next.

3.2 TOP LEVEL RISK NUMBER CRITERIA

The top level risk number criteria concern either a societal risk, an individual risk, or both. The societal risk criterion, as measured by some number or set of numbers, focuses on health and/or economic consequences to the population at large, and their associated probabilities. The individual risk criterion can be formulated to focus on the risk to a specific individual, again as measured by some number or set of numbers.

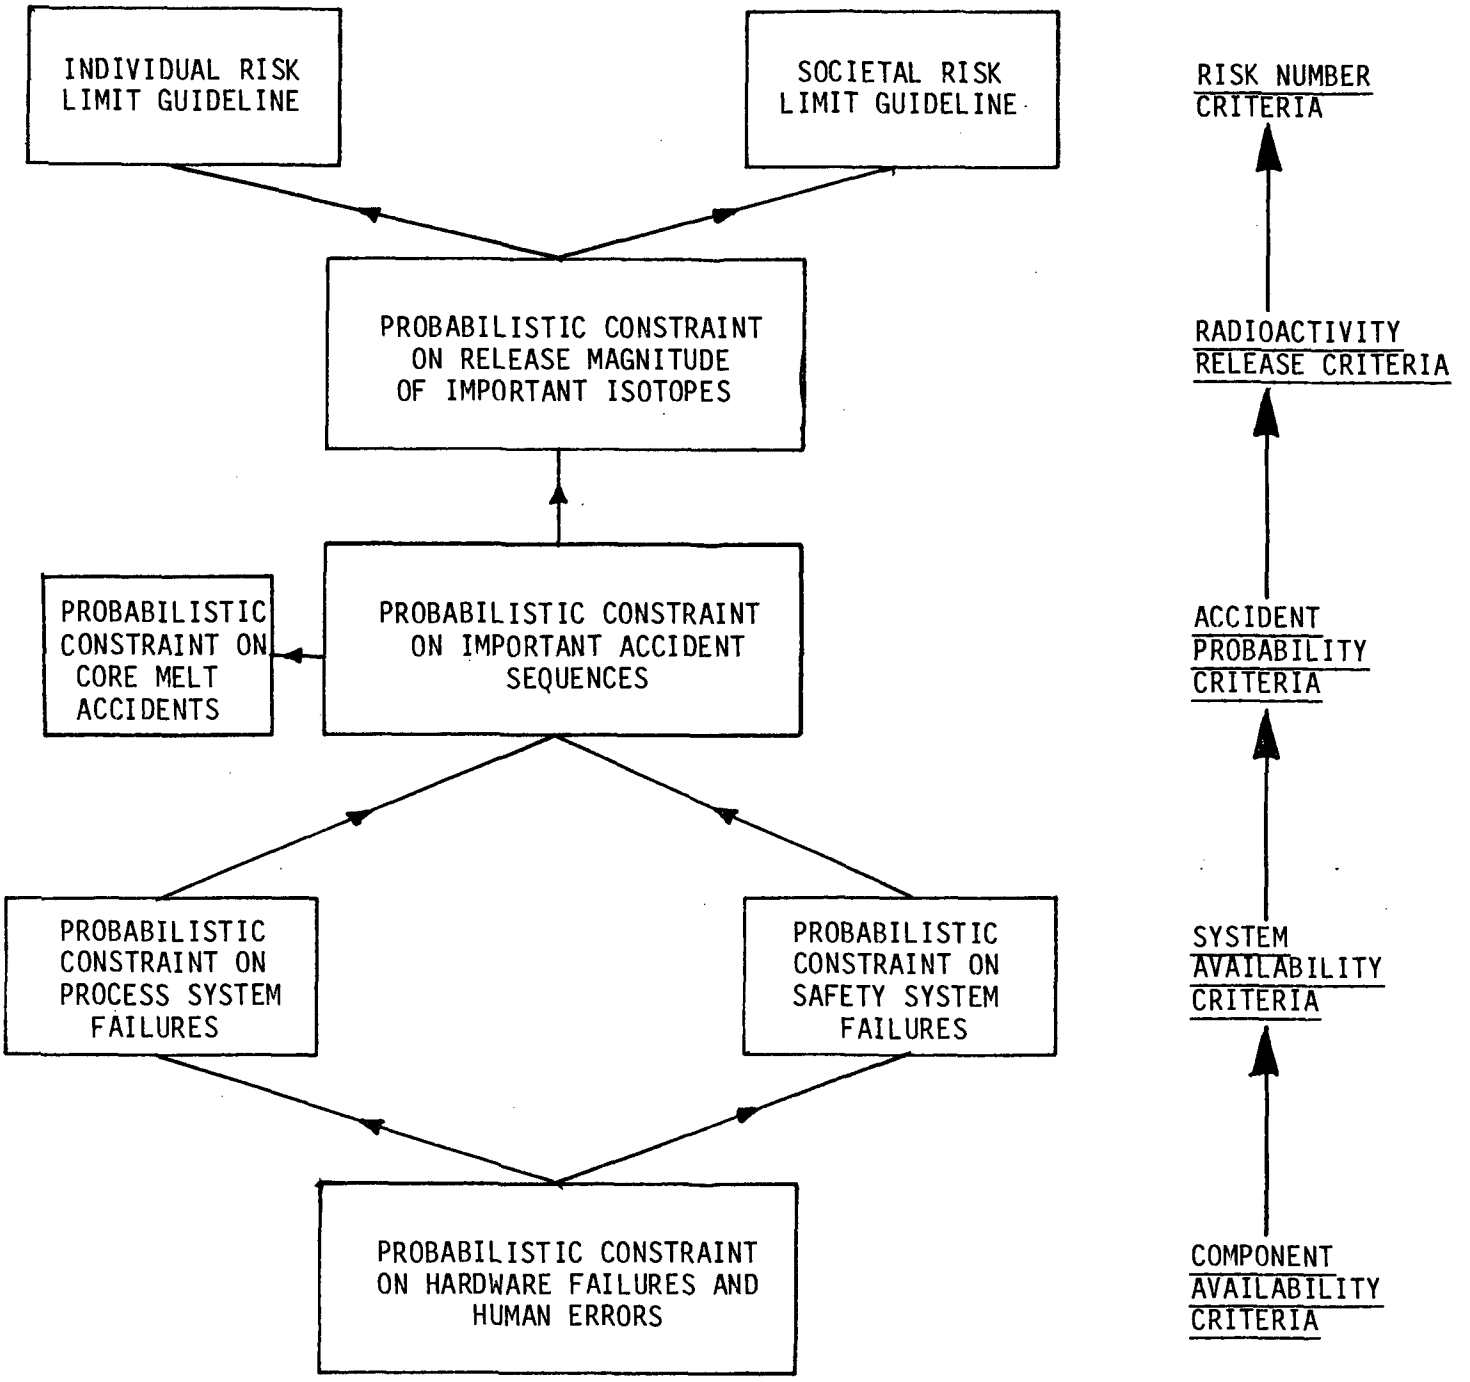

Fig. 3.1 Hierarchical Structure for Risk Criteria

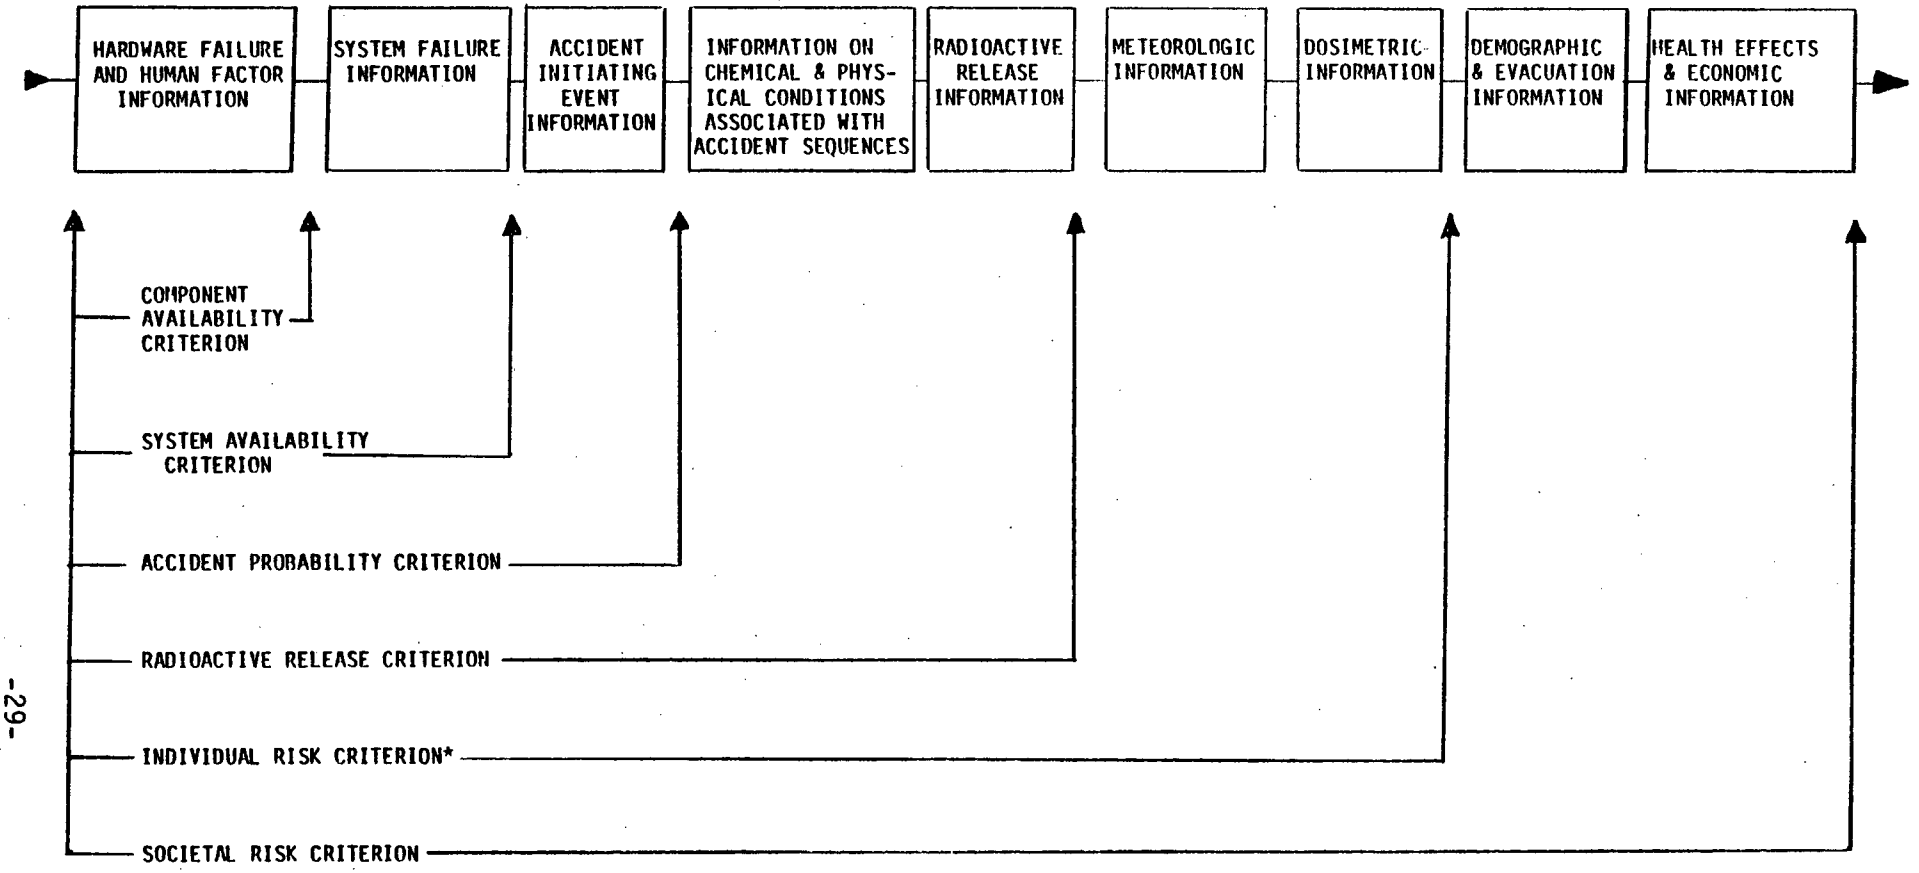

Information = [Models, data, assumptions, etc.]
*for some formulations of this criterion dose vs. health consequences information may be required.

Fig. 3.2 Information Needs of Different Criteria

Application of a top level risk number criterion involves the following process. First some specific accident scenarios that should be considered are defined. These scenarios describe which specific safety systems fail, in what mode they fail, and the postulated consequences resulting from an accident scenario (including predicted pressure, temperature, radioactivity released, etc.) In predicting the consequences to the population in terms of fatalities, injuries. or whatever other specific consequences are of interest, some models and assumptions must be adopted, regarding for example, the failure mode of the pressure vessels, the temperatures and pressures associated with the radioactive releases, the time periods involved, the way radioactivity is transported to the population, the size and distribution of the population, and the effects of the doses on the people.

The top level risk numbers involve not only the effects of doses on people but also the probabilities of these effects occurring which include the probabilities of the specific accident scenarios taking place. Estimation of the probabilities of adverse health effects requires the adoption of models and assumptions on the probabilities of systems failing, of specific amounts of radioactivity being released, of particular weather patterns occurring, and of specific numbers of people being affected by the doses. Chapter 11 presents a more detailed discussion of the evaluation process that must be performed to utilize the top level risk criteria number as well as all other levels of criteria. It is assumed here that the evaluation process has been carried and has provided numbers to be compared with the societal or individual risk criterion.

An attractive feature of top level risk criterion is that it puts the safety of a plant into a unified form by taking into account all the different elements that influence the risk due to a plant, viz., the likelihood of accident-initiating events, safety system reliability, the frequencies of various accident sequences, the integrity of containment, evacuation measures, and site-specific features such as population and meteorology. Ironically, this attractive feature is somewhat self-defeating because the uncertainties associated with the numerical assessment of each safety aspect are aggregated and this results in large uncertainties in the final evaluation of the societal risk which make it difficult to decide whether a plant has or has not

met the criterion. Moreover, the assessment of societal and individual risks of adverse health effects is sensitive to current radiobiological models, which are controversial (both to experts and to the public). Also, the state of the art in assessing the magnitude of releases from partial core melt accidents is not in general advanced enough to provide reasonable estimates.

A societal risk criterion may be constructed to constrain the risk from potential accidents to the people around a plant site, or to the U.S. population at large with all nuclear power plants considered. Any attempt to establish a criterion for all nuclear plants, existing or planned, implies that a level of total allowable risk from nuclear accidents has been allocated a priori. If the societal risk associated with each plant can be estimated, then the total risk can be found by simply adding the individual plant risks and determining whether the total satisfies the criterion. Since all nuclear plants have some societal risk, the overall societal risk criterion to be satisfied could restrict the number of nuclear power plants allowed to operate at any time.

The advantage of a societal risk criterion governing all existing or planned nuclear power plants is that it controls the total risk from such plants. It is again emphasized that what is being controlled by comparing the inferred risk level with the criterion is one measure of risk, i.e., a number or set of numbers, which certainly does not and cannot include all factors associated with the true and complete risk picture. Another advantage from a pragmatic standpoint is that older plants can have slightly higher risks than newer plants without being considered unacceptable as long as the total risk is controlled. Thus, this allowance of plant to plant variability accommodates a learning process.

On the other hand, the allowance of plant to plant variability can be considered a disadvantage if individual control on each plant is desired. Another disadvantage of a criterion governing all existing or planned nuclear plants is the difficulty of forecasting the future growth of nuclear power, and also of forecasting the possible future gains in safety or reliability that might lower the risk.

Whether for one plant or for all plants, the societal risk criterion may be formulated with the help of a variety of approaches: comparison with other non-nuclear societal risks, revealed preference, expressed preference, cost benefit or risk benefit, or combinations of the above. These approaches outlined with their advantages and disadvantages in Chapter 2, are helpful, but as stated there, an iteration process is also necessary.

One approach that can be taken to circumvent the problem of accounting for plant to plant risk variability, if it is considered to be a problem, is to define a site-specific societal risk criterion. This would ensure that any plant, no matter where it is located, be it near a large city or in a very low population area, has to operate within some guidelines on societal risk. An implication of the site-specific societal risk criterion is that it establishes a tradeoff between the engineered safety and the location of a plant, so that permission for construction of a new plant, once denied on the grounds of unacceptable design with corresponding high societal risk, may conceivably be granted if the same plant is constructed at a remote location. This means that, even if weak designs have been recognized, the site-specific societal risk criterion, or the societal risk criterion considering all plants, is not structured to deal with this issue. Because of lack of knowledge and the presence of uncertainties, an accident in which a particular system design or a high system availability is critical may not have been considered. Top level criteria do not control the individual elements contributing to risk and therefore do not address specific sources of risk such as safety system availabilities, component availabilities, and human error probabilities.

Because the top level risk criteria control only the final risk number from only the specific accident scenarios considered, the omission of any accident scenarios whose probabilities are not negligible could result in the neglect of important systems, components, and human errors. However, on the other hand, the top level risk criteria give designers and plant operators flexibility in deciding what optimal path to follow, taking into account economics and factors other than risk, while still maintaining overall control of the societal and/or individual risk by conforming with the top level risk number criteria.

An important question in the formulation of a societal risk criterion is whether or not its numerical value should include society's perception of nuclear risk. To gain public acceptability, public perception may have to be taken into account. However, if the sole purpose of the criterion is to identify high risk plants (outliers) among present reactors, incorporation of the perception factor may not be important.

Formulated in the right way, a societal risk criterion can be made understandable to the public and to public policy makers because the risk number is expressed in units that allow comparison with similar risk numbers from natural occurrences and other human activities. The societal risk criterion may be expressed in various forms, e.g., expected frequency of early or latent fatalities, etc., and/or a complementary cumulative distribution function which constrains the frequency of events that lead to X or more fatalities or injuries. The societal risk from nuclear accidents may be expressed in terms of different types of consequences, e.g., prompt fatalities, latent fatalities, genetic effects, thyroid nodules, radiation-related illnesses, property damage, etc. In this study attention has been focused primarily on two types of health consequences, early and latent fatalities, and on one economic consequence, property damage. Choices of definitions, units, and values for the top level risk criteria are discussed in Chapters 8, 9, and 10.

In lieu of societal risk as a top level risk number criterion, a criterion on individual risk can be used, aimed at assuring that an individual is not exposed to large accidental risk as measured by some risk number. The evaluation process is the same: after the consequences and probabilities are estimated for the assumed accidents, an individual risk number or set of numbers is evaluated instead of a societal risk number. A criterion on individual risk may be formulated on the basis of an "average"* person in the population at risk or a specific person at some reference location with respect to a plant site or a specific person exposed to the highest risk in the event of a nuclear accident. The individual risk criterion, formulated on the basis of

*The risk of fatality of an "average" person is operationally defined here as the societal risk measured by the expected number of fatalities per year divided by the size of the population at risk.

either of these specific persons, focuses on the consequences of an accident, regardless of the population size and distribution around a reactor site --unlike the societal risk criterion, which requires that these demographic features be considered. In essence, an individual risk criterion formulation focuses on the radiation dose and the associated health effects estimated for a person at some reference location at a certain distance from the reactor. The societal risk criterion may allow high doses at a location if no or few people are there. The differences between the societal and individual risk criteria are exemplified by the variation in the estimated consequences from the same release of radioactive material due to an accident at two different sites, one in an area where the population density is high, and the other in an area where it is low. In the first case, a large number of fatalities would be expected, and in the second case, fewer fatalities. The dose at a given point from the reactor may be the same and may pass the individual risk criterion, but the societal risk as measured by the expected number of fatalities from the accident is higher for the high population density site. Use of the societal risk criterion may lead to the construction of plants at remote locations with higher individual risk in terms of the doses allowed at given locations from the plant.

This decision as to which top level risk number criterion is to be used, societal versus individual risk criterion, both, or none, depends on consideration of each criterion's focus, and the evaluation of actions to be taken if the criterion is not satisfied. In addition, the decision is influenced by the availability of necessary models, data, and manpower to perform the evaluations and compliance assessments.

3.3 PROBABILISTIC RELEASE CRITERIA

A probabilistic criterion on the releases of various radioisotopes to the environment represents the second highest level in the hierarchy of the criteria set, as shown in Fig. 3.1. This criterion is site independent and judges the adequacy of the integrated engineered safety built into the plant. The criterion focuses on the availability of safety systems and containment

integrity and attempts to control the frequency and consequences of accidents by specifying amounts of radioactivity and associated probabilities which are unacceptable.

In theory, a release criterion could be defined by limiting the amount of radioactivity that could be released in any accident. Such a criterion might read something like, "No plant shall be operated if it is capable of releasing more than X amount of radioactivity from any given accident." This kind of criterion would certainly be welcomed, but it is generally neither feasible nor meaningful since there would always remain some non-zero probability that the release could be larger than what the criterion allowed.

In practice, a release criterion is defined by specifying either an unacceptable curve of probability versus amount of radioactivity released, or some characteristic related to the curve such as the expected amount of radioactivity released in an accident. The amount of released radioactivity, which is limited in a probabilistic sense, could be specified for different isotopes or for an appropriate sum total of all radioisotopes.

The evaluation process required to implement the release criteria is the same as that required to implement the top level risk number criteria except that the following are not required:

1. data, assumptions, and models for transport of released radioactivity to locations away from a reactor;

2. data, assumptions, and models which relate the amount of radioactive material to radiation doses and thence to adverse health effects;

3. demographic and evacuation models.

A probabilistic release criterion attempts to ensure that a plant, regardless of site-specific features such as population density, meteorology, and evacuation efficiency, has an adequate level of safety. In the past, environmental release of the isotope iodine-131[1, 2] has been isolated and focused upon in defining release criteria. Since releases to the environment of various other radioisotopes may result from an accident, a more comprehensive release criterion might require that these be accounted for, perhaps weighted in some manner according to their individual health effects. On the

other hand, if the amount of radioiodine released is considered to be a sufficient indicator of the severity of the specific accidents being evaluated, then a criterion on iodine release may be adequate. The release criterion does not take into account the number of resulting adverse health effects, which depends on the dose at various locations and on the population density and distribution. Instead, the release criterion attempts to control public risk by focusing on its source -- the amount of radioactivity released and the associated probabilities. Demonstration of compliance with the release criterion regarding assessed releases due to partial core melt is difficult because of the lack of models and data for partial core melt analyses. Also, the criterion does not directly constrain the frequency of occurrence of accidents that may be perceived as serious but do not result in significant environmental releases, e.g., the Browns Ferry Fire;[3] neither do the top level risk number criteria directly constrain or control such accidents. These non-release accidents, however, would be indirectly controlled if they were associated with the same safety system failures or initiating events as are accidents involving release.

3.4 ACCIDENT PROBABILITY CRITERIA

A criterion for the frequency of accidents may be specified on the basis of selected individual scenarios and/or certain classes of accidents, e.g., loss of coolant accidents, transient-initiated accidents, and accidents leading to core melt. Some of the accident sequences may lead to complete core melt and others to varying degrees of core damage ranging from high cladding defects to partial core melt. All these accident probability criteria depend on specific accident scenarios being postulated and being quantitatively evaluated for comparison with the appropriate criterion (as do the top level risk number criteria and the release criteria). If some accident is not hypothesized, then it will not be evaluated and controlled by comparison with the criteria. If the set of defined accident scenarios is fairly comprehensive, it has a good chance of including all the pertinent safety system failures, human errors, etc., that would be involved in the non-hypothesized accident

scenarios. Also, extrapolation to other plants of accident sequences identified as being important in one plant may not necessarily be valid because of design differences in the present generation of reactors.

This criterion attempts to ensure that the frequency of defined accidents regardless of their consequences is kept within some specified limits and considers the reliability of engineered safety systems, the effect of system interactions on plant safety, and the frequency of accident-initiating events. In implementing this criterion, the effects of consequence-mitigating systems in a nuclear power plant are not necessarily considered if they are not included in the accident scenario. Furthermore, the effects of plant siting are not included. The criterion, however, does focus on the basic sources of accidents: system failures, component failures, and human errors, which are controllable in existing as well as new plants.

The evaluation process for the accident probability criterion is considerably simpler than that required for the higher level criteria (top level risk number criteria and release criteria). It consists of first defining a set of accident sequences and then identifying the system failures, systems interactions, component failures, and human errors that lead to the defined sequences. The probabilities or frequencies of these accident sequences are estimated by using available failure rate data and reliability modeling assumptions.

3.5 SYSTEM AVAILABILITY CRITERIA

The next lower level criterion in the hierarchy of risk criteria specifies minimum standards for the availability of safety systems and/or restricts the frequency of occurrence of accident-initiating events. Safety systems are referred to in this report include both man and machine. Some systems in power plants are not in continuous operation but are required to start and continue operation for a specified time period at a certain performance level in response to some plant conditions. The term availability when applied to such a system means availability on demand for operation for the required

period and at the required performance level. This criterion level circumvents the problem of incomplete identification of the key accident sequences required for all the higher level criteria, (the top level risk number criteria, the release criteria, and the accident probability criteria.) The system availability criterion is more easily demonstrable than any of the higher level criteria, and the uncertainties in the assessment intended to show compliance with it are smaller.

The disadvantages of this criterion level include neglect of siting and lack of consideration for key combinations of safety system failures required for severe accidents, and their interactions. Moreover, the consequences of safety system failures and accidents including radioactivity released and its impact on human health and on the environment, are not considered. The basic premise of the criterion is that risk is controlled if the likelihood of safety system failures is constrained; risk to the public originates with safety system failures - if there are no safety system failures then there is no risk. On the other hand, specification of one number for the availability of a safety system does not take into account the different degrees of importance of that system in different accident sequences. Thus this criterion does not give the designer or plant operator flexibility in choosing alternatives.

The evaluation process for this criterion involves specifying the component failure, test and maintenance, and human error contributions to be included in system failure definition, and the reliability/availability models, component failure rate, test and maintenance information, and human error data that are to be used.

3.6 COMPONENT AVAILABILITY CRITERIA

The lowest level criterion in our hierarchy of risk criteria specifies minimum standards for the availability of components in process and/or safety systems. The term availability as it applies to components is used in the same sense as it applies to systems. The term 'component' includes both hardware and its man-machine interface. This type of criterion is used in the

military. The basic premise of the criterion is that risk is controlled if the possibilities of component failures or malfunctions are constrained. If no component failures occur, then an accident is not initiated, nor is a safety system challenged. Without an accident-initiating event followed by a system failure there can be no accident, and without an accident there is no risk.

One deficiency of this criterion is that specification of one number for the reliability of a particular component does not consider the different environmental conditions and stresses under which the component will be called upon to perform its function, or the different numbers of challenges to a standby component in different safety systems, or the relative importance of the same component functioning as part of different systems. Moreover, common cause failures of more than one component are not explicitly considered. In addition, a component availability criterion does not take into account the relative importance of a component in different systems, or how critical a failure of a component is as an initiator or a mitigator of an accident, or what the consequences of a component failure are.

The evaluation process required to implement the component availability criteria is simpler than that for any other criterion discussed so far. The required information includes component failure rate and human error data and maintenance data obtained from tests and from experience. In addition, component availability and maintainability models are required.

Chapter 3 References

1. Farmer, F.R., "Reactor Safety and Siting: A Proposed Risk Criterion," Nuclear Safety, Vol. 8, No. 6, November - December 1967.

2. Meleis, M. and Erdmann, R., "The Development of Reactor Siting Criteria Based Upon Risk Probability," Nuclear Safety, Vol. 13, No. 1, January - February 1972.

3. "Annotated Bibliography of Safety-Related Occurrences in Boiling-Water Nuclear Power Plants as Reported in 1975," ORNL/NUREG/NSIC -126, July 1976.

4. COMPONENT AVAILABILITY CRITERIA

4.1 INTRODUCTION

This chapter is concerned specifically with the component availability criteria, which are the lowest level criteria in the hierarchical structure of risk criteria described in the preceding chapter. The component availability criteria address unacceptable levels of availabilities for components in nuclear power plants. The word component is intended to include both hardware and its man-machine interface. The term availability refers to the state of the component as it affects the safety of a plant, i.e., if the component became unavailable or malfunctions, it could act as an initiator of an accident, or degrade the initial and/or long term performance of a safety system. The rationale behind establishing criteria on component availabilities is that risk is controlled if the unavailabilities of components, or occasions when components perform their functions inadequately are constrained.

As yet, no specific proposals for quantitative component availability criteria have been made for components used in nuclear power plants except that contained in the Standard Review Plan[1] for diesel-generator reliability testing, which establishes a reliability goal of 0.99 at a nominal 50% confidence level. Therefore, the rest of this chapter is concerned not with specific proposals but with some of the considerations involved in specifying component level availability criteria and the implications of establishing such criteria.

4.2 CONSIDERATIONS FOR SPECIFICATION OF COMPONENT AVAILABILITY CRITERIA

The first step in establishing component availability criteria is to identify the components which are to meet the criteria. In general, these are components whose failure or degraded performance could initiate an accident, or could lead to ineffective operation or failure of safety systems.

After such a component has been identified, the next step is to delineate the boundaries of the component. For example, a diesel-generator unit treated as a component might include in its boundary not only the engine and the generator but the combustion air system, fuel supply system, lubricating oil system, cooling water system up to the supply, starting energy sources, autostart controls, etc.

Once a component has been identified and its boundaries defined, the specification of the criterion might entail consideration of the type of the component. Components can be classified into four types[2]: (i) a constant availability component is defined on the basis of per demand availability independent of time (human errors per demand can be modeled as a constant availability component), (ii) A non-repairable component is one which, if it fails, is not repaired during plant operation, (iii) A periodically tested component is tested and/or maintained at intervals according to some schedule and procedure, and (iv) a monitored component is one whose condition is continuously monitored.

The specification of an availability criterion of a component may also have to include the different failure modes of a component. For example, a relief valve has two principal failure modes, failure to open on demand and failure to reseat properly, and therefore, may require two availablility criteria, each addressing a particular failure mode for the same component. Some components in safety systems, which are on standby, are required first to start and then to continue operation for a certain time period in response to some abnormal plant conditions. For such components, criteria may be specified in terms of their availability on demand and their availability during the required time period.

The specification should state the coverage of the criteria, that is, the different kinds of contributors (e.g., hardware, test and maintenance, and human errors, involved in the operation of the component) that are to be included in the assessment of component availability to show compliance with the criteria. The coverage of a criterion for a specific component would depend on, among other factors, the type of the component and the boundary defined for it.

4.3 IMPLICATIONS OF COMPONENT AVAILABILITY CRITERIA

One type of component availability criterion may be formulated to establish a single number or set of numbers for a generic class of components, e.g., pumps, motor-operated valves, etc. This type of criterion does not distinguish between components that belong to the same generic class but are exposed to different environmental conditions and stresses. Furthermore, it does not take into account the relative importance of components. Importance variability of components within the same generic class could arise for the following reasons:

- the use of the same generic class of components in different configurations within a particular system - an intrasystem component importance;
- the use of the same generic class of components in systems differing in their safety significance - intersystem component importance;
- the difference in performance capacities of individual components relative to the overall system requirements, for example, in some designs with three auxiliary feedwater pumps, one has 100% capacity and the other two 50% each;
- the difference in the potential contribution to accident initiation from failures of particular components belonging to the same generic class;
- the difference in the potential contribution to safety system performance from failures of particular components belonging to the same generic class;
- the varying consequences directly accompanying failure of particular components belonging to the same generic class, for example, failure of a relief valve to reclose properly has different consequences depending on whether the valve is on the secondary side of a steam generator or on top of the pressurizer.

To remedy some of the deficiencies of establishing a single numerical criterion at the generic component level, in principle, multiple criteria could be established for each generic class of components. A specific criterion could then be selected from among the multiple criteria and applied to components within the same generic class but belonging to different systems. The association between a generic class of components and systems could be

formulated on the basis of a system's function rather than its name to account for design variability among plants. The level of detail in specifying component reliability criteria in terms of a component's generic class and system affiliation may be extended to include the different environmental conditions in which a component operates (or is likely to). For example, components of the same generic class can be located inside or outside of containment, or components of a given type may be operating in the control room and also near a high energy pipe break.

Even if a set of component availability criteria were established on the basis of generic classes of components, their systems affiliations, and their environmental conditions of operations, it would still have deficiencies in that it would not address potential common cause failures of more than one component and would not take proper account of redundancy as it affects system availability. The component availability criteria are the most inflexible compared with all other levels of criteria in our hierarchical structure for risk criteria. They may not provide designers and plant operators flexibility and incentive to achieve high system availability by way of selecting the optimal system design and maintenance strategy. On the other hand, an advantage of establishing component availability criteria is that the evaluation necessary to show compliance with them is simpler than for all other higher levels of criteria (discussed in Chapter 3). The evaluation requirements are detailed in Chapter 11. Another advantage of component availability criteria is that they focus attention on components some of which may not be included in the evaluation of system availability. This type of criteria might also aid designers and plant operators in selecting components on the basis of inherent reliability.

Chapter 4 References

1. "Standard Review Plan for Review of Safety Analysis Reports for Nuclear Power Plants: LWR Edition," Appendix 7-A, Branch Technical Position BTP E1CSB2, Diesel-generator qualification testing, NUREG-75/087, P.B. 247611, September 1975. *

2. Vesely, W.E. and Goldberg, F.F., "FRANTIC - A Computer Code for Time Dependent Unavailability Analysis," NUREG-0193, October 1977. *

*Available for purchase from the National Technical Information Service, Springfield, VA 22161.

5. SYSTEM AVAILABILITY CRITERIA

5.1 INTRODUCTION

Assessing the general effectiveness of nuclear rectors safety systems, in terms of the risk reduction, requires criteria by which to judge its adequacy. Chapter 5 addresses the second level of the hierarchical structure as described in Chapter 3. The system availability criteria should specify the minimum standards for the availabilities of selected processes and safety systems. In this way, a system whose availabilities are lower than those specified in the standards would be considered unacceptable. This second level in the hierarchical structure focuses on integrating component failures, human errors associated with operation as well as test and maintenance that contribute to system unavailability, and common mode failures of components within the defined system boundary. In this way, the analysis encompasses a broader scope of plant safety than the component availability criteria. However, this level does not address the consequence of the system failures nor does it account for interactions between systems. The system availability criteria do not directly control the probability of accidents unless the accident involves systems to which these criteria apply.

The objective of a criterion placed on system unavailability would presumably limit the unavailability of each system while a criterion placed on accident sequences would limit their frequency of occurrence. The implications of establishing such criteria would be few if all plants were identical in their entirety. However, since they are not, the implications of establishing probabilistic criteria for system and accident sequence need to be addressed.

5.2 CRITERIA FOR SYSTEM UNAVAILABILITY

It is interesting to note that the NRC presently employs a type of risk criteria for system unavailability in the licensing process. This is the single failure criteria. The single failure criteria, being a deterministic

approach can be loosely defined as those that require that no safety system shall be designed such that a single hardware component failure will fail the system. Application of this principle has in fact resulted in a certain level of hardware safety that would likely not have been achieved if there were no criteria on system unavailability. However, it is suspected that the single failure criteria is not fully adequate to limit system unavailability; its effect on limiting overall risk has the potential to be positive, but most assuredly not sufficient. In some cases the net effect of the single failure criteria could prove to have a negative impact on the overall risk.

One level to which a probabilistic risk criteria could be applied is the individual system unavailability allowed for each of a nuclear plant's safety systems. Presumably a decision would be made concerning the allowable unavailability for each safety system, and this would comprise the set of acceptable criteria. Any safety system analyzed would be required to meet the allowable system unavailability limits to be acceptable. The analysis would include all contributions to system unavailability, i.e. hardware faults, human errors, test and maintenance contributions and common mode contributions.

Several options are available with regard to the definition of a "system". The definition employed here encompasses the more or less traditional definitions of safety systems that are in standard usage for risk analysis, e.g. low head injection system (decay heat removal system), high head injection system, auxiliary feedwater system, etc. Different reactors might employ diverse systems to accomplish a single function. For instance, one reactor might have both a fan cooling system and a containment spray system to accomplish post-accident containment heat removal, but another reactor may only have a containment spray system. Should credit be given for this diversity in establishing system availability criteria for the containment spray system of the reactor with diversity? If so, this implies that criteria will have to be generated with diversity in mind; and different criteria will apply to different reactors. If not, then reactors employing system diversity would be penalized. How to treat system diversity is one of the difficulties in developing criteria for system unavailability. The problem is further compounded if, for instance, the containment spray system also performs the function of removing radioactive materials from the containment atmosphere.

Another difficulty with this approach involves the explicit definition of a system. Some safety systems can be explicitly defined with respect to the components comprising the system while others cannot. In a risk analysis, judgements are often made concerning the components to be included in one system or another. This is justified on the basis that, since all the safety systems (and components) are to be analyzed, it is not critically important whether certain components are included in any one system definition and hence excluded from another; eventually they will all be implicitly included on the event tree sequences. However, this does not assure the completeness of the study now. If risk criteria are applied to systems, it is important that an explicit definition of the components comprising each system be identified. One can imagine a situation where, for instance, component failures are arbitrarily assigned to a system with a low estimated unavailability in order to make the high unavailability system marginally acceptable. An example might be a system designed to deliver a certain flow of water during the injection phase of an accident that contains valves that must change position. This system is initiated by a logic control system. Does one consider the valve control logic (valve driver) to be part of the system that provides flow, or part of the logic control system? Depending on the unavailabilities of the two systems with respect to their risk criteria, the choice may be made deliberately in such a way that both systems pass the criteria. Another approach is to follow a functional analysis. In this case the exact definitions of systems are not needed since the analysis includes all components that are required to fulfill a given plant function, e.g., reactivity control.

Another consideration arises when systems are required to work during both the injection and recirculation phases of an accident. Risk criteria would be required for both phases of system operation; a limit on system unavailability that applies to the injection phase, and a limit on system failure probability during the recirculation phase. If criteria were established separately for each phase, then these would not directly address the possibility of non-recovered failures during the injection phase that could make system failure during the recirculation phase more probable. For instance, an ECCS system may be required to operate during both the injection and recirculation phases of an accident. Failure of one leg of the ECCS during injection that is not recovered by the recirculation phase will result in a greater

probability of failure of the ECCS during recirculation. Combinations of injection and recirculation phase failures can be quite complex. Should risk criteria on the recirculation phase operation of systems attempt to account for every such combination? This situation is another type of system interaction, and the difficulties of obtaining a direct link between risk and system failure probability discussed in the previous paragraph apply here also.

The relationships between system unavailabilities and risk may also depend on the type of accident. Some accidents are expected more frequently than others. Also, some systems are more important for particular accident types than are other systems, e.g. for some transients containment cooling may not be required to mitigate core melt. This implies that the criteria for systems will have to be set considering accident types weighted by likelihood of occurrence. This simply adds to the complexity of obtaining the relationships between risk and the system availability criteria.

A second approach in the development of a system criteria is that of the Atomic Energy Control Board of Canada presented by Atchison [1]. Their Reactor Siting Guide practices the defense-in-depth approach, based on minimizaton of the probability of human and equipment failure and on the provision of highly reliable protection systems. Here a nuclear power plant consists of two types of systems; Process and Special Safety Systems. Process Systems are defined as systems and equipment required for the normal functioning of a nuclear power plant as a power producer and include such systems as the heat transport systems, the turbine-generator and main power output systems, the reactor and plant control systems, the refuelling and spent fuel transfer systems, and so on. Were it not for the radiation and fission products produced in the reactor these would be the only systems necessary in a nuclear power plant. During normal plant operation some process system failures could lead to an event which degrades operation enough to challenge safety systems. In this event the process system is the initiator of the accident. Therefore, by restricting select process system failures, one controls the frequency of accident initiating events.

Special Safety Systems are designed to cope with postulated failures in the process systems and, as will be explained later, also failures in the

process systems combined with unavailability of any one of the special safety systems. Included in this division are the shutdown systems which prevent the neutron chain reaction, reducing reactor power to decay heat levels. The Rector Safety Guide[1] implies that the total frequency of all process failures could be as high as once every three years, however the frequency for severe process system failures is generally considered to be 10^{-2} to 10^{-4} per year. In addition, the siting guide specifies dose limits for serious process systems, and it is this limitation that requires the overall effectiveness of the Special Safety Systems. The Siting Guide recognizes that even special safety systems are fallible. The fallibility of a special safety system can be expressed as unavailability. Since a special safety system is normally in a dormant or "poised" state, ready to operate if process parameters exceed set limits, the availability of a special safety system during normal operation must be determined by testing. Considering the practical limitations of the frequency of testing and the reliability of human operators, it is felt that the best one can claim for the availability of any special safety system is .999, that is, an unavailability (=1 - availability) of 10^{-3}, or 8 hours per year, on average. Atchinson goes on to say that the probability of a serious failure (once every 3 years) occuring at the time that the appropriate special safety system(s) is (are) unavailable is about 3×10^{-4} per year. The validity of even this simple calculation depends on one important proviso, viz. that there is no cross-linking effect whereby the factors which caused the serious failure could also disable the operation of the special safety system(s) which are intended to cope with the particular serious failure. The elimination of potential cross-links dominates considerations of nuclear power plant design, construction and operation. It may be relatively easy to show on paper that a system will have an unavailability of say 10^{-6} or less. It is much more difficult to demonstrate this unavailability in practice or to show that the system is immune from cross-linked failures. Hence, emphasis is placed on design practices which experience has shown will minimize the potential for cross-linked failures between process and special safety systems and among the special safety systems themselves.

The Reactor Siting Guide then addresses the probability of occurrence for a dual failure, that is a process system failure and the unavailability of a special safety system. This probability may be estimated by multiplying the

severe failure of a process system of 10^{-2} to 10^{-4}/year by the required unavailability of 10^{-3} for the special safety system, giving a result of 10^{-5} to 10^{-7} per year of a dual failure. Atchinson concludes that the examination of actual major accidents that have occurred in reactors shows that cross-linked failures and human errors in the form of design weaknesses or operating errors predominate over random coincidence when one is considering accidents of very low probability. This fact is not surprising when one considers the amount of effort that goes into ensuring the low probability of coincidence of purely random and unrelated failures in more than one system or component. Experience has shown that an equal effort must be expended to ensure by careful design, construction and operation that the probability of cross-linked failures, dual, are indeed low.

5.3 LIMITATIONS IN THE DEMONSTRATION OF COMPLIANCE WITH A SYSTEM AVAILABILITY CRITERION

The objective of this section is to discuss the implication of incomplete data, human error, and common mode contributions as it affects the demonstration of compliance with a system availability criterion.

5.3.1 Hardware Failure Data

Risk assessment and the results they provide draw attention to the fact that meaningful statistical analysis of the results require understanding of the failure data incorporated. An example of dealing with incomplete data and how sensitive WASH-1400 event and fault trees are to various inputs is illustrated in Table 5.1. To develop the system unavailability, WASH-1400's median and upperbound point values were initially used. An attempt was then made to incorporate into WASH-1400's reduced fault trees additional sources of data, those of IEEE 500 and the work performed by EG&G. What was evident here was that a criteria to which a system analysis will be based must incorporate a common language with respect to data. Table 5.2 shows a comparison between WASH-1400 data and that of EG&G. What is illustrated here is an incompatibility of valve descriptors which prevented the input of this data into WASH-1400 fault trees. In addition, a problem in collecting data was evident and may

Table 5.1. Sensitivity of Accident Probability for Sequences from WASH-1400

Sequence Contribution to Core Melt Probability	WASH-1400 Base Case	WASH-1400 Upper Bound with T&M Same as Base Case	WASH-1400 Upper Bound with Max. Maint. Duration Min. Maint. Interval
AD	1.80 (-6)	2.60 (-6)	6.86 (-6)
AH	8.85 (-7)	4.29 (-6)	4.49 (-6)
S_1H	2.394(-6)	8.91 (-5)	8.91 (-5)
S_1D	1.34 (-6)	1.818(-4)	1.818(-5)
S_2D	3.65 (-6)	5.88 (-5)	5.88 (-5)
S_2H	4.6 (-6)	2.34 (-5)	2.34 (-5)
S_2C	2.4 (-6)	1.2 (-5)	1.7 (-5)
V	4 (-6)	1.2 (-4)	1.2 (-4)
TML	3.28 (-6)	2.66 (-5)	2.66 (-5)
TMLB'	1.64 (-6)	1.33 (-5)	1.33 (-5)
TKQ	6.45 (-6)	2.93 (-5)	2.93 (-5)
TKMQ	1.29 (-6)	5.86 (-6)	5.86 (-6)
CORE MELT PROB.	3.37 (-5)	3.886(-4)	4.128(-4)

best be summarized by an example of how valves and other components may be classified falsely. If an MOV fails due to a torque switch failure, is this classified as a valve failure or a switch failure? When the IEEE 500 data was reviewed, the same basic problem existed. The WASH-1400 data categorized instrumentation as (amplification, annunciators, transducers, combination) with "failure to operate" or "shift calibration" errors. IEEE 500 presented in detail various types of instrumentation devices, various failures incurred, partial or full, and consideration for environment factors.

Table 5.2 Component Failure Sensitivity of WASH-1400

Valves	WASH-1400	Median	Error Factor	EG&G Geom. Mean/ Error Factor	EG&G W.C.F.
MOV	Fails to Operate	$1x10^{-3}$/D	3	9(-3)7	9(-3)7
	(Plug) failure to remain open	$1x10^{-4}$/D	3	------	3(-6)1
	External leakage or rupture	$1x10^{-8}$/Hr.	10	9(-3)7	1(-6)2
SOV	Fails to Operate	$1x10^{-3}$/D	3		
AOV	Fails to Operate	$3x10^{-4}$/D	~3	2(-2)3	2(-2)4
	(Plug) failure to remain open	$1x10^{-4}$/D			1(-5)1
	External leakage or rupture	$1x10^{-8}$/Hr.	10	8(-6)10	
CHECK	Failure to open	$1x10^{-4}$/D	3	6(-3)2	
	Reverse leak	$3x10^{-7}$/Hr.	~3	3(-6)5	
	External leak (rupture)	$1x10^{-8}$/Hr.	10	4(-6)1	
MANUAL	Failure to remain open (Plug)	$1x10^{-4}$/D	3		see below
RELIEF	Failure to open/D	$1x10^{-5}$/D	3		see below
	Premature open/Hr.	$1x10^{-5}$/Hr.	3		see below
ORIFICE FLOW METERS	Rupture	$1x10^{-8}$/Hr.	10		see below
VACUUM	Failure to operate	$3x10^{-5}$/D	~3		see below
	Rupture	$1x10^{-8}$/Hr.	10		
MANUAL	Failure to operate			2(-3)1	
	Leak externally			6(-7)1	
PWR PRIMARY SAFETY	Premature open			2(-5)2	
	Fail to open			6(-2)3	
BWR PRIMARY RELIEF	Fail to open			3(-2)3	3(-2)3
	Fail to reset			6(-3)3	7(-3)2
	Premature open			2(-5)2	2(-5)3
REMOTE & MOV	Fail to operate			9(-3)7	9(-3)7
	Leak externally			1(-6)2	
	Plugged				1(-5)1

For the low pressure recirculation system, it was possible to incorporate the EG&G data into one area of the reduced fault trees. This produced the following results, a system unavailability of 1.21×10^{-2} as the medium value and 1.42×10^{-1} as the upper bound value, compared to WASH-1400 results of 8.85×10^{-3} and 4.3×10^{-2} respectively. Input was incorporated into double failures only since it is component dominant, single failure and common mode failure are human related and Test and Maintenance is only performed during refueling. In conclusion, when system analysis comes into play with a risk criterion, it is increasingly evident that the handling of failure rates and the development of fault trees will play a major role when it comes to interpretation. What is needed is a balance between exactness and simplicity in order to reduce the complexity of a nuclear power plant without oversimplifying in order to meet the requirements for the development of a criteria.

5.3.2 Human Error

In the review of WASH-1400's fault trees, areas of uncertainty exist in an acceptable methodology to determine the human-hardware contributions to the top event of a fault tree. A study of the sensitivity of human errors was performed by P. Samanta[3]. The author states that the RSS event tree/fault tree methodology takes human intervention into account explicitly. A reading of the RSS and the sensitivity assessments performed by Kelly et al and Parkinson show that human error plays an important role in reactor safety. But unfortunately, as acknowledged by RSS and pointed out by the Lewis Report the human error performance data base is weak in many aspects. The RSS data base was developed from non-reactor relevant experiences.

Hence, it is important to assess the impact of the changes in human error rates (this term is loosely used in this report to represent the unavailability contribution of the human error) at every stage in risk assessment. Such analysis will reveal particular aspects that are more vulnerable to human errors and will also provide important information regarding our ability to reduce the risks due to human errors in a nuclear power plant. The impact of human errors categorized into generic classes with regard to their time of occurrence, location within the plant and the type of action involved

also needs to be evaluated for a better understanding of the problem. One could properly allocate the available resources, if the relative importance of various generic classes of human errors could be assessed. Importance of individual human errors need to be measured in order to identify those errors that require more attention. Also, identification of those errors that need additional attention given a particular type of accident could be very useful to the operators.

Figures 5.1 and 5.2 represent the dependence of core melt probability on human error rates. Changes in core melt probability are studied by changing all human error rates and also by varying the minumum human error rate considered.

Since large numbers of sequences contribute to core melt probability, the potential for the sensitivity of core melt probability is limited. The core melt probability shows significant increase with the increase in human error rates; however, it does not show a similar decrease with the decrease in human error rates. This pattern is largely followed until the minimum human error rate (M.H.E.R.) is increased to 10^{-3}. At this point core melt probability shows much more sensitivity, indicated by the increased gap between the curves (2) and (3), compared to that between (1) and (2) in Figure 5.1. This is explained by the fact that with M.H.E.R. = 10^{-3}, the base case human error rates start dominating the hardware failures contributing to the same event. Loosely, it could also be said that baseline core melt probability (M.H.E.R. = 10^{-5}) has comparable contribution of human and hardware failure rates, since this probability is almost twice the core melt probability with H.E.R. = 0.

The sensitivity of core melt probability obtained in this study differs from Parkinson's[4] results. The factor increase of 73.6 in core melt probability due to a factor increase of 30 in all human error rates, observed by Parkinson, seems unrealistic. Among the dominant accident sequences contributing to core melt, only S_2C-δ increases by a factor as large as 91.6 for a factor increase of 30 in all human error rates. Since this sequence contributes only about 7% to the core melt probability and the sensitivities of other systems are much less, the factor increase of 23.82, in core melt probability for a factor increase of 30 in all human error rates is believed to be the correct one.

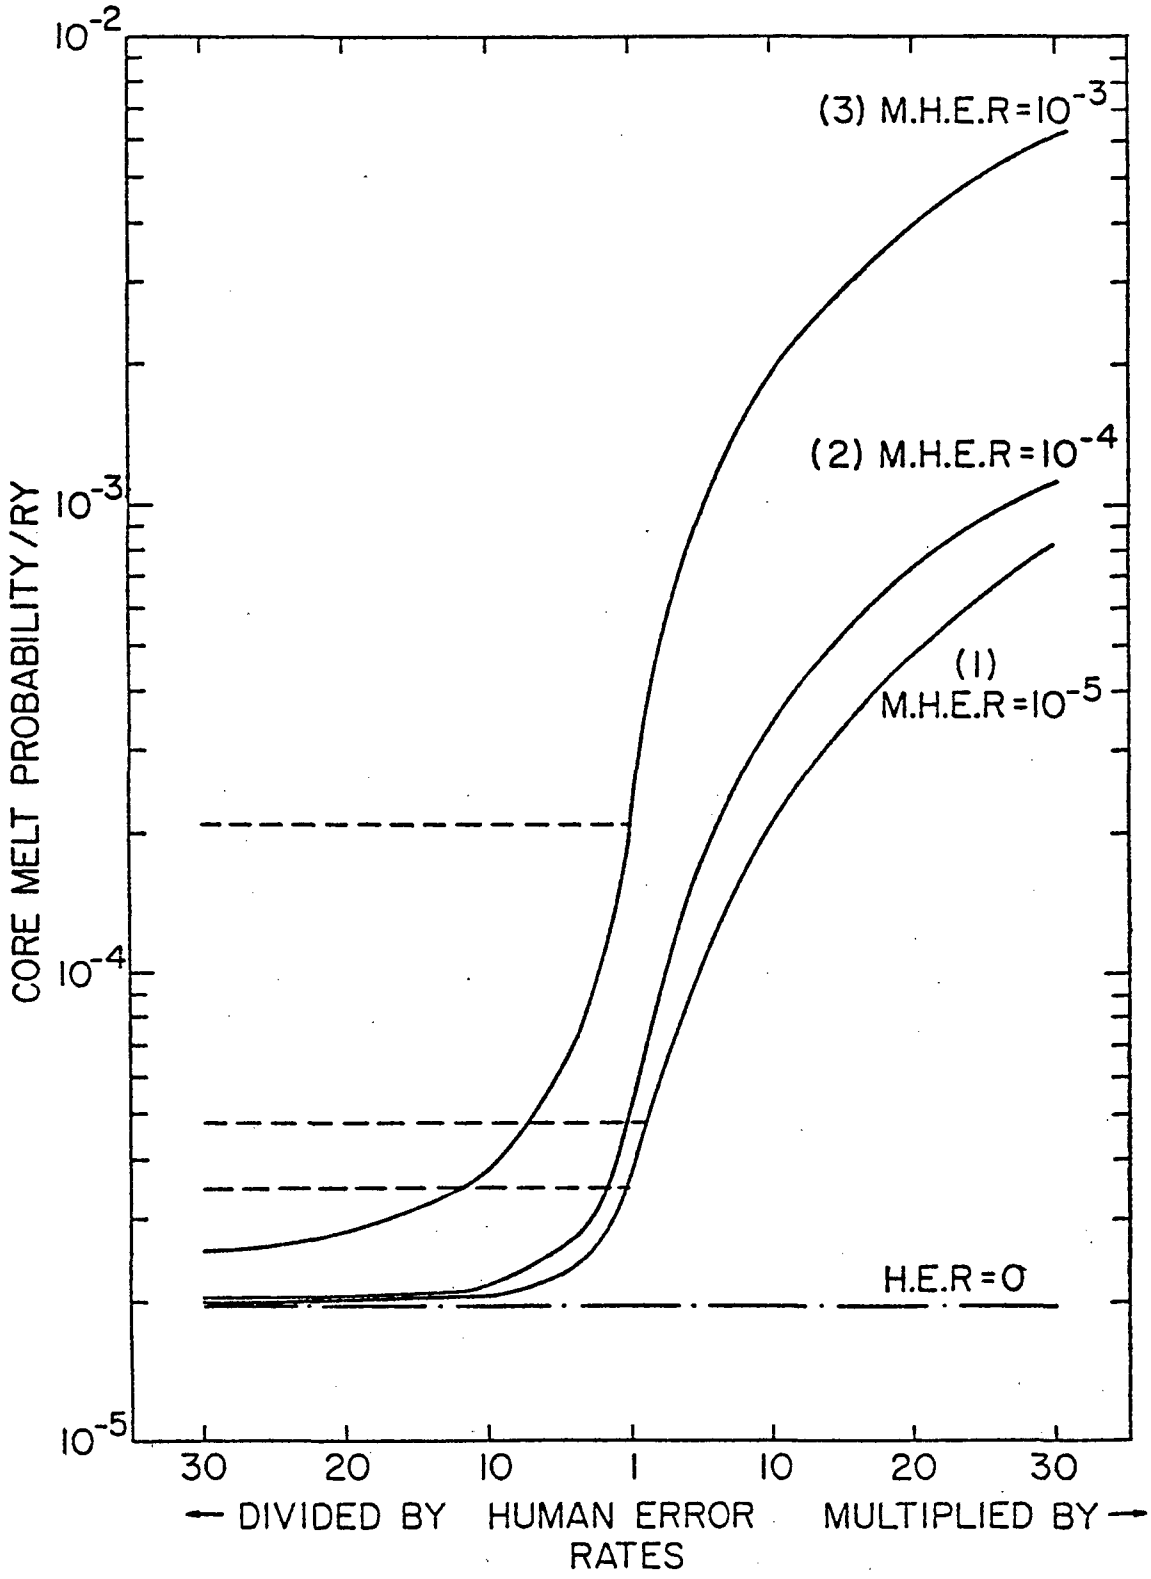

Figure 5.1. Changes in core melt probability due to changes in all human error rates.

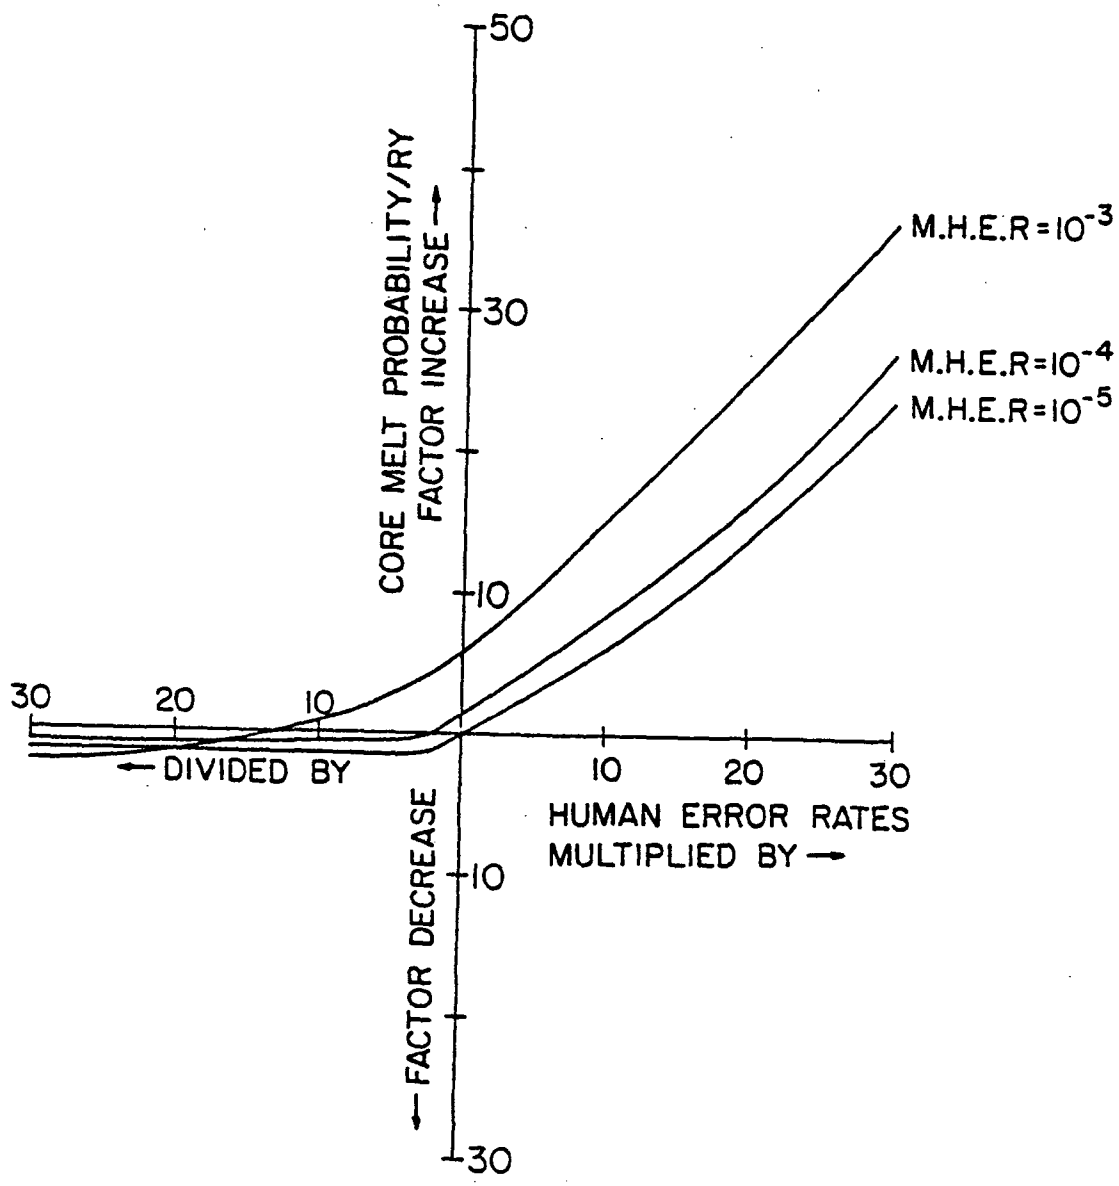

Figure 5.2. Sensitivity of core melt probability to human error rates.

5.3.3 Common Mode

System analysis requires the recognition and technique for the accountability of common mode failures between systems and components. Two types of common mode failures exist, those due to failures of several similar components due to failure of a common interfacing function or system, and those that involve failure due to some common defect such as calibration, production line or manufacturing defect. Various methods exist to account for these failures. One such method is to consider the total coupled case and and the total uncoupled case. The coupled case deals with the case when one redundant component fails, the others fail, while the uncoupled case deals with the idea that components fail completely independently from one another.

Chapter 5 References

1. Atchison, R.J., "Nuclear Reactor Philosophy and Criteria," presented on July 18, 1979 to the Select Committee on Ontario Hydro Affairs, Summer schedule of hearing on the safety hearings on the safety of nuclear reactors, July 1979.

2. Lewis, H.W., et al, "Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission," NUREG/CR-0400, September 1978.*

3. Samanta, P.K., Hall, R.E., and Swoboda, A.L., "Sensitivity of Risk Parameters to Human Errors in Reactor Safety Study for a PWR," NUREG/CR-1879, BNL-NUREG-51322, January 1981.**

4. Parkinson, W.J., "Sensitivity Analysis of the Reactor Safety Study," M.S. Thesis, Massachusetts Institute of Technology, February 1979.

*Available for purchase from the National Technical Information Service, Springfield, VA 22161.

**Available for purchase from the NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission, Washington, DC 20555, and/or the National Technical Information Service, Springfield, VA 22161.

6. ACCIDENT PROBABILITY CRITERIA

6.1 INTRODUCTION

The accident probability criteria are placed at a level in our hierarchical structure for risk criteria that is above the system availability criteria but below the release criteria (see Section 3.1). In essence, the accident probability criteria address unacceptable probabilities for different kinds of accidents. They attempt to ensure that the frequencies of some defined accidents, regardless of their consequences, are not above some specified unacceptable level. An accident probability criterion can be specified on the basis of accident sequences and/or classes of accidents. Accidents can be classified in many ways, for example, on the basis of their initiating events such as large pipe breaks in the reactor coolant system, on their potential for core damage which may range from high cladding defects to core melt, or on their potential for releasing various amounts of fission products.

6.2 PROPERTIES AND FORMS OF ACCIDENT PROBABILITY CRITERIA

In general, an accident is caused by an accident-initiating event followed by failures of one or more systems. Consequently, the accident probability criteria focus on the frequencies of accident-initiating events, the availabilities of safety systems, and the effects of operator intervention and system interactions on plant safety. Although an accident probability criterion focuses on these potential sources of accidents, it is not concerned with how reliable safety systems or components should be, or how a plant is designed and operated as long as the probability of accidents is not above some specified unacceptable level.

The control exercised on the frequency of a certain class of accidents by an appropriate criterion depends on specific accident sequences belonging to the same class being postulated and quantitatively evaluated. If some accident sequences are not hypothesized, then these will not be evaluated and

controlled by comparison with the criterion. If the set of postulated accident sequences is fairly comprehensive, it has a high possibility of including all the relevant safety system failures, accident initiating events, human errors, etc. that would be involved in the non-hypothesized accident sequences. In this situation, the frequencies of these non-hypothesized accident sequences would be indirectly controlled by the criterion.

Let us suppose that an accident probability criterion has been defined on the basis of some important accident scenarios and this is applied to all plants. The effectiveness of such a criterion in controlling the frequencies of important accident scenarios with respect to any specific plant may be questionable since extrapolation to other plants of accident scenarios identified as being important in one plant may not necessarily be valid because of design differences in the present generation of reactors.

The frequency of accidents which result in small or no release of radioactivity to the environment can be directly controlled by an appropriate accident probability criterion. These non-release accidents cannot be directly controlled by other levels of criteria that exist in our hierarchical structure. Although these accidents may not be important from the standpoint of public health and safety, they may be perceived as being serious by the public. In addition, these non-release accidents may be precursors to more severe accidents. Therefore, by restricting the probability of these accidents by an appropriate criterion, the frequency of more severe accidents may be controlled. In spite of the insignificant adverse health effects due to these accidents, there is an incentive on the part of a utility to restrict their probability since they can impose large financial penalties.

According to WASH-1400,[1] for the present generation of LWRs, risk to the public is dominated by core melt accidents. One way to reduce public risk is to restrict the frequencies of core melt accidents. However, it is to be noted that not all core melt accidents cause injury to the public. It is inferred from WASH-1400 that only about 1 in 100 core melt accidents in a PWR is capable of causing one or more fatalities. In the rest of this section, it is assumed that an accident probability criterion has been defined on the basis of accidents that lead to core melt.

The evaluation necessary for the demonstration of compliance with a core melt probability criterion requires first, identification of accident sequences that lead to core melt, and then an assessment of their frequencies. These frequencies are summed and compared with the numerical value specified by the criterion to determine whether the criterion has been met. Different core melt probability criteria can be applied to restrict the probability of core melt accidents depending on their level of severity. For example, a core melt probability criterion can be established for those accidents that result in greater than x per cent liquified fuel, or those that lead to environmental releases of a certain radioisotope greater than y curies. Demonstration of compliance with these criteria would require that the appropriate measures of severity, in addition to the frequency of core melt accidents, be evaluated.

The specification of a criterion for the frequency of core melt accidents, or for that matter any class of accidents, should state its coverage. "Coverage" here means the different kinds of contributors (e.g., hardware failures, human errors, common mode failures, accidents induced by earthquakes, etc.) that are to be included in the evaluation of the frequencies of accidents to demonstrate compliance with a criterion. For example, two core melt probability criteria can be established, one to be satisfied by considering only hardware failures and the other by considering both hardware failures and human errors. Just such a proposal is presented in Section 6.4 of this chapter. The principal advantage of considering only hardware failures in one case stems from the recognition that there exists at present substantially large uncertainties in the estimates of human errors. Thus, by considering only hardware failures which are known with better certitude than human errors, one is reasonably assured that a satisfactory hardware safety level has been achieved.

Even if the same value of a core melt probability criterion is applied to all plants, it would allow plant-to-plant variation in the risk to the public. The extent of this variation depends on plant design (especially the design of consequence-mitigating systems) and site-specific features (e.g. population size and distribution, meteorology and public protection measures).

Forms of Core Melt Probability Criteria

Two types of core melt probability criteria can be formulated. One specifies a limit on the frequency of core melt accidents as it applies to a given reactor in any given year while the other specifies a limit on the total frequency of these accidents in any given year regardless of the number of reactors. Accordingly, the former type of criterion is called a R-Y criterion* and the latter, Y criterion.* The R-Y criterion by definition restricts the probability of core melt accidents for a given reactor in any given year, i.e., per Reactor-Year (R-Y). On the other hand, the Y criterion restricts the probability of core melt accidents considering all reactors operating in any given year.

The same or different numerical values of a R-Y criterion can be applied to all plants. If the same numerical value is used, i.e., all reactors have the same goal for any given year, then the decision on the safety of a plant would not take into account additional factors other than the probability of core melt accidents which affect the public risk of a plant, for instance, the effectiveness of consequence-mitigating systems and the site specific features such as the size of the surrounding population. These additional factors can be accounted for if different numerical values are specified by a R-Y criterion for different reactors. If different values of a R-Y criterion are used, then slightly higher frequencies may be permitted for reactors which are old, have very effective consequence-mitigating systems, or are located in remote areas. Use of the same numerical value of a R-Y criterion would not distinguish between old versus new plants. The use of two different numerical values of a R-Y criterion, one pertaining to an old plant and the other to a new plant, might be desirable since this accommodates a learning process. In addition, retrofitting of an old plant to conform with the same standard which also applies to a new plant may not have a favorable cost-benefit ratio. However, if the same numerical value of a R-Y criterion is applied to all

*It may be noted that both the R-Y and the Y criteria can be formulated on the basis of a given year instead of any given year. If this is done, the specified numerical values of these criteria would be different from one year to another. This type of formulation is not considered any further in this study.

plants, it ensures that the frequencies of core melt accidents in any plant regardless of its age or potential consequences are not above a certain unacceptable level. The second type of core melt probability criterion, the Y criterion, is discussed next.

It may be recalled that the Y criterion attempts to control the total frequency of core melt accidents from all reactors operating in any given year. Since all nuclear reactors can be associated with some frequency of core melt accidents, the Y criterion to be satisfied could restrict the number of reactors allowed to operate at any time. There can be two procedures which may be followed to implement a Y criterion. The first procedure consists of apportioning the total frequency allowed by the Y criterion among operating reactors. The second procedure consists of summing the frequency of core melt accidents associated with each reactor and then determining whether or not the total satisfies the Y criterion level. The essential difference between the first and the second procedures is that in the first, the Y criterion in its application focuses on the frequency of core melt accidents that is associated with each reactor, whereas in the second no direct control is exercised on this frequency as it pertains to any one reactor. In the first procedure, the apportionment of the total allowed frequency among reactors can be done equally or unequally. In the case of equal apportionment, the numerical value of the criterion as it applies to one reactor is the same for all reactors in a given year, i.e., all reactors have the same goal in a given year. This is somewhat similar to the R-Y criterion that specifies the same numerical value for all reactors in any given year. However, there is a difference because the numerical value obtained by equal apportionment of the Y criterion is dependent on the number of reactors that are operating in a given year, whereas the R-Y criterion is independent on the number of reactors. Consequently, even though the goals are the same for all reactors, these goals change from one year to another. Therefore, a reactor may pass the criterion in one year but fail in some other year. In the case of unequal apportionment, the numerical value of the criterion as it applies to a particular reactor is not the same compared to other reactors. Therefore, all reactors do not have the same goal in a given year. In this case too, these goals change from one year to another.

The second procedure which can be followed to implement a Y criterion is only concerned with the total frequency of core melt accidents from all reactors in any given year. No specific numerical values need to be met by an individual reactor. Therefore, this procedure allows plant to plant variation in their frequencies of core melt accidents. Thus, some plants (e.g., old plants) may have slightly higher frequencies than others (e.g., new plants) without being considered unacceptable as long as the total frequency is not above the specified limit of the Y criterion. This allowance of variability can be considered a disadvantage if individual control on each plant with regard to its frequency of core melt accidents is desired.

6.3 CORE MELT PROBABILITIES ALLOWED BY THE CRITERIA FOR THE FREQUENCY OF CORE MELT ACCIDENTS

This section assesses the core melt probabilities allowed by the two forms of the criteria, the R-Y and the Y criteria, discussed in Section 6.2 for various numerical values. The results of these assessments are presented in the form of tables and parametric plots. These facilitate comparison of different proposed values of the criteria by examining the allowed core melt probabilities in the lifetime of a reactor in the next decade and through the year 2000.

Core Melt Probability in the Life of a Reactor Allowed by a R-Y Criterion

A criterion for the frequency of core melt accidents for a given reactor in any given year allows certain probabilities of occurrence of core melt accidents in the life of a reactor. These are dependent on the value specified by the criterion. These allowed probabilities of core melt accidents up to any point in time t in the life of a reactor considering different numerical values specified by the criterion can be obtained from the following expression,

$$P_{cm} \text{ per reactor year} = 1 - e^{-\lambda_{cm} t}$$

where, λ_{CM}= the value specified by the criterion. Table 6.1 presents the allowable probabilities at the end of a reactor lifetime of 40 years considering different values of the criterion.

TABLE 6.1

Core Melt Probabilities Allowed at the End of Reactor Lifetime by the Criterion for Core Melt Frequency (R-Y Criterion)

Value specified by the criterion for the frequency of core melt accidents [per Reactor-Year]	Core melt probability allowed at the end of reactor lifetime (40 yrs) by the R-Y Criterion
10^{-2}	33%
5×10^{-3}	18%
10^{-3}	4%
5×10^{-4}	2%
10^{-4}	0.4%
5×10^{-5}	0.2%
10^{-5}	0.04%
10^{-6}	0.004%

Core Melt Probability from 1980 to the Year 2000 Allowed by a R-Y Criterion

In the previous paragraph, the probabilities of core melt accidents in the life of a plant allowed by a R-Y criterion were assessed. In a similar manner, the core melt probabilities allowed by a R-Y criterion from some reference point in time up to any future time can be assessed considering different forecasts for nuclear power. Consequently, these allowed core melt probabilities are dependent on the projections of nuclear power and the numerical value specified by a R-Y criterion. In order to estimate these allowed probabilities, the reactor-years of experience that are projected to accumulate in the future from all existing and planned rectors need to be assessed. Such an assessment is shown in Figure 6.1. This figure shows a range in the projected number of reactor-years from 1980 to the year 2000 based on different growth rates of nuclear power (high, low, medium) and

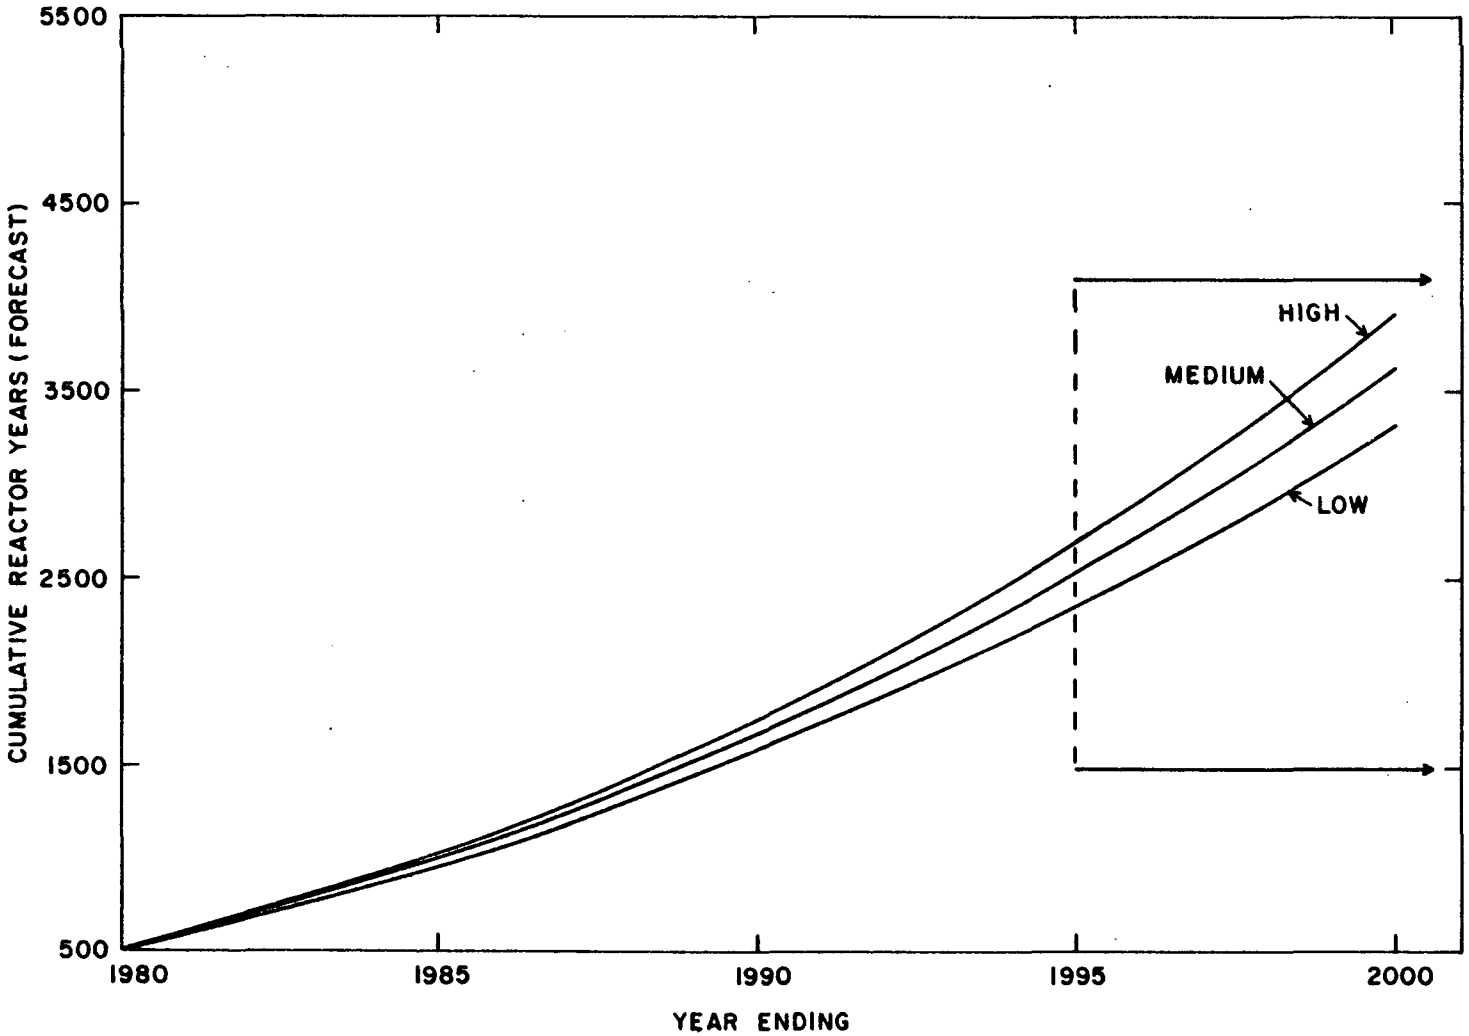

Fig. 6.1 Estimated light water reactor-years that are projected to accumulate up to the year 2000.

assuming that existing LWRs would be operational till the year 2000. A detailed explanation of how these three forecasts were derived is given in Appendix A. The allowed core melt probabilities from 1980 to 2000 implied by a R-Y criterion can be obtained from the following expression:

$$P_{cm}(T_y - T_{1980}) = 1 - \exp\left[-\lambda_{cm}\,(T_y - T_{1980})\right]$$

where,

λ_{cm} = value specified by the R-Y criterion

T_y = reactor-years that are forecasted to accumulate at the end of calendar year y

T_{1980} = reactor-years of experience accumulated up to the beginning of 1980 (assumed to be equal to 440 reactor-years as calculated in Appendix A)

$P_{cm}(T_y - T_{1980})$ = allowed core melt probability implied by the R-Y criterion from the beginning of 1980 up to the end of calendar year y.

Figure 6.2 shows the allowed core melt probabilities from 1980 till the beginning of the year 2000, $P_{cm}(T_y - T_{1980})$ versus y, for different values specified by the R-Y criterion based on the low projection of cumulative reactor-years as shown in Figure 6.1. Similar assessments can be performed considering different projections for cumulative reactor-years, for instance, the "high" and the "medium" estimates that are shown in Figure 6.1. Table 6.2 presents the core melt probabilities allowed by a R-Y criterion in the next decade (end of 1979 to end of 1989) and till the year 2000 (end of 1979 to end of 1999) assuming "high", "low", and "medium" projections of nuclear power. This table also shows the implication of a R-Y criterion based on the hypothetical scenario of moratorium on the deployment of new reactors from 1979 onwards. For the moratorium scenario, it is assumed that there were 65 LWRs operating at the end of 1979 and these would continue to operate till the year 2000 (see Appendix A).

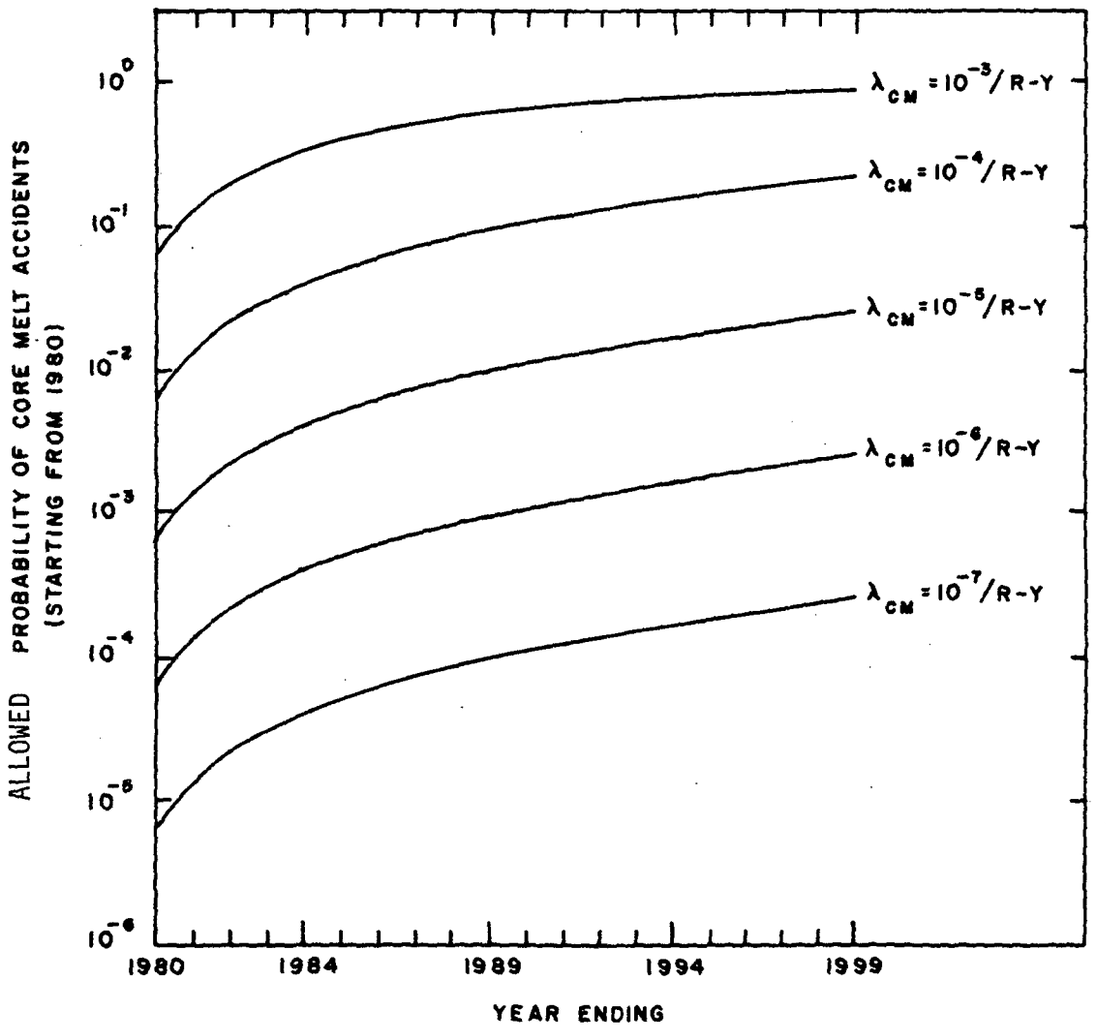

Fig. 6.2 Allowed core melt probabilities from the year 1980 onwards considering different values of the R-Y type of core melt probability criterion.

TABLE 6.2

Core Melt Probabilities Allowed in the Next Decade And Till the Year 2000 from 1980 Onwards by the Criterion for Core Melt Frequency (R-Y Criterion)

Cumulative Reactor-Years Projected to Accumulate starting from 1980		Per cent Core Melt Probabilities Allowed by a R-Y Criterion Considering the following numerical values specified by the Criterion		
		10^{-5}/R-Y	10^{-4}/R-Y	10^{-3}/R-Y
IN THE NEXT DECADE	Moratorium: 650	0.6	6.3	47.8
	Low Projection: 1011	1.0	9.6	63.6
	Medium Projection: 1088	1.1	10.3	66.3
	High Projection: 1150	1.1	10.9	68.3
TILL THE YEAR 2000	Moratorium: 1300	1.3	12.2	72.7
	Low Projection: 2677	2.6	23.5	93.1
	Medium Projection: 2952	2.9	25.6	94.8
	High Projection: 3204	3.2	27.4	95.9

Core Melt Probability Over a Period of Time Allowed by the Y Criterion

Like the R-Y criterion, the Y criterion which restricts the total accident probability in any year regardless of the number of reactors allows certain probabilities of core melt accidents over a period of time. These allowed core melt probabilities are dependent on the numerical value which specifies the limit on the total frequency per year of core melt accidents considering all reactors operating in the U.S. in any given year. Consequently, these allowed core melt probabilities are not dependent on the number of reactors operating in any given year. This is unlike the core melt probabilities allowed over a certain period of time by the R-Y criterion which are

dependent on the number of operating reactors. These allowed probabilities estimated over 10, 20, and 30 years, are shown in Figure 6.3 as a function of the numerical value of the Y criterion. These probabilities are obtained by using the following expression:

$$P_{cm} = 1\text{-}\exp\ (-\Lambda_{cm}T)$$

where,

Λ_{cm} = numerical value of the Y criterion

T = time period in years over which P_{cm} is to be estimated.

6.4 SOME RECENT PROPOSALS FOR CRITERIA SPECIFYING THE FREQUENCY OF CORE MELT ACCIDENTS

This section presents four proposals for core melt probability criteria. The first three criteria can be classified as R-Y type criteria and the fourth as a Y criterion.

Burns[2] states that the frequency of core melt accidents during the life of the nuclear industry in the U.S. should be low since these accidents tend to dominate the risk of the present generation light water reactors. The numerical value of the criterion for the acceptable frequency of core melt accidents recommended by Burns is 5×10^{-7}/R-Y. This numerical value was based on the requirement of a 95% chance of no core melt accidents in the entire lifespan of the nuclear fission industry in the U.S., which he conservatively assumed to be 100,000 reactor-years (an average of 300 reactors over an estimated 300 year life of the nuclear industry). However, Burns states in his paper[2] that his recommended numerical value of 5×10^{-7}/R-Y may be stricter than it needs to be in light of the WASH-1400 prediction that 1 in 50 core melt accidents is capable of causing early fatalities. Consequently, he suggests an alternative numerical value of the criterion as one accident in 40,000 reactor-years (i.e., 2.5×10^{-5}/R-Y) taking into account the reliability of containment systems and the small likelihood of violent steam and hydrogen explosion. Burns does not recommend this alternative value of the criterion unless it is shown that the recommended value of 5×10^{-7}/R-Y is

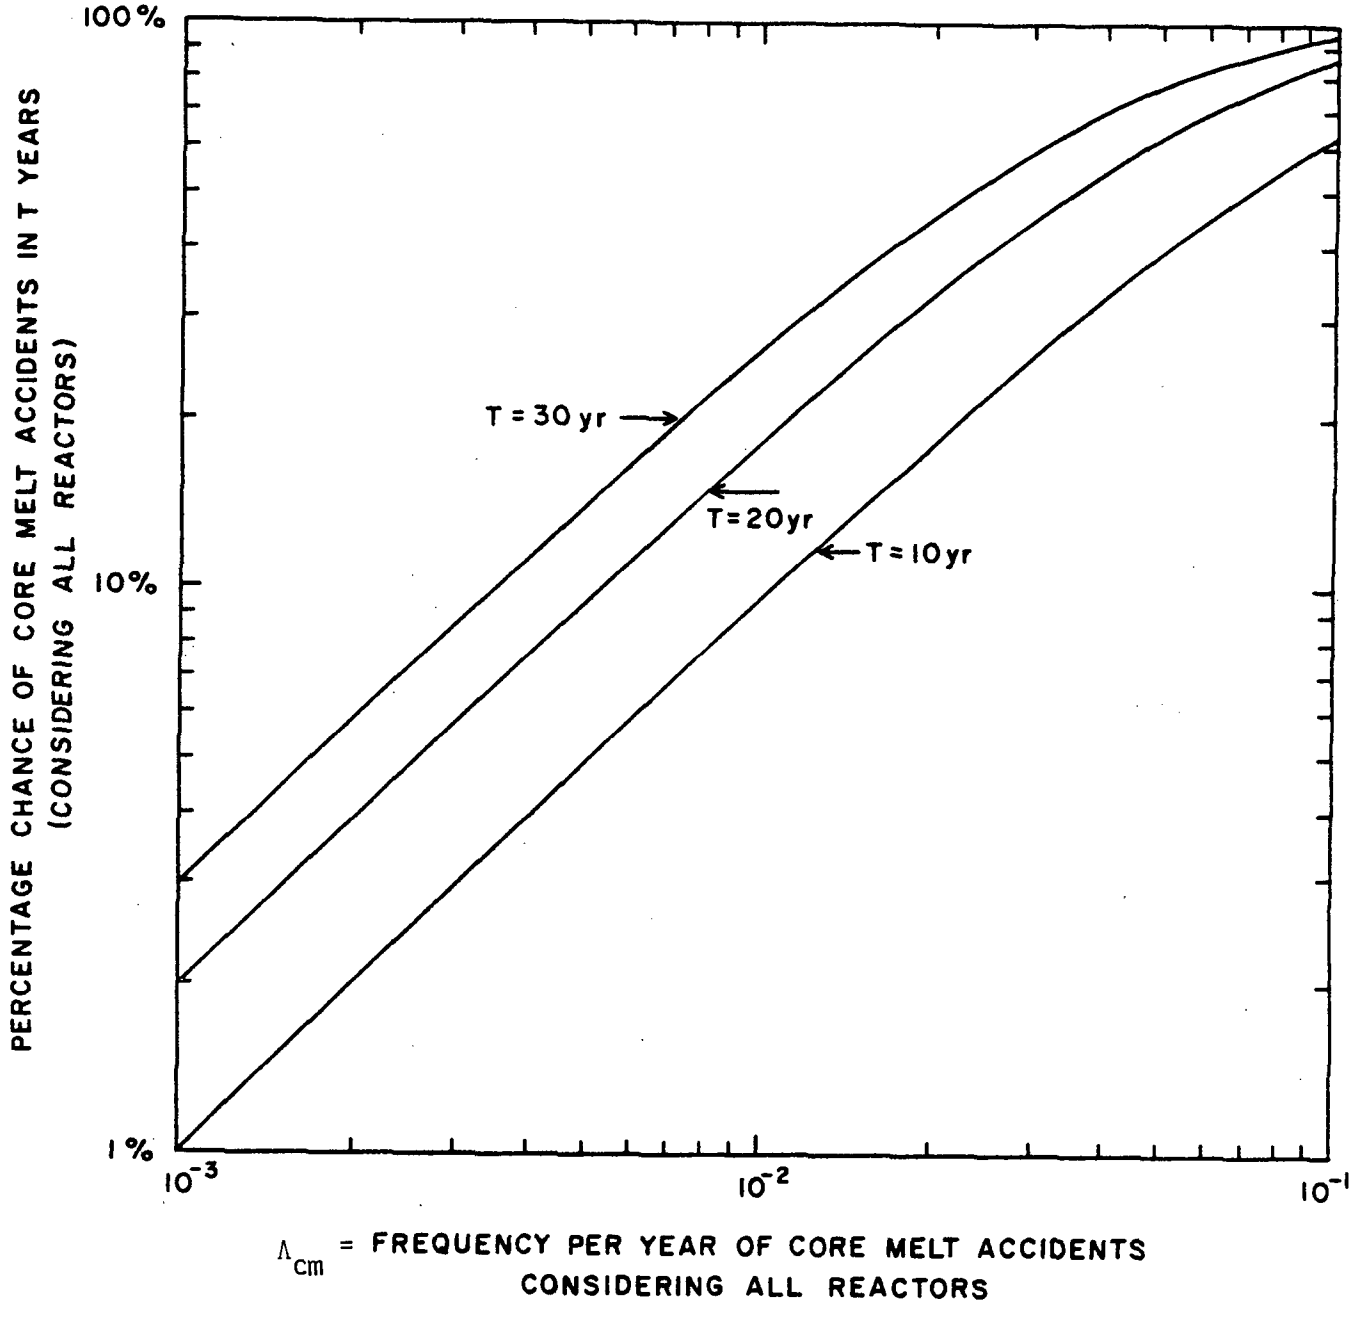

Fig. 6.3 Allowed core melt probabilities over the time periods 10, 20, and 30 years as a function of the values of the Y type of core melt probability criterion.

not achievable. As an implication, it may be observed from Figure 6.2 that Burns' recommended value allows about a 0.05% chance of core melt accidents in the next decade.

Wall[3] of the Electric Power Research Institute (EPRI) proposed a criterion for the frequency of core melt accidents based principally on three considerations. They are:

(i) frequency of core melt accidents for current U.S. LWRs is within the range of 4×10^{-6} to 6×10^{-4}/R-Y. (He obtained this range by assuming an error factor of about 12 in relation to the "best estimate" value of the frequency of core melt accidents, equal to 5×10^{-5}/R-Y, that was assessed in WASH-1400);

(ii) improvement of core melt frequency by a factor of 10 or more over what currently exists is not economically worthwhile;

(iii) risks from nuclear power plants as assessed in WASH-1400 are lower compared to the risks from alternate sources of generating electricity.

Based on the above considerations, Wall proposed that a probablistic goal in the range of 10^{-5}/R-Y should be acceptable. However, he stated that for less severe accidents higher values of the goal should be acceptable. As a tentative guideline for the frequency of accidents involving significant cladding degradation (e.g., Three Mile Island type accidents), Wall suggested a frequency value of 10^{-4}/R-Y.[4] From Figure 6.2, it may be observed that Wall's goal for core melt accidents allows about a 1% chance of occurrence in the next decade.

Vesely[5] has suggested two types of criteria for the frequency of core melt accidents. One criterion focuses on the design of the plant as it affects the frequency of core damage accidents while the other combines these same aspects of design with the manner in which a plant is operated. Accordingly, they are called design criterion and operational-cum-design criterion. Each of these criteria is divided into two levels, an unacceptable level and a warning range. The first level of the operational-cum-design criterion considers a point estimate of the frequency of core damage accidents, λ_{cd}, for a plant in excess of 10^{-3}/R-Y to be unacceptable. The second level defines a warning range of $10^{-4} \leq \lambda_{cd} \leq 10^{-3}$, where an outright decision on unacceptability is not made. In this case, the decision would be made on a close scrutiny of the merits and demerits of the individual plant in question.

The warning range is intended to account for the calculational uncertainties that are inherent in the evaluation of the frequency of core damage accidents. For the purpose of deciding the adequacy of a design, an assessment, based only on hardware, which includes a specific set of accident initiating events, system failures, and component failures would also be judged on a two-level basis:

- unacceptable level

$$\lambda_{cd}(H) > 1 \times 10^{-5}/R\text{-}Y$$

- warning range

$$1 \times 10^{-6}/R\text{-}Y \leq \lambda_{cd}(H) \leq 1 \times 10^{-5}/R\text{-}Y$$

In comparison, the assessment to be performed to show both design and operational safety adequacies includes hardware failures, testing contribution, human errors, and common mode failures.

In order to determine whether or not the above criteria can be met, the Surry plant was used as a test case. Calculations based on WASH-1400 have shown that the hardware contribution to the frequency of core melt accidents is near the unacceptable level specified by the design criterion. However, with respect to the operational-cum-design criterion, it is below the warning range.

The Director of Nuclear Safety Analysis Center, E.L. Zebroski, proposed the following formulation of a National Nuclear Safety Goal[6]:

- Considering the actual population of civilian reactors in the U.S., accidents which reach the state of core melt should have a probable frequency of no more than one such occurrence in 30 years.

- Reactor safety systems and containments shall be maintained and operated so that even if core melt occurs there should be less than one chance in 1000 that there is a radiation release which leads to a dose of 1R or more to any member of the public.

From the above formulation, it is interpreted that Zebroski's criterion specifies a limit on the total frequency per year of core melt accidents (equal to 1/30 per year) considering all reactors operating in the U.S. in any given year. Based on this interpretation it may be observed from Figure 6.3 that the limit specified by Zebroski allows a core melt probability no higher

than 25% over a period of 10 years. Since Zebroski's safety goal for the frequency of core melt accidents is dependent on the actual population of reactors, the safety goal is implicit for any given reactor in any given year.

Let the frequency of core melt accidents for a given reactor in a given year be denoted by λ_{cm}. Therefore, different values of the safety goal for $_{cm}$ would be obtained by considering different sizes of reactor populations. A minimum value of the safety goal for core melt accidents is obtained by assuming that no new reactors would be put into operation for the next 30 years and that the present population of 65 operating reactors would stay on line. Under this assumption, the number of reactor years that would accumulate in the next 30 years is no higher than 1950. Furthermore, if it is assumed that all reactors have the same goal, then a value of the safety goal for λ_{cm} of about 5×10^{-4}/R-Y is obtained. If some growth in nuclear power over the next 30 years is considered, this would lead to values of the safety goal for λ_{cm} lower than 5×10^{-4}/R-Y. Since considerable uncertainties exist in the projected growth of nuclear power in the next 30 years, values of the safety goal for $_{cm}$ inferred from the projected growth in the next 10 years might be more relevant. These values of the safety goal for λ_{cm} are all about 3×10^{-4}/R-Y based on high, low, and medium projected growth rates in nuclear power over the next 10 years that are shown in Figure 6.1.

6.5 ALLOWABLE TIME PERIOD FOR CORRECTIVE ACTIONS AND DERIVATION OF SHORT TERM GOALS

An issue related to the implementation of a standard for the frequency of severe core damage accidents from the point of view of a regulatory agency is the development of guidelines for allowable time periods within which corrective actions assuring compliance with the standard should be completed. In general, the allowable time periods should be such that the probability of core damage accidents within this period is small. Accordingly, the allowable time period should be shorter for a plant which is assessed to have a high frequency of core damage accidents compared to a plant with a lower assessed frequency. However, if an old plant is found to have a high frequency, then

corrective actions may not be necessary because the chance of severe accidents within the few years of residual plant lifetime may be small enough to warrant continued operation without extensive remedial actions. On the other hand, in the case of a new plant with a frequency as high as the aforementioned older plant, corrective actions may be required because the chance of an accident is greater over the longer residual plant lifetime.

Rowsome[7] has provided an initial formulation for inferring backfit deadlines based on the following hypothesis:

"We might accept up to a 0.1% chance/unit of a significant accident, i.e. one of the seriousness of TMI or worse, in the interval between the discovery of vulnerability to a short term fix, another 0.1% chance while we decide upon a midrange policy on backfits, and a third 0.1% chance for the rest of the service life of a unit."

The implications and ramifications of the above hypothesis will be investigated in the rest of this section to illustrate an approach which provides a guideline for allowable time periods for corrective actions. Prior to the discussion a few variables need to be defined. These are:

$\lambda_{cd}(D)$ = significant accident frequency per reactor-year assessed for a plant. (The discovery of this assessment can take place at any time within the service life of the plant which is assumed to be 40 years).

$\lambda_{cd}(G)$ = short term fix goal of significant accident frequency per reactor-year. (This goal is implicit in Rowsome's hypothesis. It is assumed here that in the application of this hypothesis an explicit goal would be specified. This goal is to be achieved by corrective actions which are to be completed within a prescribed time period, t_s.)

t_d = time in years from start of reactor operation when $\lambda_{cd}(D)$ is discovered, $0 \leq t_d < 40$.

t_s = allowable time period in years between time of discovery and completion of corrective actions.

Rowsome's inferred allowable time periods for corrective actions, t_s, for different values of discovered frequency of significant accidents, $\lambda_{cd}(D)$ are obtained from the following inequality:

$$t_s \leq \frac{\ln (0.999)}{\lambda_{cd}(D)}$$

The value of t_s determined from the above expression assures that the chance of a significant accident in a plant until the time of completion of corrective actions is 0.1%. It is based on the assumption that the assessed frequency, $\lambda_{cd}(D)$, is constant over the allowable time period, t_s.

From a plot of t_s versus $\lambda_{cd}(D)$ shown in Fig. 6.4, it is seen that for:

$$\lambda_{cd}(D) = 10^{-4}/\text{R-Y}, \; t_s \leq 10 \text{ years}$$

$$\lambda_{cd}(D) = 10^{-3}/\text{R-Y}, \; t_s \leq 1 \text{ year}$$

$$\lambda_{cd}(D) = 10^{-2}/\text{R-Y}, \; t_s \leq 0.1 \text{ years}$$

Thus, the higher the discovered value of the frequency, the shorter is the allowed fixing period.

If the discovered value of frequency $\lambda_{cd}(D)$, is less than or equal to $2.5\text{x}10^{-5}$/R-Y, the allowed time period exceeds the 40 year service life of a plant. Consequently, $\lambda_{cd}(D) = 2.5\text{x}10^{-5}$/R-Y, is the minimum discovered frequency for a significant accident that would require corrective actions. This threshold value of $\lambda_{cd}(D) = 2.5\text{x}10^{-5}$/R-Y which calls for corrective action is applicable only in cases when the discovery of $\lambda_{cd}(D)$ is made at the start of the service life of a plant. If the discovery is made at any point in the service life of a plant other than at the start, then the threshold value of the frequency which requires fixing is greater than $2.5\text{x}10^{-5}$/R-Y. This implication is illustrated in Fig. 6.5, from which it may be observed that a plant assessed to have $\lambda_{cd}(D) = 2\text{x}10^{-4}$/R-Y at 35 years into its service life does not require fixing, compared to an assessed value of $\lambda_{cd}(D) = 2.5\text{x}10^{-5}$/R-Y, if the discovery had been made at the start of

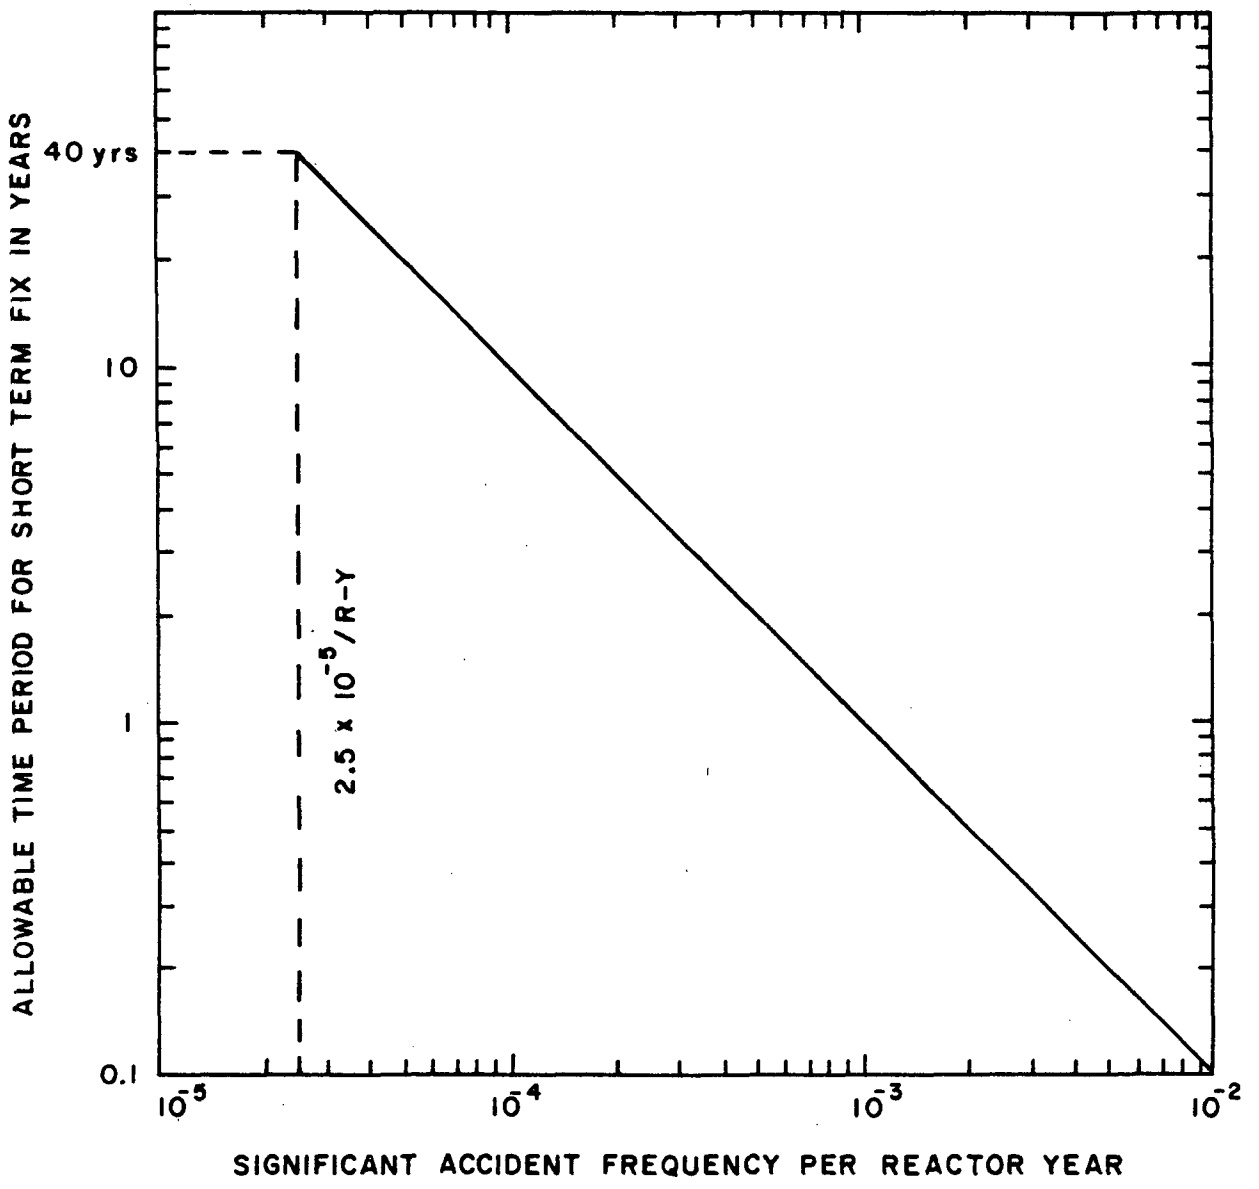

Fig. 6.4 Time allowed for short term fix as a function of the discovered value of the frequency of significant accidents.

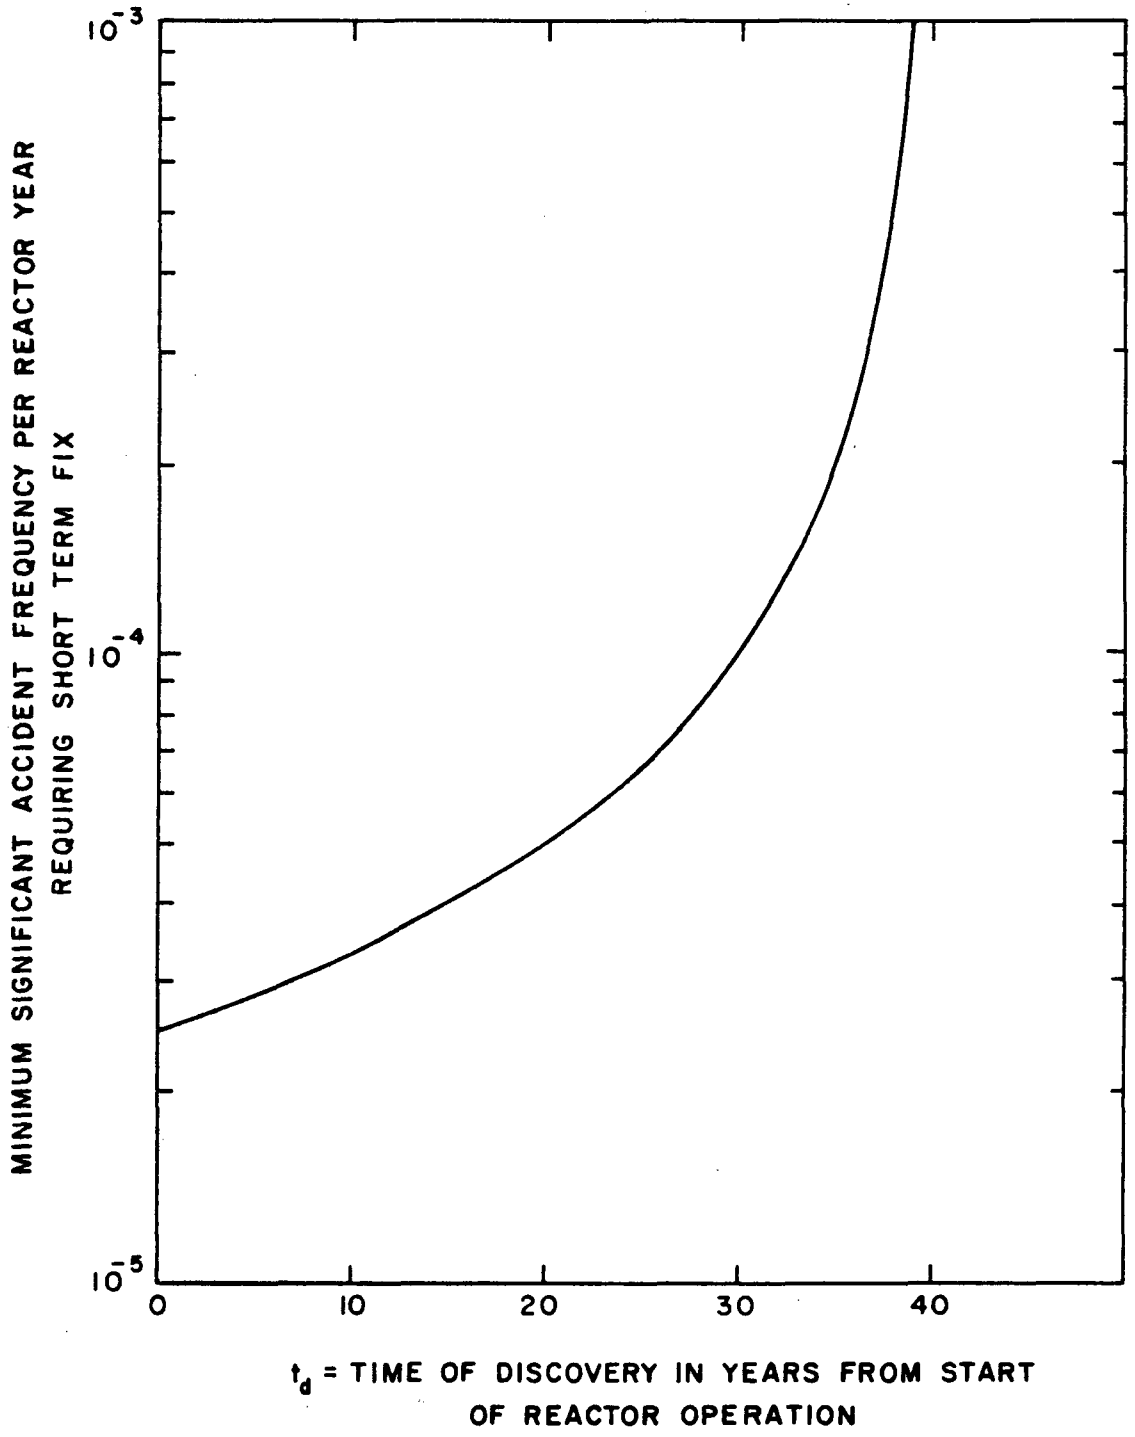

Fig. 6.5 Minimum discovered value of the frequency of significant accidents requiring short term fix as a function of the time of discovery.

plant operations. Although this prevalence of a seemingly double standard to resolve the issue of "fixing" versus "not fixing" may appear disconcerting to some, it stands to reason because an old plant which has few years remaining in its service life can live out its life without exceeding the 0.1% chance of a significant accident.

It has been stated earlier that the implementation of the approach under discussion would require the specification of the short term fix goal ($\lambda_{cd}(G)$) for the frequency of a significant accident so that corrective actions can be aimed towards meeting the goal. Rowsome's hypothesis does not provide explicit goals that are to be achieved at the end of the allowable time period for corrective actions. However, short term goals, $\lambda_{cd}(G)$, are implied by the hypothesis which states that the chance of a significant accident in a plant up to 0.2% in the time interval between the completion of the short term fix and the end of plant life might be acceptable. The value of $\lambda_{cd}(G)$, for a plant which is assessed at t_d years from initial operation as having a frequency of a significant accident equal to $\lambda_{cd}(D)$ may be obtained from the following expression

$$\lambda_{cd}(G) \leq \frac{-\ln (0.998)}{40-(t_d+t_s)}$$

where, t_s = allowable time period for corrective actions.

$$= \frac{-\ln (0.999)}{\lambda_{cd}(D)}$$

The explicit values of $\lambda_{cd}(G)$ given by the above expression are based on the assumption that no additional goals other than $\lambda_{cd}(G)$ are established in the time interval between the completion of short term fix and the end of plant life. In other words, a plant which has achieved the short term goal level of $\lambda_{cd}(G)$ continues to operate at this constant level till the end of its life. From the above expression it may be observed that the inferred level of the goal, $\lambda_{cd}(G)$, depends on the assessed frequency, $\lambda_{cd}(D)$, and the elapsed time, t_d, measured from the start of plant operation when $\lambda_{cd}(D)$ was assessed. The level of the goal becomes less stringent as the time of discovery, t_d, is delayed for a given value of $\lambda_{cd}(D)$. On the other hand, for a given value of t_d, the goal level decreases with the

decrease in the discovered value of the frequency of significant accidents. A plot of the short term fix goal as a function of the sum of the time of discovery and the allowable time period for corrective action is shown in Fig. 6.6

Figures 6.4, 6.5, and 6.6 may be utilized to determine whether or not short term fix is necessary, if so, what the values are for the short term fix goal and the allowable time period for achieving this goal. Consider, as an example, that a discovery of $\lambda_{cd}(D) = 10^{-4}$/R-Y was made at the start of a plant operation, i.e., $t_d = 0$ years. From Fig. 6.5, it is observed that short term fix is required; from Figs. 6.4 and 6.6, it is determined that the allowable fixing period is 10 years within which the frequency of significant accidents is to be reduced from 10^{-4}/R-Y to 6×10^{-5}/R-Y. However, if the discovery was made 10 years into the service life of the plant instead of at the beginning, the short term fix goal becomes 10^{-4}/R-Y, which is equal to the original discovered value of the frequency of significant accidents. If the discovery is delayed for 20 years, the goal becomes 2×10^{-4}/R-Y, a factor of 2 higher than the original discovered value. Thus, from these examples, it appears that use of Rowesome's hypothesis to simultaneously determine explicit values of the goal level and the allowable time period for corrective action under the stated assumptions could give rise to anomalous situations where the short term fix goals are less stringent than the original assessments.

Fig. 6.7 provides the short term fix goals for the frequency of significant accidents and the allowable time periods for completion of short term corrective actions as a function of the assessed frequency of significant accidents for a new plant where the time of discovery, $t_d = 0$. From this figure it may be observed that if the assessed frequencies lie between 10^{-2} to 10^{-4}/R-Y the short term fix goals are nearly constant at an approximate level of 6×10^{-5}/R-Y to be achieved in a period ranging from 0.1 to 10 years. However, if the assessed frequency lies between 7.5×10^{-5}/R-Y and 2.5×10^{-5}/R-Y, the goal values are lower than the assessed values. Any frequency which is assessed to be lower than 2.5×10^{-5}/R-Y does not matter since the allowable time period for corrective actions exceeds the service life of the plant. These anomalous cases can easily be corrected by imposing

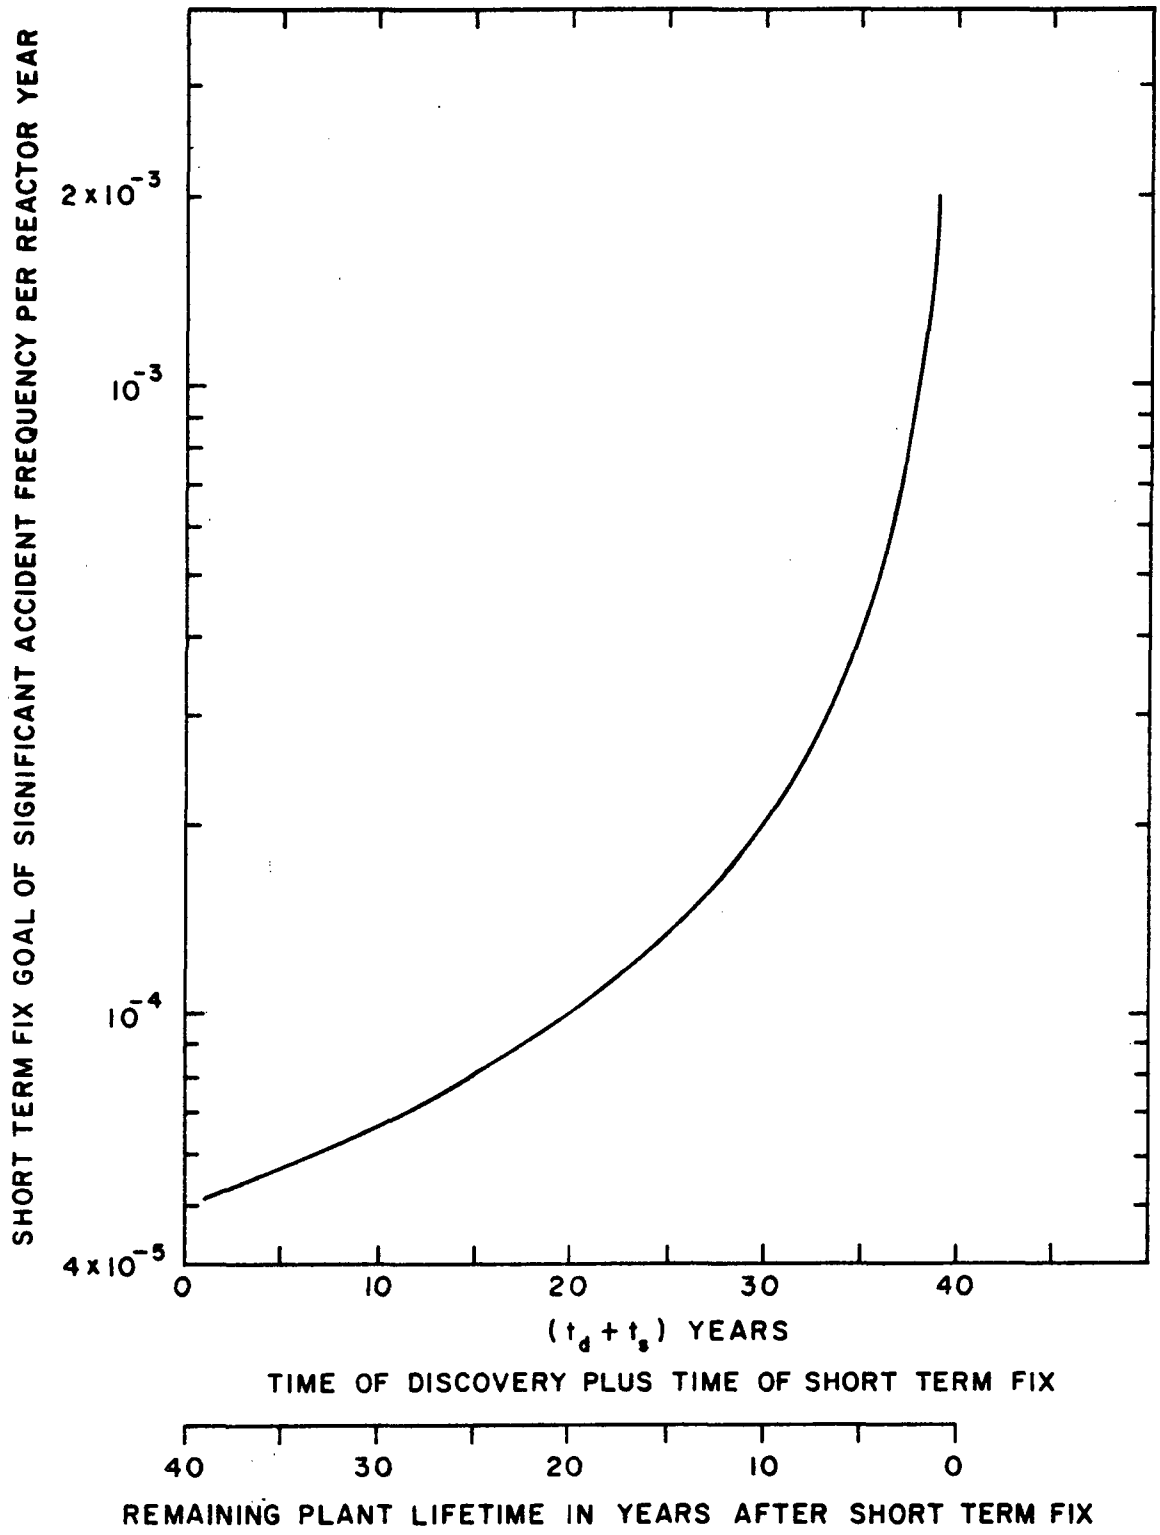

Fig. 6.6 Inferred short term fix goal for the frequency of significant accidents as a function of the sum of time of discovery and the allowed time period for short term fix.

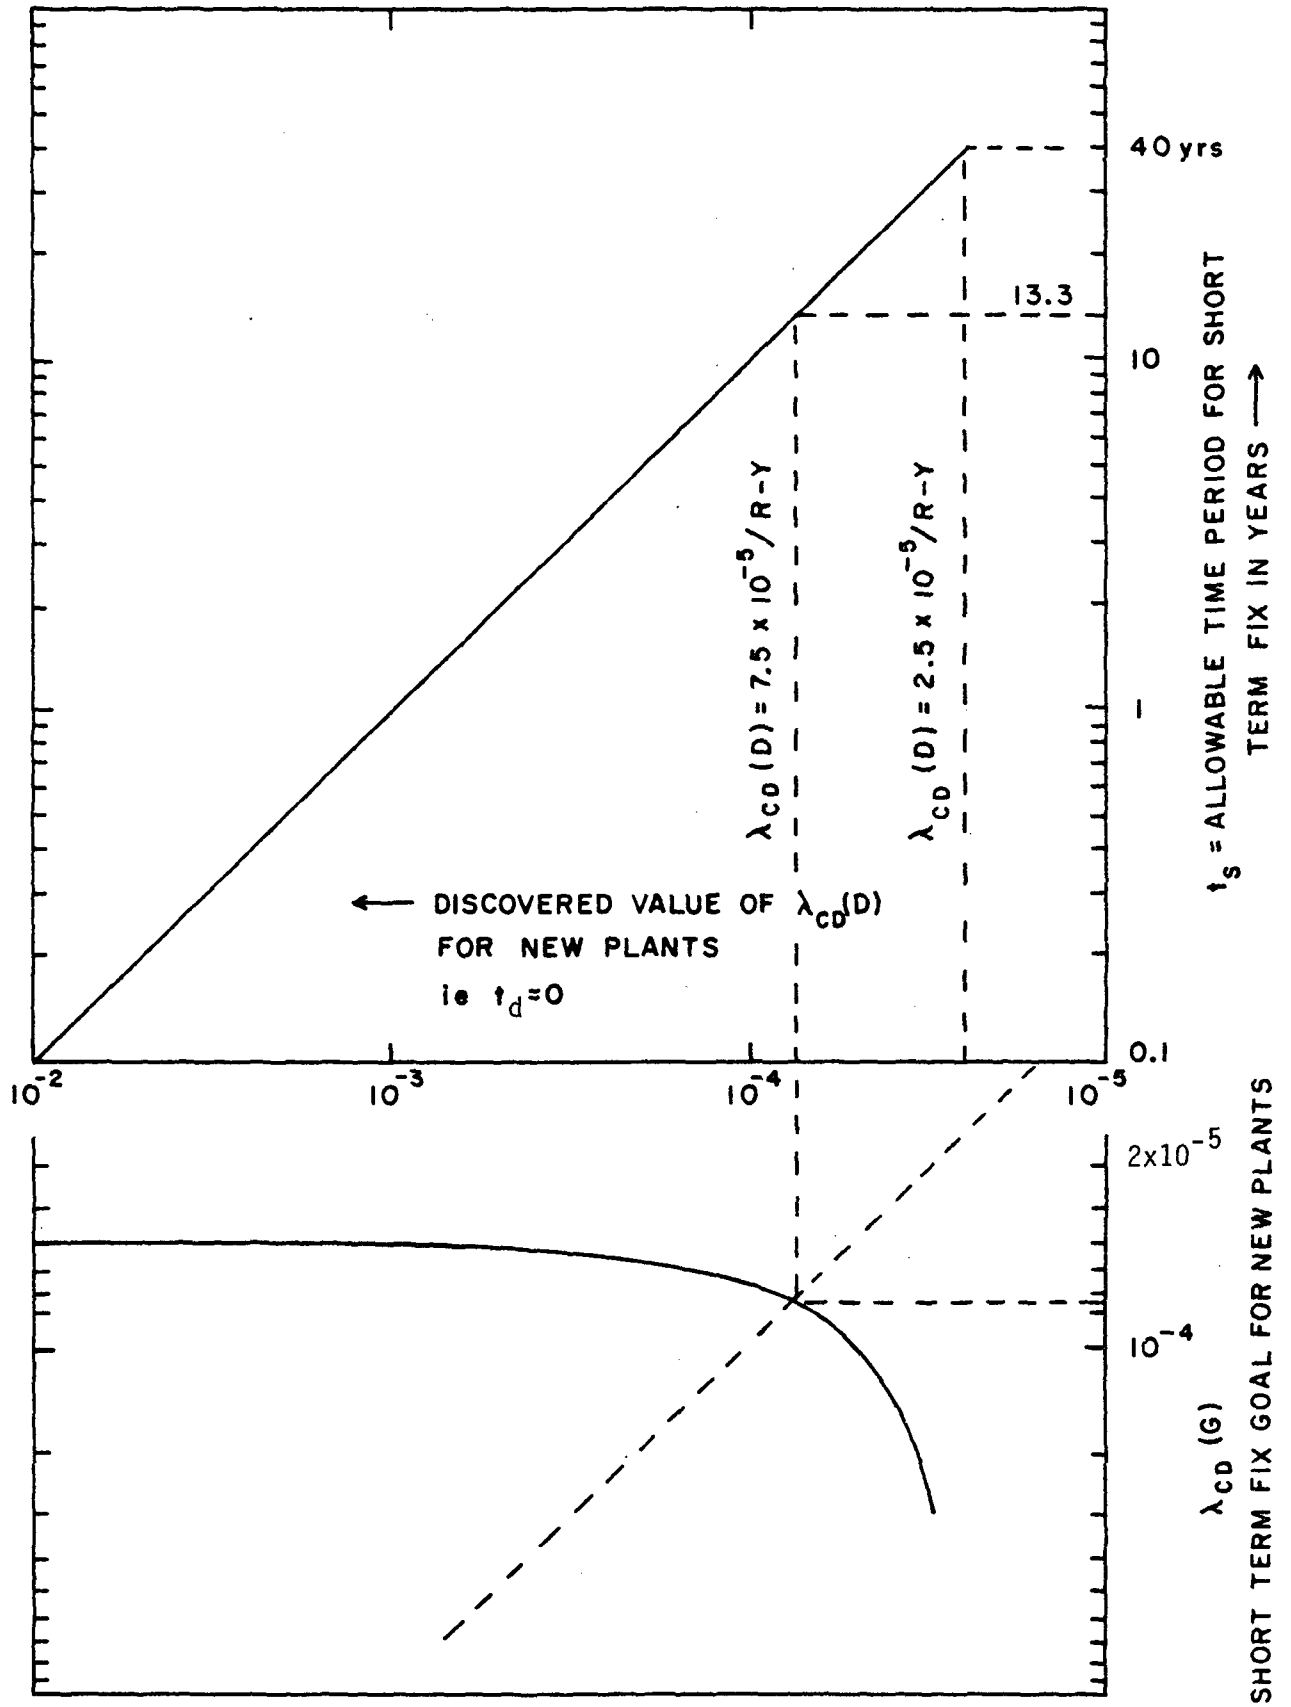

Fig. 6.7 Time allowed for short term fix and inferred short term fix goals for new plants.

additional constraints such as specifying additional goals to be achieved after the short-term fix goals are met or by specifying the short term goals independently of the discovered frequency and at the time of discovery. On the whole, Rowesome's conceptual framework for deriving allowable time periods for corrective actions possesses attractive features, and it is recommended that further work be done based on his approach.

Chapter 6 References

1. "Reactor Safety Study, An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants," Appendix VI, USNRC, WASH-1400, NUREG-75/014, October 1975.

2. Burns, R.D., "Proposed Criteria for Risk Assessment," Energy Division, Los Alamos Scientific Laboratory, (in response to a letter from W.E. Vesely of the USNRC), 1979.

3. Wall, I.B., "Probabilistic Risk Assessment in Nuclear Power Plant Regulation," Paper presented at the Second International Seminar, Structural Reliability of Mechanical Components & Subassemblies of Nuclear Power Plants, Berlin, West Germany, August 1979.

4. Wall, I.B., "What ATWS Risk Would be Cause for Concern?" Presented to ACRS ATWS-Subcommittee.

5. Vesely, W.E., "Numerical Risk Criteria," Office of Nuclear Regulatory Research, Presented to the ACRS RPA Subcommittee meeting held on April 30, 1980.

6. Zebroski, E.L., "A Proposed National Nuclear Safety Goal," Presented at the Seventh Energy Technology Conference in Washington, D.C., March 24-26, 1980.

7. Memorandum from F.H. Rowsome, NRC, to R.M. Bernero, NRC, "Backfit Deadlines Based on Hypothetical Interim Acceptable Risk Numbers," July 11, 1980.

7. RELEASE CRITERIA

7.1 INTRODUCTION

This chapter is concerned specifically with the release criteria. These criteria focus on the frequencies of various amounts of different radioisotopes that could be released to the environment from accidents in a nuclear power plant. These criteria are placed in the hierarchical structure for risk criteria at a level that is above the accident probability criteria but below the top level risk number criteria (i.e. the societal and individual risk criteria). The term release criterion, as used in this report, is probabilistic in nature because it takes into account the frequencies of different amounts of releases. In the preceding three chapters, types of criteria were reviewed that directly controlled only the probabilities of certain kinds of events regardless of their consequences. These events were component failures, system failures, and accidents. In contrast, the release criteria control explicitly the probabilities of events with regard to their consequences in terms of the amounts of radioactive material that could be released to the environment from a particular plant in a given period of time.

7.2 PROPERTIES OF RELEASE CRITERIA

A release criterion is not concerned with the availabilities of components and systems, the frequencies of accidents, or the integrity of the containment under different accident conditions as long as the amounts of radioactive releases and their associated frequencies are not above some unacceptable criterion level. This level is determined by a number or a set of numbers. The release criteria address only two characteristics associated with the releases of a particular radioisotope, the frequencies and the magnitudes. Other characteristics associated with an environmental release (e.g., duration, elevation, enthalpy, etc.) are not addressed by the release criteria, although, these characteristics act in conjunction with some site specific features (e.g., meteorology) to affect public risk.

A release criterion when applied to a particular plant attempts to judge the engineered safety built into the plant regardless of its site specific features such as the surrounding population density and distribution, meteorology, and evacuation measures. Although, these site specific features affect public risk, a release criterion attempts to control public risk by concentrating on its source (i.e., the amounts of radioactive releases to the environment and their associated frequencies). The implication of a release criterion in terms of its ability to control site specific societal risk is exemplified by considering the variation in the estimated societal risk of early fatalities from two plants at different sites that have identical releases. If one plant is located in an area where the population is higher than the other, then the site specific societal risk of early fatalities, in particular, is higher for the plant that is surrounded by the larger population. Both these plants may comply with a release criterion, however, the societal risk of early fatalities allowed by the release criterion is higher for the plant that is surrounded by the larger population. The site specific societal risk of latent fatalities is not as sensitive as the societal risk of early fatalities to the immediate population size around a plant site.

If the amounts of allowable releases of a radioisotope specified by a criterion were expressed in curies and this criterion was used to judge the engineered safety of all plants, then as such it would not account for the plant to plant variation in the inventories of different radioisotopes that is available for release in case of accidents. It is possible to overcome this weakness by expressing the amounts of allowable releases in a release criterion as fractions of a plant's radionuclide inventory that is available for release.

In the past, release criteria have been defined on the basis of airborne release of isotope iodine-131.[1, 2] The selection of a radioiodine isotope was based on the premise that large quantities of volatile radioiodine could be released from reactor accidents constituting the predominant hazard to the general public. In addition to iodine-131, other radionuclides may be released to the environment from an accident. Some of these radionuclides are more important than others from the standpoint of public health hazard. In theory, a set of release criteria could be defined,

each criterion constraining the release of a particular radionuclide which is capable of causing significant adverse health effects. However, if release criteria were established for a large number of radionuclides the use of such criteria could be cumbersome. In addition, a set of release criteria could lead to inflexibility and cause difficulties in making decisions on the safety of a plant based on the releases of various radionuclides. This is exemplified in a situation where the assessed releases (both in amount and frequency) of some radionuclides are much lower than what the criterion allows while for others they are higher; however, when all the pertinent radionuclides are considered the public risk may not be unacceptable. The difficulty in the demonstration of compliance with a release criterion based only on one radionuclide largely depends on the accuracy with which the predicted release magnitudes of this nuclide can be assessed. This difficulty is more pronounced if compliance is to be demonstrated independently for each of the several radionuclides since it assumes that reasonable estimates of release magnitudes of each radionuclide under various accident conditions can be obtained. At present, reasonable estimates of releases for each of these radionuclides may be unobtainable because of lack of data and realistic models for predicting release magnitudes under various accident conditions. In an attempt to account for pertinent radionuclides that could be released in an accident, instead of focusing only on iodine-131, and to alleviate some of the difficulties associated with establishing a set of release criteria (i.e., a criterion for each of the several radionuclides), a comprehensive release criterion could be established which might require that different radionuclides be weighted in some manner according to their individual health effects.* On the other hand, if the amount of iodine-131 released is considered to be a sufficient indicator of the severity of the accident being evaluated, then a criterion on radioiodine may be adequate.

It may be noted parenthetically that in the accident at Three Mile Island only 15 Ci[3] of iodine-131 is estimated to have been released to the environment. Recent studies[4, 5] have suggested that for accidents in

---

*Hall, R.E. et al, "A Risk Assessment of a Pressurized Water Reactor for Class 3-8 Accidents," BNL-NUREG-50950, October 1979.

light water reactors that are physically realizable and where water is available, the amount of radioiodine that could be released to the atmosphere is much smaller than what was previously calculated.

Since release criteria are concerned with the amounts of radioactivity released and their associated frequencies, the frequencies of accidents that result in no or small releases (e.g., the Browns Ferry Fire) are not directly controlled by these criteria. Nevertheless, release criteria would indirectly control the frequencies of these non-release accidents if they were associated with the same safety system failures or initiating events as are accidents which result in releases to the environment. The frequencies of these non-release accidents could be controlled by the accident probability criteria.

7.3 FORMS OF RELEASE CRITERIA AND THEIR IMPLICATIONS

In this section different ways of expressing a release criterion for controlling releases to the environment from a plant are presented and then these are examined with regard to their implications. The following three forms of release criteria will be examined:

1. As a number which specifies the expected amount of radioactive release per plant year of operation that is unacceptable.

2. As a limiting curve which specifies the amounts of radioactive releases and their associated frequencies that are unacceptable.

3. As a limiting complementary cumulative distribution function which specifies the frequencies of equalling or exceeding specific amounts of releases that are unacceptable.

All of the above forms of criteria can be defined for different radioisotopes or for an appropriate sum total of important radioisotopes. The amount of release can be expressed in curies or as a fraction of core inventory that is available for release due to accidents from a particular plant. To simplify the discussion on the forms of release criteria, it is assumed that these are defined on the basis of one radioisotope and the amount of release is measured in curies.

First Form of Criteria

Demonstration of compliance with a criterion expressed in this form is shown when the appropriate number evaluated for a plant is no greater than the number specified by a criterion. This is expressed mathematically as:

$$x = \sum_{i=1}^{n} f_i c_i \not> x^*$$

where,

x = the expected amount of release of the radioisotope of interest in curies per plant year of operation,

i = the index for an event that result in a release,

n = the total number of release events,

f_i = the estimated frequency per plant year associated with the i'th event,

c_i = the assessed amount of release in curies associated with the i'th event, and

x^* = the number specified by a criterion.

It is clear from the above expression that this form of criteria is not concerned with the frequency or the amount of release associated with any particular release event as long as the expected amount of release per plant year of operation considering all release events that are being evaluated is not above some unacceptable level. Since the average release per year is controlled by this form of criteria, it does not take into account the extra concerns associated with events that could result in large releases but have low frequencies of occurrence. Uncertainties in the evaluations of the frequencies (i.e., f_i's) and the amounts of releases (i.e., c_i's) can be addressed by interpreting the number specified by a criterion at various levels of confidence.

Second Form of Criteria

The second form of criteria will be discussed by considering, as an example, the release criterion that was proposed by Farmer.[6, 7]

Farmer proposed a release criterion by specifying upper limits for the amounts of releases of iodine-131 at ground level due to accidents as a function of their probabilities of occurrence. The intent of Farmer's approach is to control the risk of individual accidents by constraining simultaneously the amounts of releases and their associated frequencies. This is achieved by comparing the frequency per year of an accident and its associated iodine-131 release at ground level in curies for a given plant with Farmer's "limit line." The limit line serves as a criterion to judge the acceptability of risk of any type of accident. Fig. 7.1 shows three different Farmer types of limit lines. Although, these limit lines are represented as continuous curves, they are aimed at judging the acceptability of risk of an individual accident. The accident being evaluated in question is characterized by its frequency of occurrence and its release of iodine-131, which is represented by a point in the frequency/release plane of Fig. 7.1. If points representing accidents fall below the limit line, they are considered acceptable. Points falling above the line require that corrective actions be taken to reduce the magnitude of the release and/or the associated frequency to bring these points below the line.

The three limit lines shown in Fig. 7.1 have slopes of -1, -4/3, and -3/2 in a log-log plot for releases greater than 1000 curies. A slope of -1 reflects an inverse relationship between the frequency and the amount of release. Therefore, every point on a line with a slope equal to -1 denotes the same risk i.e., the product of the frequency and the release is the same for all points. This relationship does not hold for points on lines with slopes other than -1. Slopes smaller than -1, such as -4/3 and -3/2 account for the increased risk perceived by society for remote events resulting in large releases than frequent events resulting in small releases even if the same numerical risk is assessed for these events. Therefore, Farmer's limit lines with slopes smaller than -1 reflect societal aversion to events that result in large releases. Because interpolation of these limit lines to small releases, in the region of less than 1000 curies, would have allowed high frequencies, Farmer made the slope of these curves greater than -1 to control the frequencies of events that result in small releases. Accordingly, the limit lines specify that a release of 10 curies of iodine-131 from any plant should not occur with a frequency greater than 10^{-2} per year.

An attempt is made in Fig. 7.2 to show the ability of existing reactors to comply with Farmer's release criteria. This figure shows the three different limit lines in comparison to the frequency and the amount of environmental release of iodine-131 pertaining to the nine "release categories" for a PWR and five "release categories" of a BWR that were defined in WASH-1400.[8] A straight line fit in a log-log plot is used to represent the functional behavior of all the limit lines for releases between 10 and 1000 curies,[10] instead of a curved line as shown in Fig. 7.1. It should be pointed out that "release categories" do not represent individual accident sequences nor are all the releases at ground level. In addition, frequencies associated with release categories may have been over-estimated due to the smoothing technique employed in WASH-1400. In light of the accident at Three Mile Island, where only about 15 curies[3] of iodine-131 were released to the environment even though a 0.39[9] estimated core fraction inventory was released from the core suggests that perhaps the release fractions utilized in WASH-1400 were on the conservative side.

In mathematical terms, Farmer's limit lines are expressed by the following functions:

$$P(c)\begin{cases} = 0.1725\ c^{-0.746} & \text{(Ref. 10)} \qquad 10 \leq c < 10^3 \\ = c^{-1} \text{ or } 10c^{-4/3} \text{ or } 31.6c^{-3/2} & 10^3 \leq c \leq 10^7 \end{cases}$$

The dimension of P(c) is the frequency per year per event at c curies. If f(c) represents the frequency per year per curie release, then following the prescription in Ref. 7, a relationship between f(c) and P(c) is obtained as

$$f(c) = (\ln 10)^{-1}\ \frac{P(c)}{c}$$

By utilizing the function f(c) some implications of the limit lines in terms of certain risk measures are presented in Table 7.1.

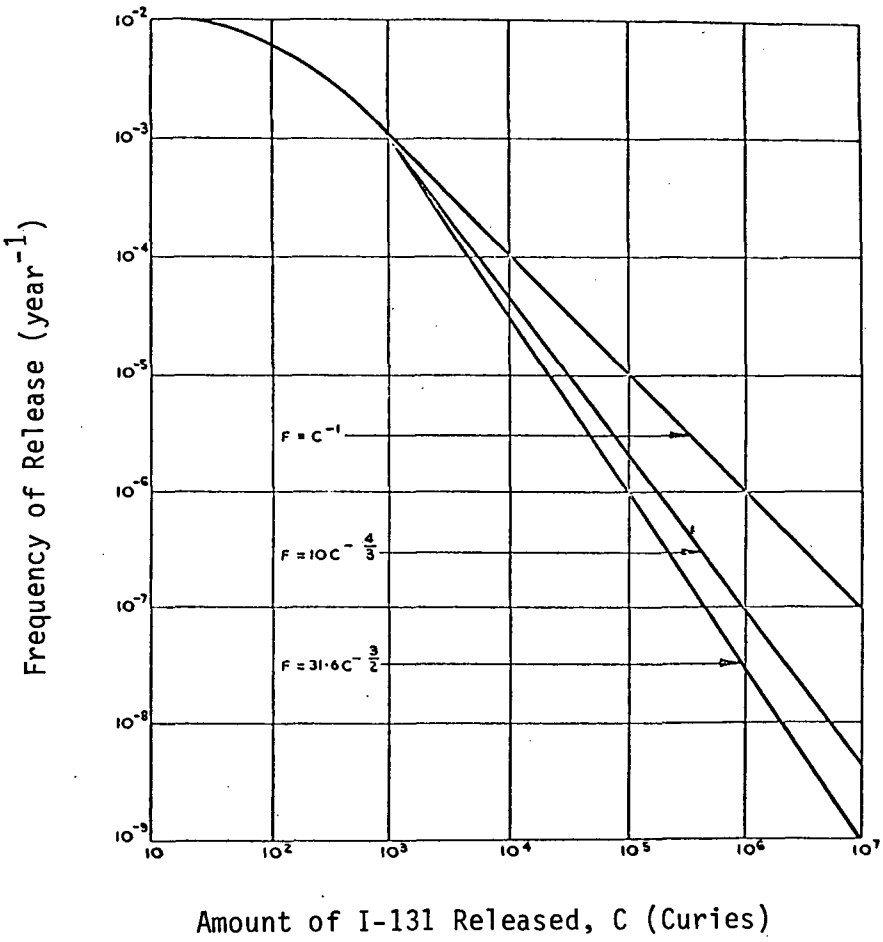

Fig. 7.1 Farmer's release frequency limit lines.

Table 7.1. Some Risk Measures of Farmer's Limit Lines in Terms of I-131 Release

Limit Line Slopes $10^3<c<10^7$	Frequency of Accidents with Releases > 10 Ci plant-year^{-1}	Average Release Given Accidents with Releases > 10 Ci $\bar{}$ [Ci]	Risk Per Annum ci/plant year
Slope = -1	1.8×10^{-2}	288	5.18
Slope = -4/3	1.8×10^{-2}	134	2.42
Slope = -3/2	1.8×10^{-2}	113	2.03

An advantage of expressing release criteria in the form of a limit line is that it provides a simple decision tool as to whether or not the risk of release from any accident is unacceptable. However, it does not constrain the risk from all accidents that are to be evaluated unless some assumptions are made regarding the form of the probability density function of accidents over the release range. The function f(c), as shown above, is derived by assuming a constant density function over the logarithmic scale of releases. In other words, the number of accidents which result in releases between 10^3 and 10^4 curies is the same as those which result in releases between 10^4 and 10^5 curies. The weakness of the limit line in constraining the total risk from all accidents can be illustrated by considering a cluster of points representing accident sequences in the accident frequency/release plane of Fig. 7.2, which may fall below a limit line and thereby comply with a release criterion. However, the total risk obtained by adding the risk from each accident might be considered unacceptable if there is a large number of points. This weakness can be corrected by specifying that frequencies of accidents that result in similar amounts of releases (e.g., accidents that result in releases within a decade) be summed to demonstrate compliance with a limit line.

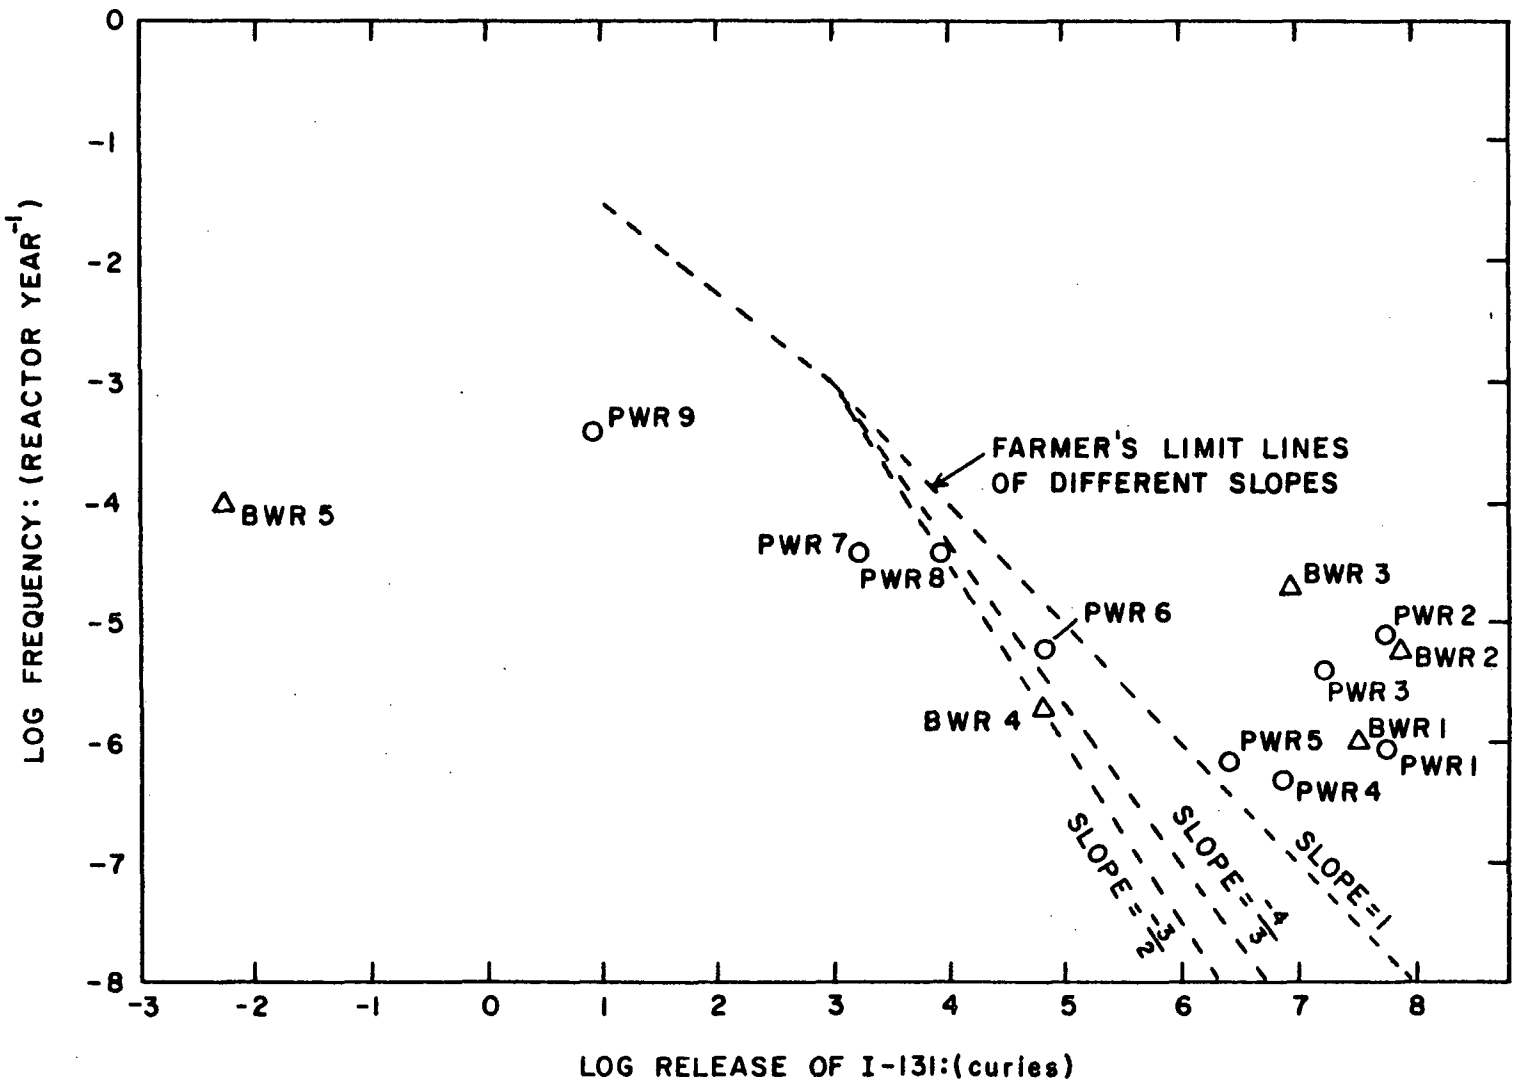

Fig. 7.2 Comparison of the frequencies and release magnitudes of PWR and BWR release categories of WASH-1400 with Farmer's limit lines.

In the preceeding discussions it was assumed that in order to show compliance with a limit line, an accident sequence is represented by a point which defines its frequency and the amount of release. However, there are uncertainties associated in the estimates of both the frequency and the amount of release. Therefore, the limit line has to be defined appropriately to account for these uncertainties.

Third Form of Criteria

A release criterion can also be expressed as a limiting complementary cumulative distribution function (CCDF), denoted by $F^c(c)$, which specifies the frequencies of equalling or exceeding certain amounts of releases that are unacceptable. Fig. 7.3 portrays three hypothetical release criteria expressed in the form of CCDFs in relation to the assessed CCDFs of a PWR and a BWR that were based on "release categories" of WASH-1400. These hypothetical release criteria, $F^c(c)$ as a function of c, were derived from limit lines shown in Fig. 7.2 by evaluating the following integral:

$$F^c(c) = \int_c^{c_{max}} f(c')dc'$$

$$= (\ln 10)^{-1} \int_c^{c_{max}} \frac{P(c')}{c'} dc'$$

where, c_{max} is assumed to be equal to 10^8 curies of iodine-131 and the functions f(c) and P(c) have been defined earlier. This form of criteria specifies that the evaluated amount of release equal to or greater than c(curies) should not occur with a frequency per plant-year greater than $F^c(c)$, for all c within a defined region.

Criteria expressed in the form of a CCDF have the ability, unlike the limit line, to constrain the risk of release from all accidents that are required to be evaluated because the area under the CCDF curve represents the risk of release in terms of curies per plant-year of operation. In addition, it is suitable for handling the uncertainty in the estimated amount of environmental release from any accident. A weakness of these forms of release criteria is that they lack the ability to identify a specific accident that may have a high risk of release.

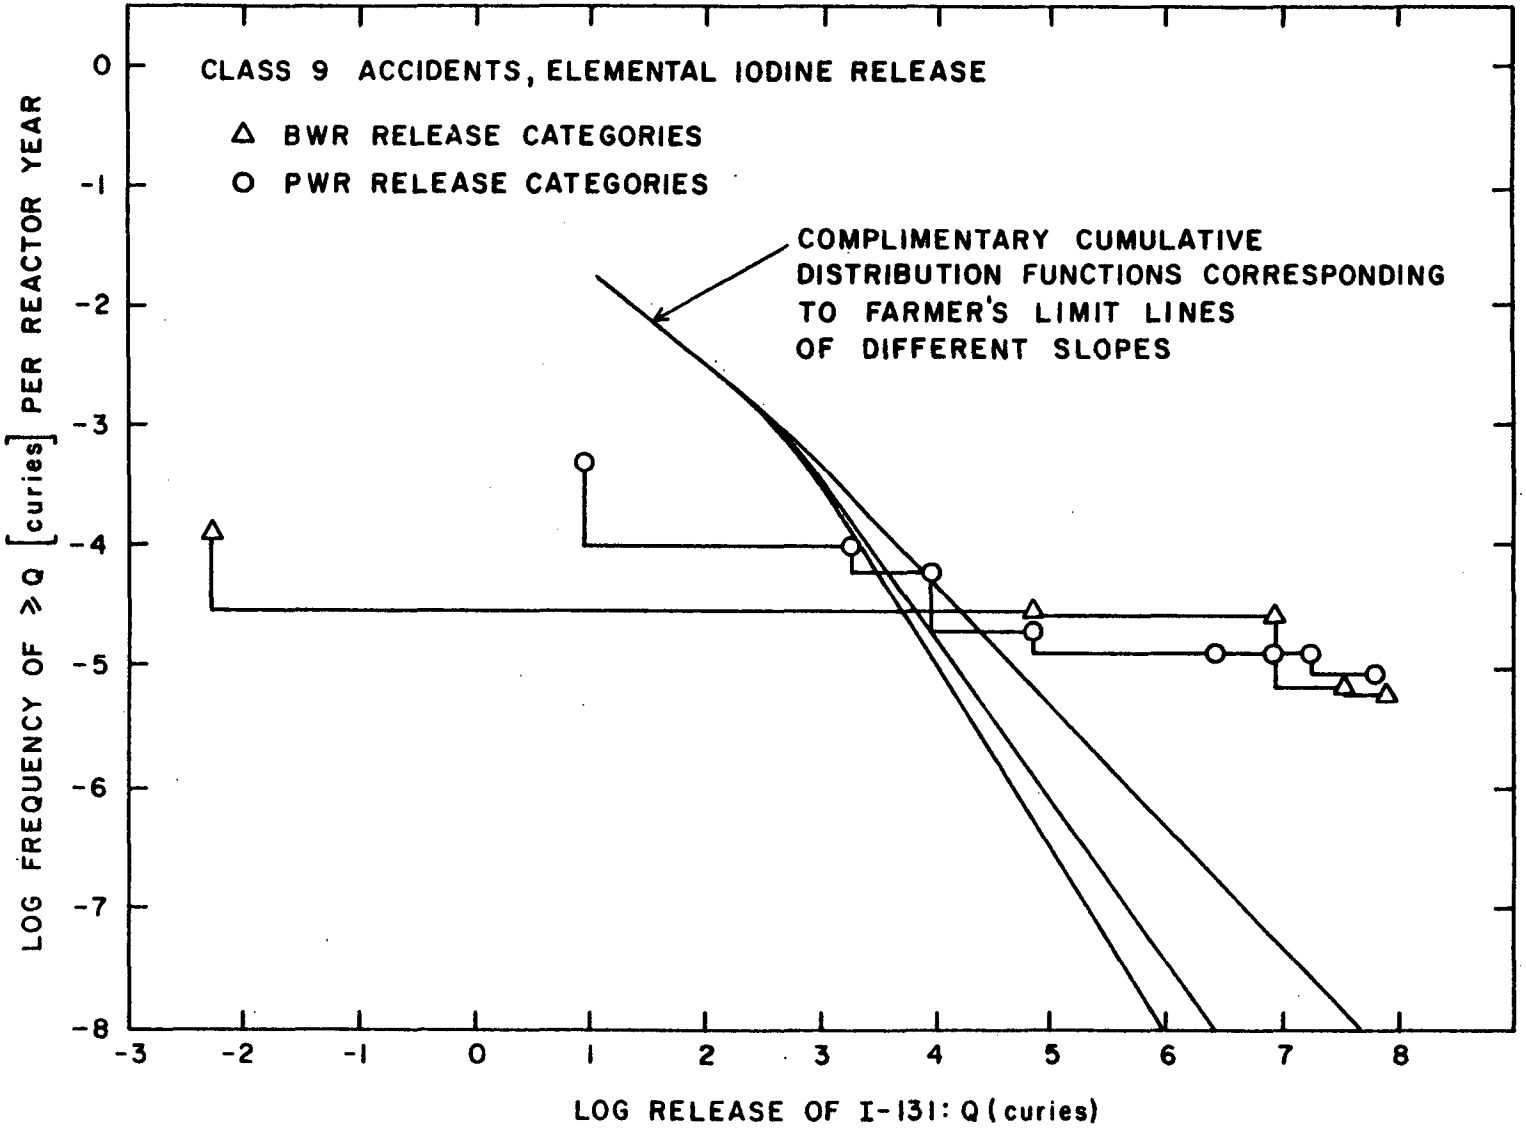

Fig. 7.3 Comparison of the complementary cumulative distribution functions (CCDFs) of BWR and PWR release categories of WASH-1400 with CCDFs of Farmer's limit lines.

Chapter 7 References

1. Farmer, F.R., "Reactor Safety and Siting: A Proposed Risk Criterion," Nuclear Safety, Vol, 8, No. 6, November - December 1967.

2. Meleis, M. and Erdmann, R., "The Development of Reactor Siting Criteria Based upon Risk Probability," Nuclear Safety, Vol, 13, No. 1, January - February 1972.

3. "Three Mile Island, A Report to the Commissioners and to the Public," NUREG/CR-1250, Vol, 1, 1980.*

4. Campbell, D.O., Malinauskas, A.P., and Stratton, W.R., "The Chemical Behavior of Fission Product Iodine in Light Water Reactor Accidents," (to be published in the May 1981 issue of Nuclear Technology).

5. Levenson, M. and Rahn, F., "Natural Limits on the Dispersal of Radioactivity in Nuclear Accidents," Electrical Power Research Institute (EPRI), November 1980, (Available from Nuclear Power Division, EPRI, 3412 Hillview Avenue, Palo Alto, California 94304).

6. Farmer, F.R., "Siting Criteria - A New Approach," Safety and Reliability Directorate, United Kingdom Atomic Energy Authority, Warrington (U.K.), SM-89/34, 1967.

7. Beattie, J.R., Bell, G.D., and Edwards, J.E., "Methods for the Evaluation of Risk," Safety and Reliability Directorate, United Kingdom Atomic Energy Authority, Warrington (U.K.), AHSB(S) R 159, 1969.

8. "Reactor Safety Study, An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants," Appendix V, USNRC, WASH-1400, NUREG-75/014, October 1975.**

9. "Three Mile Island, A Report to the Commissioners and to the Public," NUREG/CR-1250, Vol. 2, Part 2, 1980.*

10. Farmer, F.R., "Letters to the Editor," Nuclear Safety, Vol. 13, No. 5, September - October 1972.

*Available for purchase from the NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission, WAshington, DC 20555, and/or the National Technical Information Service, Springfield, VA 22161.

**Available free upon written request to the Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

8. INDIVIDUAL RISK CRITERIA

8.1 INTRODUCTION

These criteria specified by some number or set of numbers focus on the risk to an individual from accidents in nuclear power plants. The harm to an individual in the event of a nuclear accident may take various forms, e.g., early death, delayed death, morbidity, genetic effects, and economic loss. In this chapter, consideration of individual risk will be confined to the risk of death. As stated in Chapter 3, the individual risk criterion forms the top level of the hierarchical structure for risk criteria.

The purpose of an individual risk criterion is to ensure that any member of the public is not exposed to large accidental risks. In other words, the additional risk faced by an individual due to nuclear accidents should be small as compared to non-nuclear risks.

8.2 PROPERTIES AND FORMS OF THE INDIVIDUAL RISK CRITERIA

A criterion on individual risk is not concerned with the availabilities of systems, the frequencies of accidents, and the frequencies of various amounts of radioactive releases as long as the individual risk is not above some unacceptable level. Thus, an individual risk criterion gives more flexibility than other types of criteria (i.e., system availability criterion, accident probability criterion and release criterion) to designers and plant operators to achieve better safety and economy while still maintaining an overall control of the individual risk from nuclear accidents. However, it does not directly control the frequencies of accidents that are accompanied by very small or no environmental radioactive releases since they result in negligible risk to an individual.

An individual risk criterion, when formulated on the basis of a person located at some reference point associated with a plant site, focuses on the consequences and the frequencies of an assumed set of accidents and its impact

on the person's risk. It constrains only the risk to a hypothetical person regardless of the population size and distribution. Consequently, the risk to the population at large in the vicinity of a plant is not directly controlled by an individual risk criterion.

The desirability of an individual risk criterion can be illustrated by comparing the interplay between the societal and individual exposures to risk from accidents postulated to occur in two plants (say Plant A and Plant B) at two different sites. The societal risk of Plant A may be the same as Plant B, and both may comply with a site specific societal risk criterion. If this situation arises because Plant A is located in an area surrounded by a much lower population than Plant B and the radioactive releases (both in amount and frequency) are higher for Plant A than Plant B, then the estimated risks to persons located equidistant from each plant would be higher for Plant A than Plant B. In the absence of an individual risk criterion, the above situation may be allowed even though the individual risk is higher for Plant A. Therefore, use of a societal risk criterion without an individual risk criterion may lead to the construction of plants at remote locations with higher individual risk being permitted at a certain distance from a plant.

In the preceding paragraphs, a type of individual risk criterion that is specified on the basis of a person located at some reference point with respect to a plant site was discussed. Another type of individual risk criterion can be established. It constrains the risk to an average individual in the U.S. population at large from all nuclear power plants, existing or planned. To show compliance with this type of criterion requires that the assessed average individual risk from all plants operating at any time is not above the numerical value specified by the criterion. The assessed average individual risk can be obtained by dividing the sum of the societal risk from all plants by the appropriate population at risk. A disadvantage of this criterion is that it does not control the risk to a person who is exposed to a higher risk than the average individual.

Forms of the Individual Risk Criteria

A criterion for individual risk can be expressed in the form of limits on the magnitudes of the radiation doses to a defined individual and their associated frequencies. An example of such a criterion is the proposed

Canadian safety requirements for licensing of their CANDU type of nuclear reactors.[1] The numerical values of the Canadian individual risk criteria are shown in Table 8.1. The limits on the radiation doses are specified on the basis of the doses to the whole body and to the thyroid gland as they pertain to any offsite individual. An estimate of the average annual dose to the whole body allowed by this criterion for one reactor is about 9 mrem per year.* This estimate is obtained by assuming that the allowed doses to the whole body lie at the midpoint of the reference dose intervals as prescribed in Table 8.1. Using an approximate conversion factor of 10^{-4} cancer deaths per rem, the allowed annual individual risk of death is of the order of 10^{-6} per year.

Individual risk criterion can also be expressed in the form of a limit on the annual risk of death from accidents for a defined individual. Limits can be specified for the risk of early death and/or latent death. Such a criterion might read something like, "The risk of early death from accidents in a plant shall be no higher than 10^{-6} per year for a person located at the site boundary."

8.3 IMPLICATIONS OF A CRITERION FOR ANNUAL INDIVIDUAL RISK OF DEATH

The objective of this section is to present assessments of the following implications of a criterion for annual individual risk of death:

- lifetime risk of death of an individual
- loss of life expectancy.

The results of these assessments are presented in the form of parametric plots. These facilitate comparison of different numerical values which might be proposed for a criterion for individual risk of death with regard to the above implications. The implications assessed in this section are pertinent only to an individual to whom the criterion applies.

---

*In reality, the average annual dose allowed is less than 9 mrem because the criterion requires that compliance be shown for the case of serious process system failures without taking into account the effects of safety systems.

TABLE 8.1. PROPOSED SAFETY REQUIREMENTS FOR LICENSING OF CANDU NUCLEAR POWER PLANTS[1]

Table 8.1a. Proposed Reference Values

Reference Dose Interval (Sv*)		Reference Value for the Sum of the Predicted Rates of Occurrence of Failures** within the corresponding Reference Dose Interval
Whole Body	Thyroid	(Per Reactor Unit Per Annum)
0-0.0005	0-0.005	10^{-1}
0.0005-0.005	0.005-0.05	10^{-2}
0.005-0.05	0.05-0.5	10^{-3}

Table 8.1b. Proposed Reference Values

Reference Dose Interval (Sv*)		Reference Value for the Sum of the Predicted Rates of Occurrence of Failures** within the corresponding Reference Dose Interval
Whole Body	Thyroid	(Per Reactor Unit Per Annum)
0.05-0.1	0.50-1.0	10^{-4}
0.1-0.3	1.0-3.0	10^{-5}
0.3-1.0	3.0-10.0	10^{-6}

* 1 Sv = 100 rem

**In Table 8.1a, "failures" refers mainly to Serious Process Failures, defined as those failures which , in the absence of any Special Safety System action, could lead to exposure of an individual off-site to a radiation dose greater than 0.0005 Sv.

In Table 8.1b, "failures" refers mainly to Serious Process Failures combined with the inability of any one of the Special Safety Systems to perform its function.

Lifetime Risk of Death of an Individual

A criterion for individual risk of death allows some probability of death over the lifetime of an individual. The lifetime individual risk of death as a function of some values of individual risk of death per year is shown in Figure 8.1. The lifetime risk is calculated by assuming that an individual is exposed at a constant level of risk (as specified by a criterion) over his or her entire lifetime. The individual lifetime assumed for this calculation is 72.8 years. This is the average life expectancy at birth in the U.S. estimated for the year 1976.[2] From Figure 8.1 it may be observed that if the value of an individual risk criterion is, for example, 10^{-5} per year, the allowed lifetime risk is 7×10^{-4}.

Loss of Life Expectancy

A criterion for the individual risk of death per year may be interpreted as a limit on the additional mortality rate of an individual due to nuclear accidents over and above the normal mortality. In the context of competing risk analysis, the additional mortality rate will decrease the life expectancy of an individual. Thus, a criterion which allows some additional mortality rate can be expressed in terms of the loss of life expectancy. The expected length of life due to the additional risk allowed by a criterion is given by the following expression,[3]

$$\int_0^\infty \exp\,[-\lambda y]\,[1-G(y)] \approx \mu - \frac{\lambda K}{2}$$

where

λ = mortality rate from accidents allowed by the individual risk criterion. (Since the value of the criterion is the same for any year, the allowed mortality rate is considered to remain constant over the lifetime of an individual.)

Y = random variable associated with average life expectancy

$G(y)$ = cumulative probability distribution function of Y

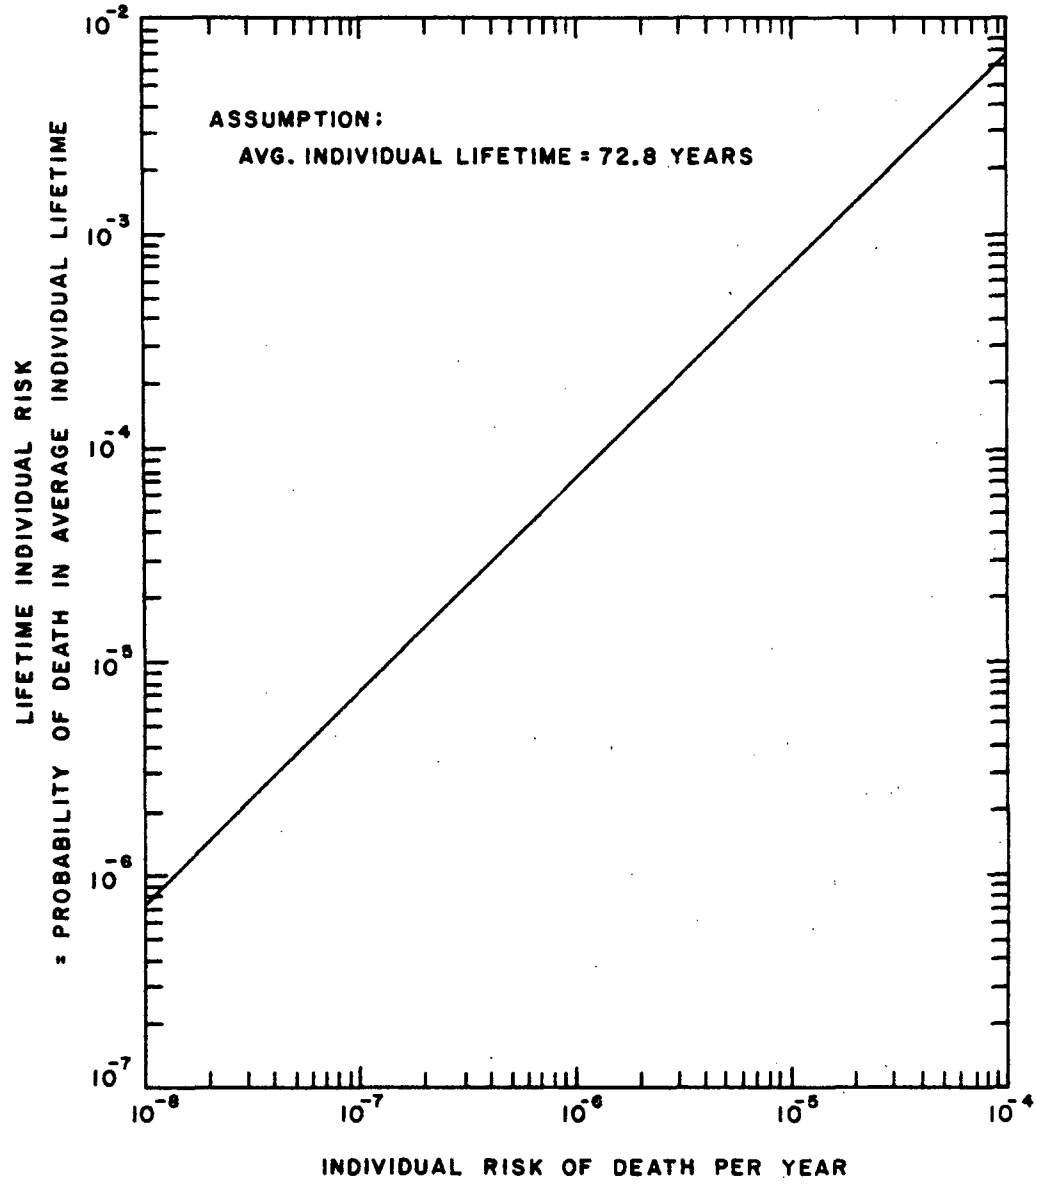

Fig. 8.1 Lifetime individual risk allowed by a criterion for the individual risk of death per year.

$g(y) = dG(y)$ = density function of Y

μ, k = first and second moment of $g(y)$

From the above expression, it is seen that the additional risk allowed by the criterion decreases the average life expectancy from μ years to $(\mu - \lambda K/2)$ years. Therefore, the reduction in life expectancy allowed by the criterion that specifies a value which is equal to λ is given as

$$\frac{\lambda K}{2}$$

The reduction in life expectancy in days as a function of various values of the additional mortality rate allowed by an individual risk criterion based on $\mu = 72.8(\text{year})^{(3)}$ and $K = 5625.41\ (\text{year})^{2(3)}$ is shown in Figure 8.2. From this figure it may be observed that if the value of the individual risk criterion is considered to be 10^{-5} per year, then it can be said to allow a loss of life expectancy equivalent to 10 days for an average individual in the U.S. It is emphasized that the allowed reduction in life expectancy is relevant only to the person to whom the criterion applies. With this important qualification in mind, the reduction in life expectancy allowed by an appropriate criterion may be compared with loss of life expectancy in days due to various accidents for an average person in the U.S. that was estimated by Cohen and Lee[4]. These estimates are shown in Table 8.2. However, it would not be correct to compare the loss of life expectancy allowed by a criterion which applies to, for instance, the maximum exposed individual with those given in Table 8.2 because the table shows the loss of life expectancy based on the risk of an average person. A comparison of the allowed loss of life expectancy implied by a criterion that is defined on the basis of an average individual in the total U.S. population at risk from all nuclear power plants with those given in Table 8.2 is more appropriate.

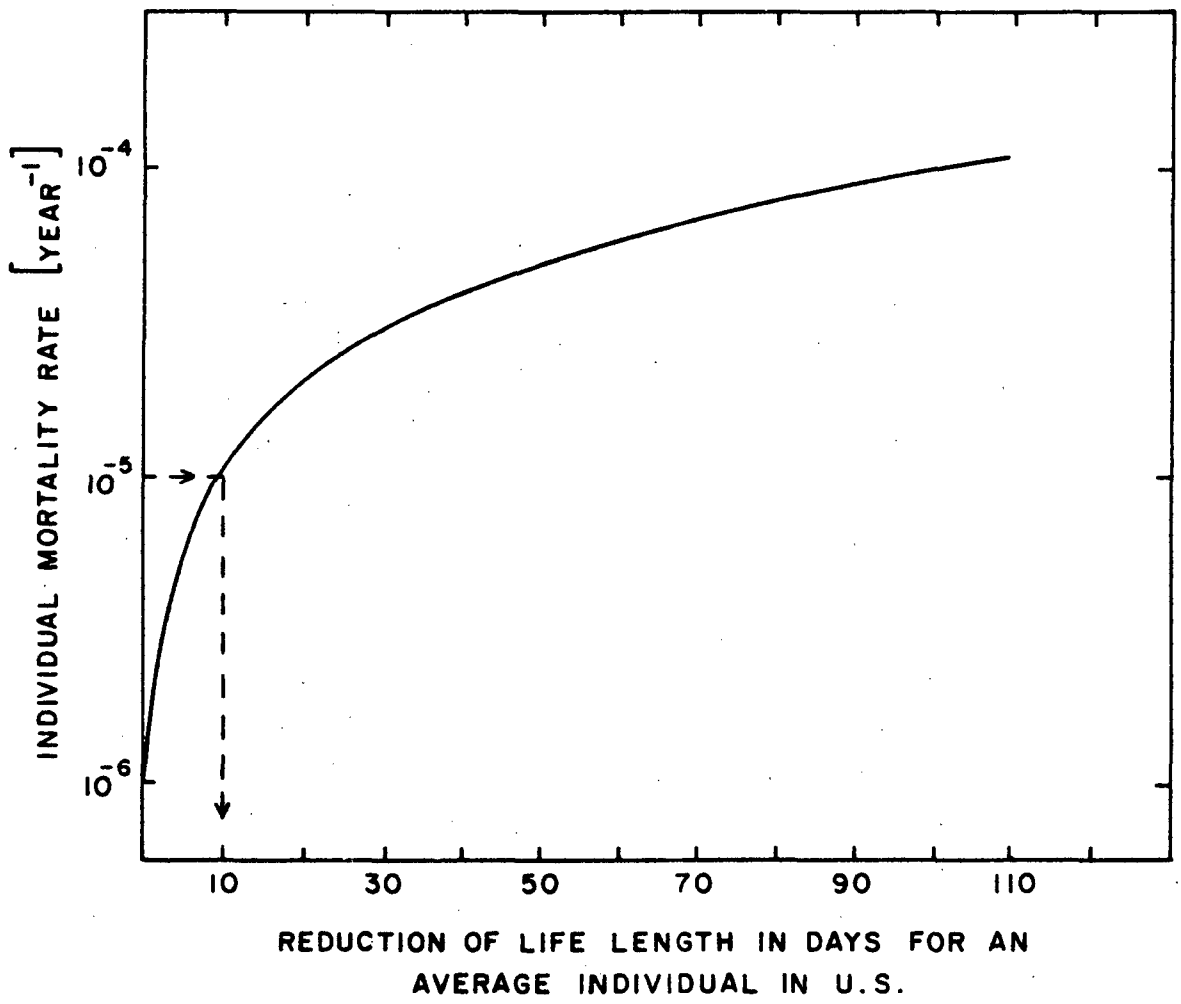

Note: The Y axis represents various values of the criterion for individual risk.

Fig. 8.2 The reduction in life expectancy allowed by a criterion for the individual risk of death per year.

TABLE 8.2. $\Delta E(M,Q)$, LOSS OF LIFE EXPECTANCY IN DAYS FOR AVERAGE AMERICAN DUE TO VARIOUS TYPES OF ACCIDENTS[4]

	AGE RANGE				
TYPE OF ACCIDENT	0-55	55-70	70-85	85- ∞	0- ∞
All accidents	505	19	31	19	435
Motor vehicle	106	7.1	6.7	3.0	207
Pedestrian	18	1.6	2.6	1.4	37
Pedalcycle	3.2	0.08	0.05	0.02	5.1
Accident in home	42	4.1	13	9.3	95
Falls	7.7	2.7	14	11	39
Drowning	23	0.9	0.6	0.2	41
Fire, -burns	12	1.6	2.3	1.2	27
Poison (sol. liq.)	7.9	0.7	0.5	0.2	17
Suffocation	6.1	0.7	0.8	0.5	13
Firearms	6.4	0.3	0.2	0.07	11
Poison (gas)	3.6	0.3	0.2	0.09	7.5

NOTE: $\Delta E(M,Q)$ = loss of life expectancy between the Mth and the Qth birth days.

Chapter 8 References

1. Atchison, R.J., "Nuclear Reactor Philosophy and Criteria," Presented on July 18, 1979 to the Select Committee on Ontario Hydro Affairs Summer Schedule of Hearing on the Safety Hearings on the Safety of Nuclear Reactors, July 1979.

2. "Statistical Abstract of the United States: 1979," U.S. Bureau of the Census, Washington, D.C. 100th Edition, 1979.

3. Thompson, W.A., "Competing Risk Presentation of Reactor Safety Studies," Nuclear Safety, Vol. 20, No. 4, July - August 1979.

4. Cohen, B.L. and Lee, I.S., "A Catalog of Risks," Health Physics, Vl, 36, pp. 707-722, June 1979.

9. SOCIETAL RISK CRITERIA

9.1 INTRODUCTION

This chapter is concerned with the societal risk criteria which address unacceptable risks of adverse health consequences to society from accidents in nuclear power plants. There can be different types of adverse health consequences, for example, early fatalities, latent fatalities, genetic effects, thyroid nodules, radiation related illnesses, etc. The societal risk criteria in conjunction with the individual risk criteria are placed at the top level of the hierarchical structure for risk criteria (see Section 3.1). A societal risk criterion is not concerned with the reliabilities of components or systems, the frequencies of various accident sequences, the integrity of the containment, the amounts of radioactive releases and their associated frequencies, the site-specific features, and public protection measures as long as the final risk to society is not above some unacceptable criterion level. Consequently, a societal risk criterion gives designers and plant operators flexibility in deciding how a plant is to be designed and operated and where it should be located while still maintaining an overall control of the societal risk. The societal risk criteria are the most flexible when compared with all other levels of criteria in our hierarchical structure for risk criteria. The reader is referred to Section 3.2 where the properties and implications of establishing societal risk criteria have been discussed. Therefore, in this chapter, we will present and review some types of societal risk criteria.

9.2 TYPES OF SOCIETAL RISK CRITERIA

As stated in Section 3.2, a societal risk criterion can be formulated to constrain the societal risk associated with a plant located at any given site for any given year (i.e. site specific societal risk criterion) or to constrain the total societal risk of all reactors operating in any given year.

Accordingly, the former is called a S-Y criterion and the latter, a Y criterion. The S-Y criterion focuses on the risk to the population surrounding a specific site while the Y criterion focuses on the risk to the total population around all sites. The S-Y criterion attempts to ensure that the societal risk associated with a plant located at any site regardless of its features (e.g., population size and distribution, evacuation measures, etc.) is not above some specified unacceptable level, i.e. all plants have the same societal risk goal in any given year. On the other hand, the Y criterion attempts to control the total societal risk of all nuclear plants operating in any given year. In order to demonstrate compliance with a S-Y criterion, the societal risk of a plant associated with a specific site is evaluated and compared with the specified criterion level. There can be two procedures which can be followed to implement a Y criterion. The first procedure consists of apportioning equally the total societal risk allowed by the Y criterion among all plants operating in any given year, i.e., all plants have the same societal risk goal in a given year. However, these goals could change from one year to the other because they depend on the number of plants considered in apportioning the Y criterion value. In contrast, the value of a S-Y criterion which is the same for all plants does not change from one year to the other. The second procedure of implementing a Y criterion consists of summing the societal risk associated with each plant and then determining whether or not the total satisfies the Y criterion value. Thus, the second procedure allows plant to plant variability in the societal risk. The S-Y and Y societal risk criteria are analogous to the R-Y and Y forms of core melt probability criteria that were discussed in Section 6.2.

The consequence component of a societal risk criterion can be expressed in terms of either some adverse health effects (e.g., early fatalities, latent fatalities, etc.) or in terms of the radiation dose integrated over the appropriate population at risk (i.e., person-rems). The dose can be expressed in terms of a particular organ (e.g., dose to the thyroid gland) and/or to the whole body.

A societal risk criterion can be expressed in various forms. It can specify: (1) a limit on the expected consequence per unit time, (2) limits on the magnitudes of consequences and their associated frequencies per unit time,

(3) limits on the frequencies per unit time which equal or exceed certain magnitudes of consequence (i.e., a limiting complementary cumulative distribution function). The unit time may be per year if it is a Y criterion or it may be per reactor-year if it is a S-Y criterion. In the rest of this section, different proposals are presented which relate to the above discussion.

The Canadian Reactor Siting Guide[1] specifies maximum population dose limits and their associated frequencies. This is shown in Table 9.1. The

TABLE 9.1

OPERATING DOSE LIMITS AND REFERENCE DOSE LIMITS FOR ACCIDENT CONDITIONS SPECIFIED BY THE CANADIAN REACTOR SITING GUIDE[1]

Situation	Assumed Maximum Frequency	Meteorology to be Used in Calculation	Maximum Individual Dose Limits	Maximum Total Population Dose Limits
Normal Operation		Weighted according to effect, i.e frequency times does for unit release	0.5 rem/yr whole body 3 rem/yr to thyroid (a)	10^4 man-rem/yr 10^4 thyroid rem/yr
Serious Process Equipment Failure	1 per 3 years	Either worst weather existing at most 10% of time or Pasquill F condition if local data incomplete	Same as annual doses for normal operation	
Process Equipment Failure Plus Failure of any Safety System	1 per 3×10^3 years	Either worst weather existing at most 10% of time or Pasquill F condition if local data incomplete	25 rem whole body 250 rem thyroid (b)	10^6 man-rem 10^6 thyroid rem

(a) For other organs use 1/10 ICRP occupational values
(b) For other organs use 5 times ICRP annual occupational dose

population dose is expressed in terms of the dose to the whole body and to the thyroid gland. The frequencies are specified on the basis of per reactor-year of operation. The Siting Guide can be interpreted as a site specific societal risk criterion.

Kinchin[2] of the United Kingdom Atomic Energy Authority (UKAEA) proposed societal risk criteria for early and latent fatalities per year applicable to a single nuclear reactor. These criteria as shown in Figure 9.1 are expressed in the form of a complementary cumulative distribution function. This figure also shows the risk of an average reactor as assessed in WASH-1400. The criterion for early fatalities is set at a factor of 30 below the criterion for latent fatalities per year to reflect the increased concern for early fatalities. Kinchin's criterion can also be interpreted as being site specific.

In contrast to the above criteria, Levine proposed an interim safety goal[3] for the total societal risk which is a combination of the risks of early and latent fatalities from all nuclear reactors. His safety goal, which is shown in Figure 9.2 is expressed in the form of a complementary cumulative distribution function (CCDF). This goal is set at one tenth the level of the lowest non-nuclear risk found in WASH-1400, i.e., risk of death for persons on the ground from aircraft crashes. This figure also shows the weighted sum of the societal risk of early and latent fatalities expressed as a CCDF for 100 nuclear power plants based on the results of WASH-1400 for an average reactor. Levine obtained the weighted sum by using a factor of 30 to indicate the relative importance of early fatalities with respect to latent fatalities.

Since Levine's goal is intended to be achieved by considering all nuclear reactors, it can be classified as a Y criterion. Both Kinchin's and Levine's proposals take into account society's increased concern with events that lead to large consequences.

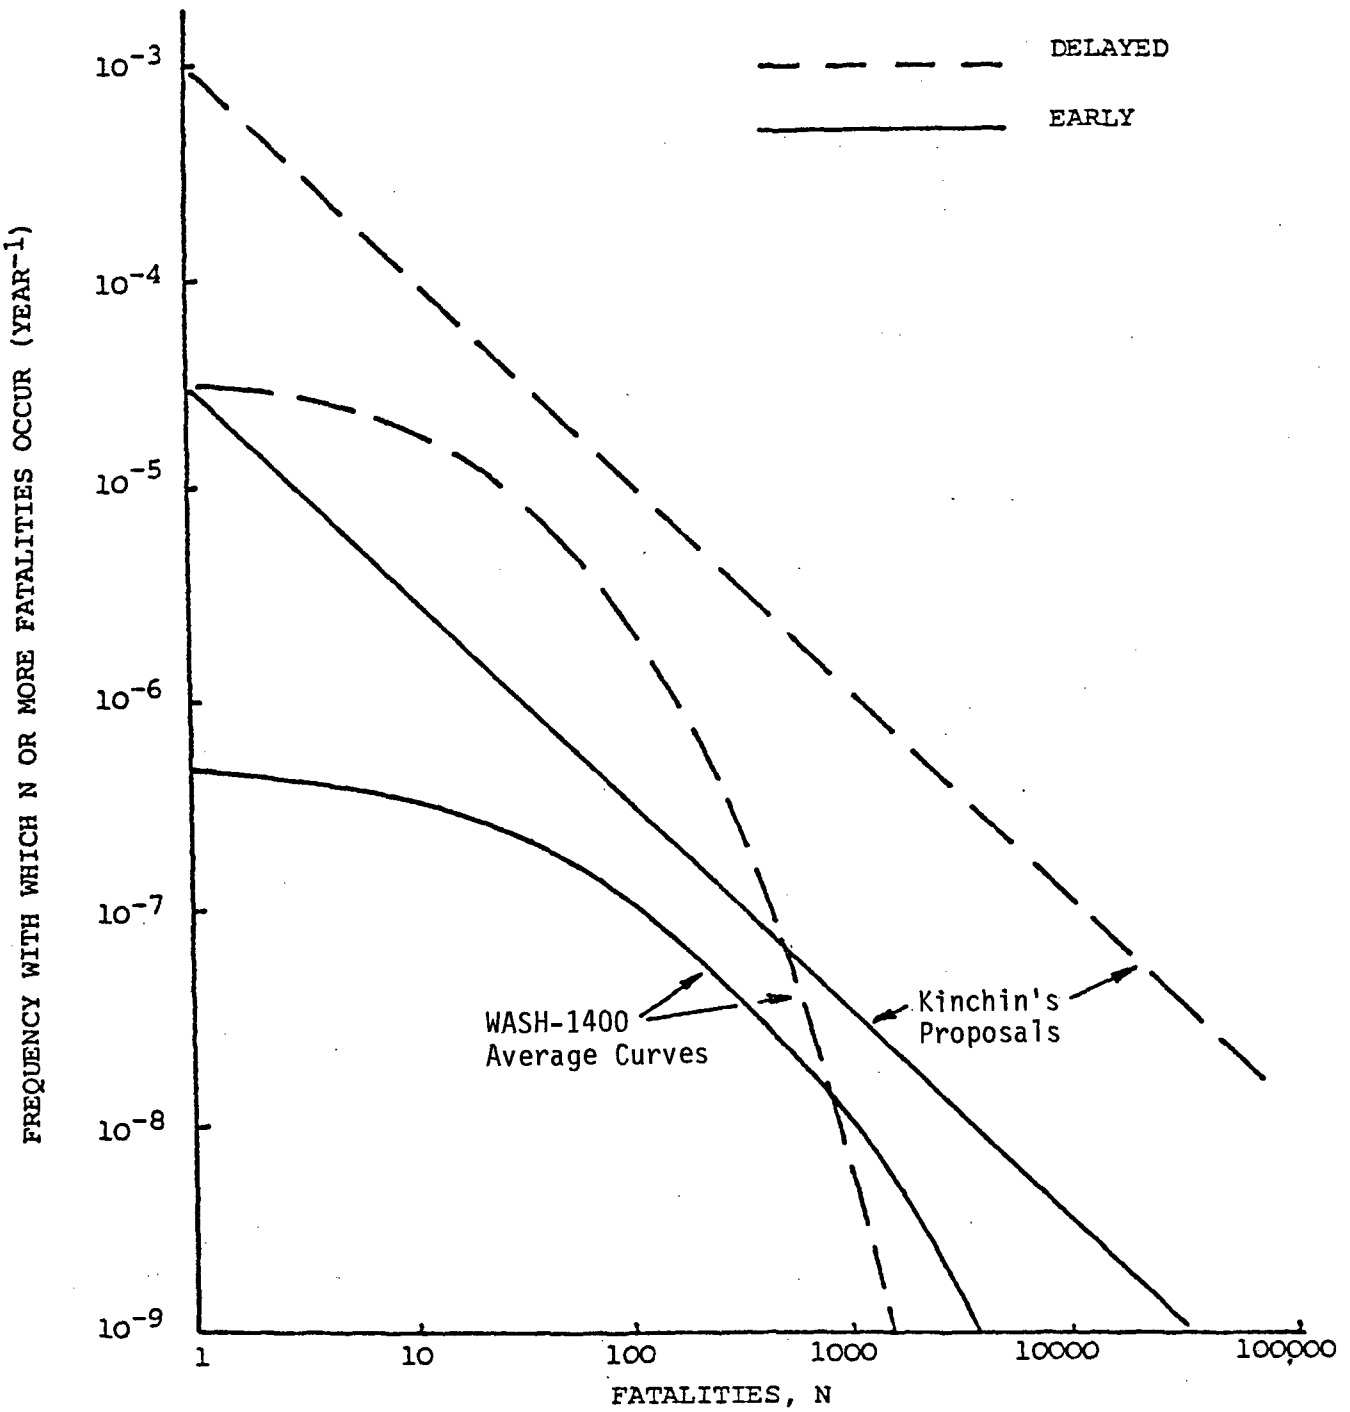

Fig. 9.1 Societal risk criteria for a nuclear reactor proposed by G.H. Kinchin of the UKAEA.

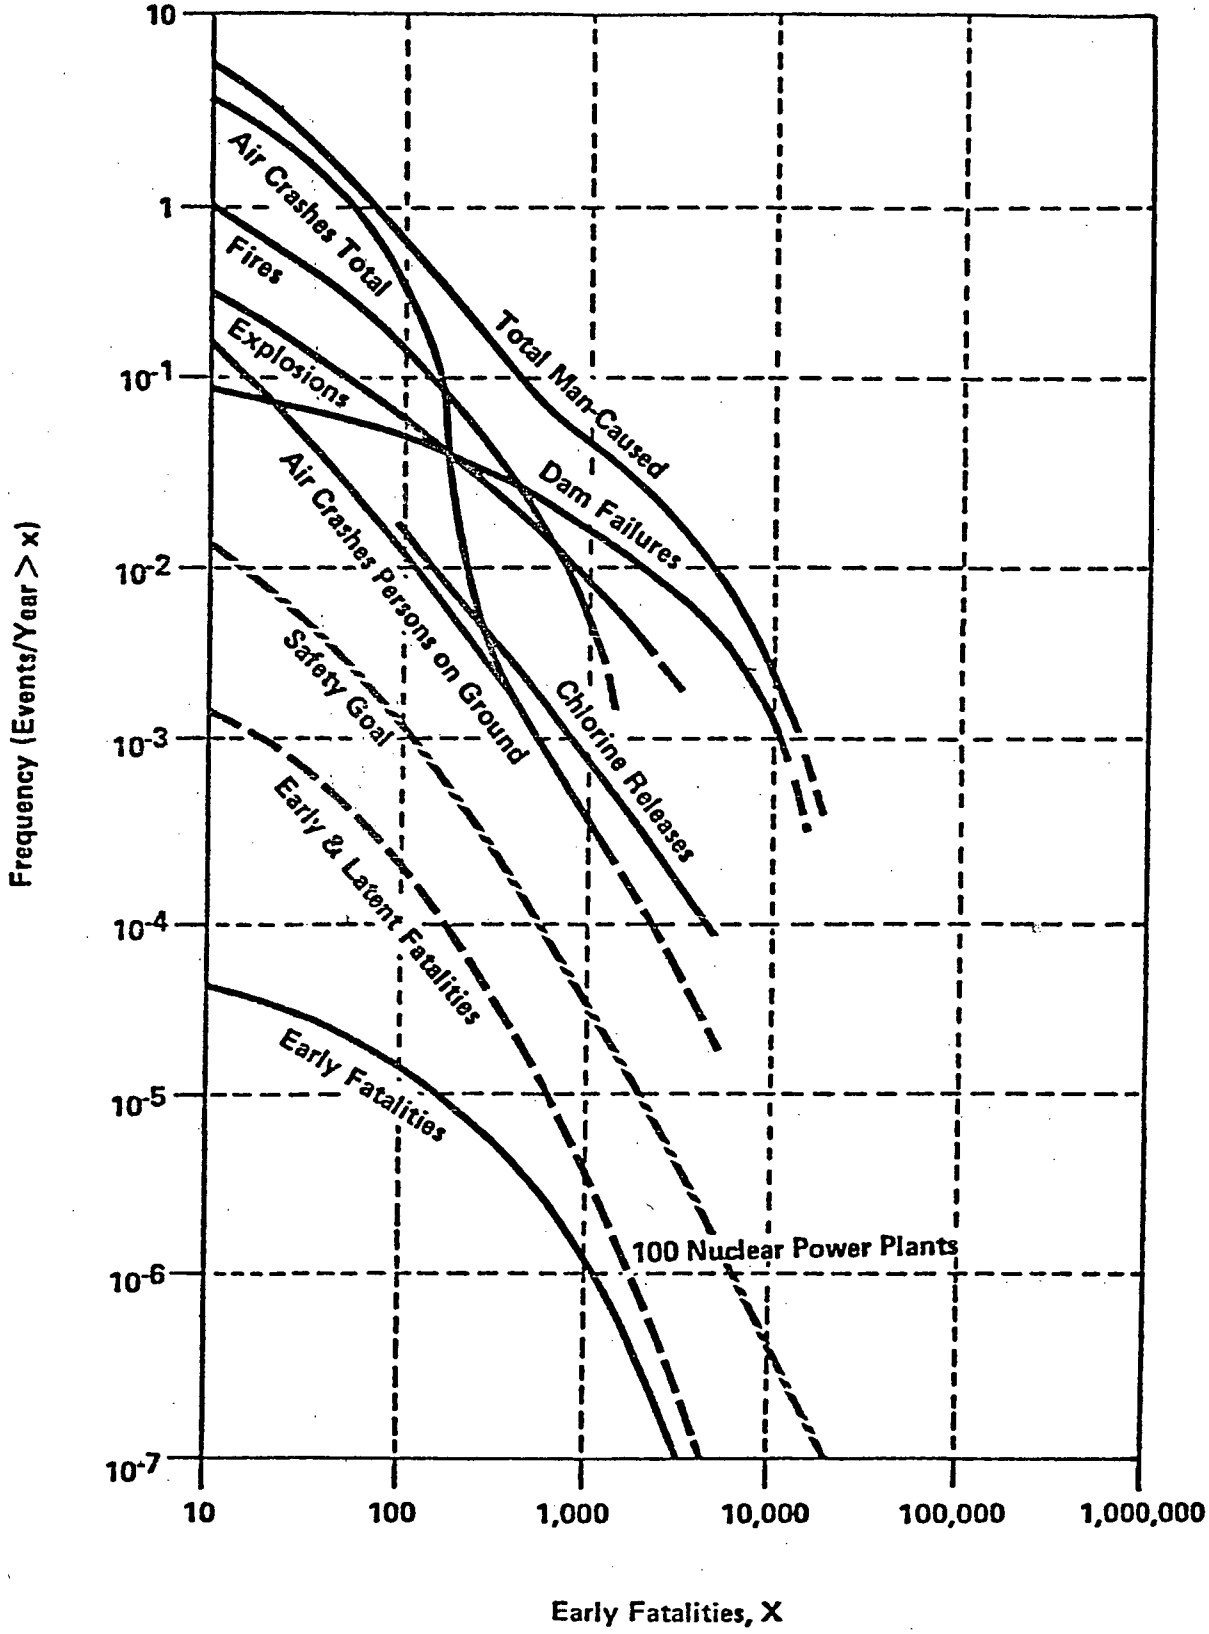

Fig. 9.2 Safety goal proposed by Levine.

Chapter 9 References

1. Atchison, R.J., "Nuclear Reactor Philosophy and Criteria," Presented on July 18, 1979 to the Select Committee on Ontario Hydro Affairs Summer Schedule of Hearings on the Safety Hearings on the Safety of Nuclear Reactors, July 1979.

2. Kinchin, G.H., "Design Criteria, Concepts and Features Important to Safety and Licensing," Presented at the ANS/ENS International Meeting on Fast Reactor Safety Technology, Seattle, Washington, August 1979.

3. Levine, S., "TMI and the Future of Reactor Safety," Presented at International Public Affairs Workshop - Atomic Industrial Forum, Stockholm, Sweden, June 1980.

10 PROPERTY DAMAGE RISK CRITERIA

10.1 INTRODUCTION

In the preceding two chapters, criteria were discussed that address health risks to society or to an individual due to accidents in nuclear power plants. In this chapter, the property damage risk criterion will be reviewed. This criterion addresses the economic risk to society from accidents in nuclear power plants. As stated in Chapter 3, the property damage risk criterion in conjunction with criteria that address health risks to society or to an individual form the top level of the hierarchical structure for risk criteria. Property damage risk criterion can be specified to focus only on the risk of property damage outside of the plant boundary or it may include in-plant property damage.

The adverse health impact on the public from a potential serious reactor accident results from exposure to airborne radioactive material and material deposited on the ground, and from ingestion of contaminated food. The principal action taken to mitigate the adverse health consequences to the public in the path of a radioactive cloud is evacuation. This mitigatory measure minimizes the early exposure. Mitigatory measures to minimize long term exposure to radioactive material deposited in the environment may consist of decontamination of land and structures, interdiction of land, i.e., denial or restriction of its use, impoundment of contaminated crops and milk, etc. The cost of the health mitigating measures have been assessed in WASH-1400[1] and results, as given in Fig. 10.1, are presented in terms of a Complementary Cumulative Distribution Function (CCDF) for property damage in terms of 1974 dollars. The cost due to property damage as assessed in WASH-1400[1] was based on property damage outside the plant and included the following:

1. Evacuation costs
2. Temporary relocation costs
3. Land and structure decontamination cost

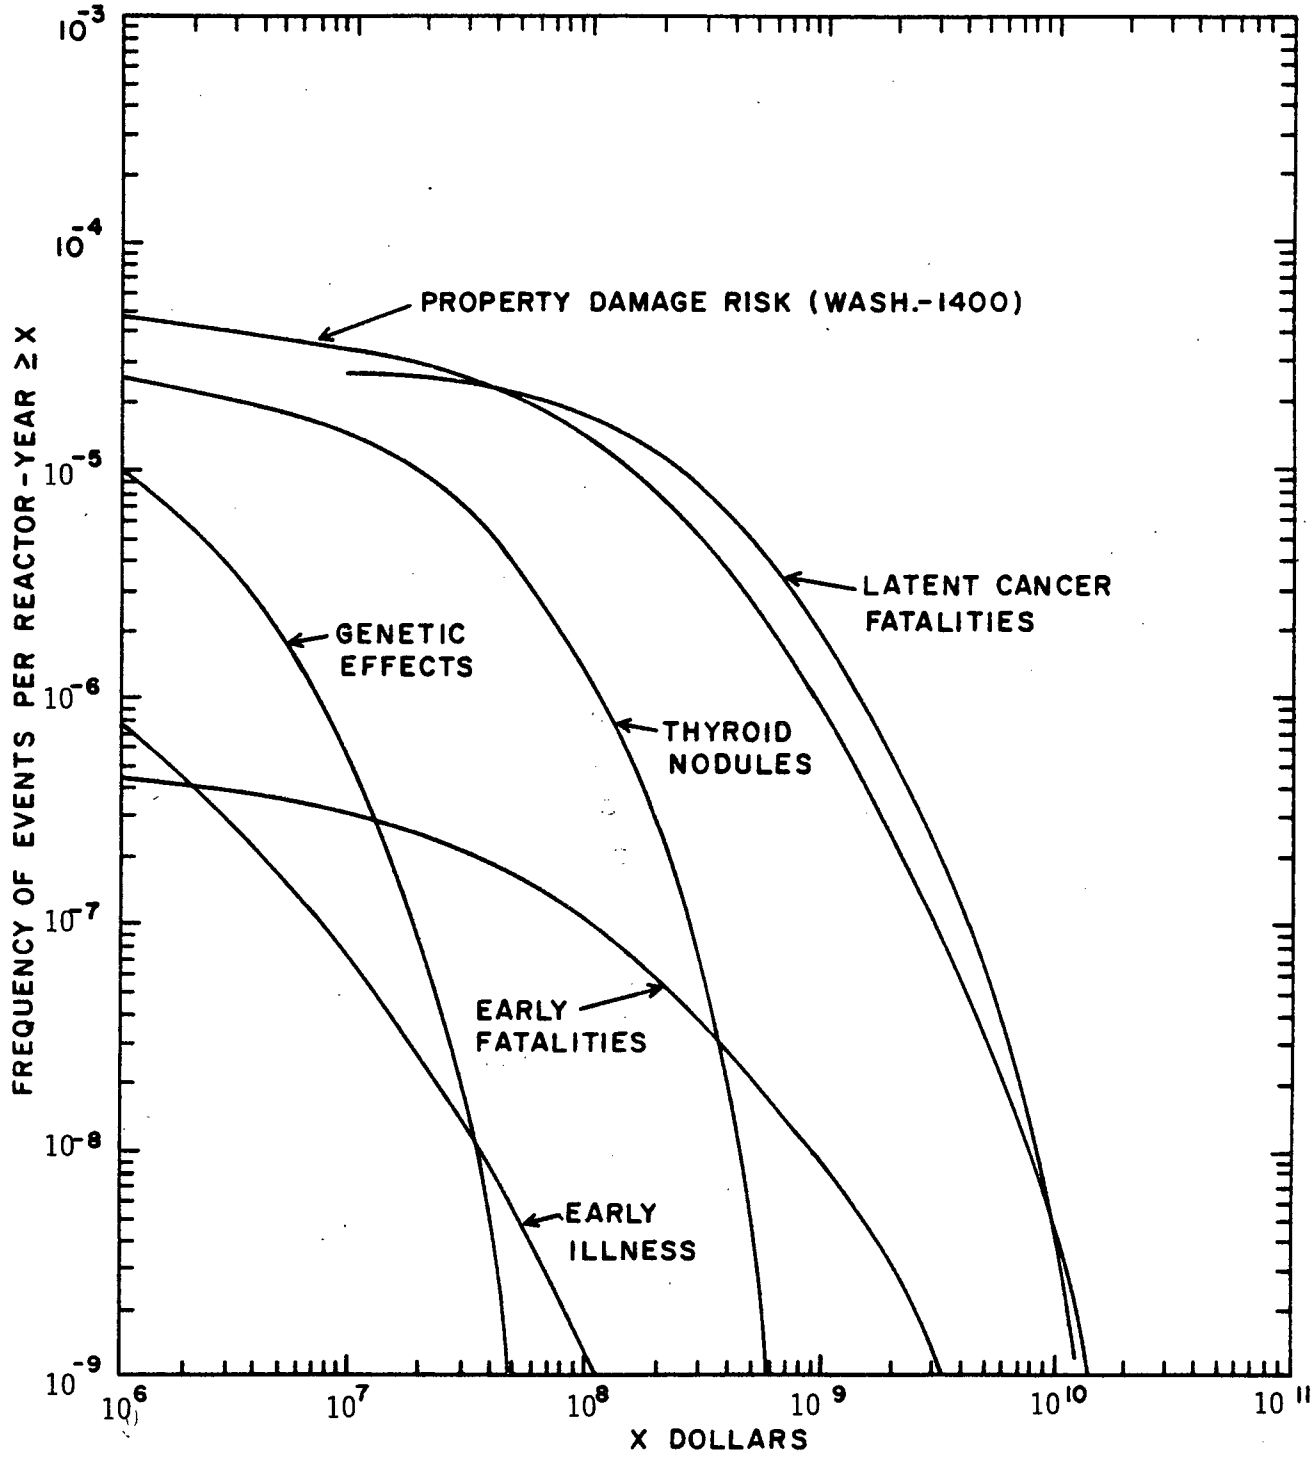

Fig. 10.1 Complementary cumulative distribution functions for total property damage and other consequences in terms of their dollar values.

4. Cost due to condemned milk and crops

5. Costs associated with loss of productive use of land and its improvements.

The type of property damage due to a major nuclear accident is different in comparison to other causes such as floods, fires, hurricanes, etc. Nuclear accidents do not result in physical damage to the property but have the potential for contamination to levels that may be considered unacceptable for long term occupation.

10.2 MOTIVATIONS FOR ESTABLISHING A CRITERION ON PROPERTY DAMAGE RISKS

The question of whether or not a risk criterion on property damage is desirable may be addressed by comparing the importance of societal risk of property damage with other types of consequences such as early fatalities, early illness, latent fatalities, thyroid nodules, and genetic effects. Results of such a comparison obtained from Ref. 2 are presented in Table 10.1. In this

TABLE 10.1

COMPARISON OF THE IMPORTANCE OF WEIGHTED ATTRIBUTES
(from Reference 2)

Attribute	Societal[a] Risk	Adjusted Societal Risk***	Weight	Weighted Risk
Early Fatalities	3×10^{-3}	3×10^{-5}	1×10^{6}	30
Early Illness	2×10^{-1}	2×10^{-3}	2.5×10^{3}	5
Delayed Fatalities	7×10^{-2}/yr	2×10^{-2}	3×10^{5}	6300
Delayed Illness	7×10^{-1}/yr	2×10^{-1}	2.5×10^{3}	500
Genetic Effects	1×10^{-2}/yr	3×10^{-3}	1×10^{4}	30
Property Damage	2×10^{6} **	7×10^{4}	1	7×10^{4} *

(a) Societal Risk for 100 Reactors as per WASH-1400.
* Sum of in-plant and out-of-plant property damage risk.
** Out-of-plant property damage risk for 100 Reactors.
*** Societal Risk per plant year.

table, the weighted risk of different types of consequences or attributes were obtained by weighing the risk of each type of consequence that was assessed in WASH-1400[1] with appropriate "weights." In essence, these "weights" are monetary values assigned to each type of consequence to reflect their importance. The risk of property damage shown in column 3 of Table 10.1 was obtained by summing the plant property damage risk (equal to 5×10^4 dollars per plant year) and the public property damage risk (equal to 2×10^4 dollars per plant year). The plant property damage risk was calculated by multiplying the frequency of core melt accidents (5×10^{-5} per reactor year) with the cost due to cleanup and repair for each core melt accident ($\$10^9$). If one accepts the different "weights" used in the above mentioned analysis,[2] it may be inferred that property damage is the attribute of major concern because the sum of the weighted risk of adverse health effects (equal to 6865) is only 10 percent of the weighted risk of property damage. In parenthesis, it should be noted that the weighted risk of different types of consequences varies from one plant to another. Consequently, the weighted risk, as it applies to a particular plant, can be different from that portrayed in Table 10.1. In Appendix B, analyses are presented which show the variability in weighted risks of different types of consequences considering design and site differences. Based on these analyses and assumptions, it appears that the plant property damage risk is higher than the out-of-plant (public) property damage risk and the weighted risk of other types of adverse health consequences. However, for some plants the risk of public property damage may be smaller than the combined weighted risks of early and latent fatalities, a situation different from that portrayed in Table 10.1.

Another way of discerning the relative risk importance of health consequencs in relation to property damage risk is to convert the complimentary cumulative distribution functions of health consequences, as portrayed in WASH-1400, to dollars. This comparison presented in Fig. 10.1 was obtained by using the same weights as shown in Table 10.1.

A weakness associated with justifying any inferences that may be drawn from the above mentioned analyses is that it is sensitive to the assumed weights. If the weights of all the adverse health effects were increased by a factor of about 10, the sum of the weighted health risk would be equal to the

property damage risk of 7×10^4. Only an increase in the weights of the different health attributes by a factor of less than 3 is necessary to equal the property damage risk of $\$2 \times 10^4$ (this figure excludes power plant property damage cost). In light of the large disparities in the dollar value of latent and early fatalities that are operative in society today, as discussed in the next section, assignment of dollar values to other nonfatal adverse health attributes becomes a complex issue. Table 10.1 from Ref. 2 construes the following relative importance of each adverse health attribute.

Attribute	Relative Importance
Early Illness	1
Delayed Illness	1
Genetic Effects	4
Delayed Fatalities	120
Early Fatalities	400

There is another important reason, beside the high weighted risk, that could call for a consideration of a constraint on public property damage risk. This reason would arise if the probability of public property damage exceeding the insurance coverage is considered to be high. However, the maximum public property damage insurance cannot be explicitly determined since there exists no specific insurance that covers only public property damage. This is because public property damage and other kinds of public liabilities (e.g. liabilities arising from adverse health effects) that could result from nuclear accidents are covered, as a whole, by the insurance provided under the Price Anderson Act.[3] The Price Anderson Act limits the aggregate public liability of the reactor operator and others who might be at fault to $560 million. This type of insurance coverage is different from the type that is available for plant property damage. In the case of plant property, there exists a specific insurance coverage for damages to plant property. At present, the maximum plant property insurance coverage available from commercial companies is $300 million.[4] Based on the results of WASH-1400,[1] as shown in Fig. 10.1, it may be observed that there is a finite probability for accidents to occur that could result in public liabilities in excess of the maximum insurance coverage. In the event that losses suffered by the public exceeds the coverage, the liability rests upon the claimants. Consequently, a viewpoint for developing a criterion to constrain public property damage risks

is to require that the frequency of all accidents with a potential to exceed the insurance coverage be made low. This concept is discussed further in Section 10.4.

To put the matter of public liability exceeding insurance coverage in perspective, it needs to be recognized that there are nonnuclear technological hazards (e.g., dam rupture, liquified natural gas (LNG) accidents, fires, etc.) and natural disasters (e.g., floods, hurricanes, earthquakes, etc.) wherein it is conceivable that the cost of property damage exceeds the insurance coverage. Solomon and Okrent[5] studied an overview of situations in large technological systems in society where the resultant liabilities exceeded the insurance coverage, and concluded that, "the presence of severe de-facto limits on liability appears to be part of the fabric of our society."

10.3 PROPERTY DAMAGE RISK - A COST BENEFIT VIEWPOINT

In this section, the manner in which a cost-benefit approach can aid in the formulation of an implicit criterion for property damage is examined and then its principal limitations are discussed.

The cost of property damage from major nuclear accidents is the cost which is incurred to minimize the adverse health effects. The benefit is the reduction in morbidity and/or mortality in the exposed population. It should be realized, that costs associated with property damage are dependent, in part, on the definition of the acceptable standard for radiation exposure. This in turn determines the physical boundaries and the extent of decontamination and interdiction. From the cost-benefit point of view, any criterion on the unacceptability of property damage risks should be set at a level such that the benefits equal the costs. In other words, the total cost, which is a sum of the cost of mitigation of adverse health effects (i.e., property damage cost) and the cost that can be associated with residual adverse health effects after mitigatory actions have been taken, should be minimized. This approach leads to an optimal de-facto standard of radiation exposure from economic considerations alone, which could be used to define the boundaries and the extent

of interdiction and decontamination. This concept is illustrated in Fig. 10.2. If the standard is overly stringent, the costs of mitigatory measures could be very high but the cost of the residual adverse health effects would be low. Alternately, a lenient radiation standard belies a high cost associated with adverse health effects and a low cost of mitigatory measures. The point where the total cost is minimized corresponds to what may be called an "economic optimal radiation standard." A criterion determining unacceptability of property damage risk could be hypothetically based on this economic optimal radiation standard. The principal limitations that are associated with following such a cost-benefit approach are discussed next.

Overestimating the cost of mitigatory measures undertaken to reduce the adverse health effects could result in the determination of an economic optimal radiation standard that may be considered lenient. Conversely, overestimating the cost of adverse health effects could result in an optimal radiation standard that may be considered stringent. The other major problem in the cost-benefit approach is that monetary values need to be explicitly assigned to different types of health effects. From the subsequent discussion it appears that there is no national consensus on such monetary values.

The notion of assigning monetary values to adverse health effects may appear to some as immoral or unethical. However, governmental agencies, commercial organizations, and even individuals employ some variation of cost-benefit approaches towards decision making by ascribing some sort of monetary values to human lives. Generally, individual decision making procedures are implicit, non quantitative and based on heuristic considerations. Thus, one may decide to buy low priced tires rather than the type that cannot blow out, or decide against the required frequency of medical check-ups.[6]

A retrospective look at the societal expenditures of averting a statistical death or deferring a premature death reveals a large dispersion and implies a poor consensus of opinion as to what monetary values should be assigned to adverse health effects. This is substantiated by an analysis by Cohen[6] where dollar values were derived from societal actions taken to avert fatalities. A summary of the result reproduced from Ref. 6 is shown in Table 10.2. Okrent[7] has highlighted the inconsistency in the implicit

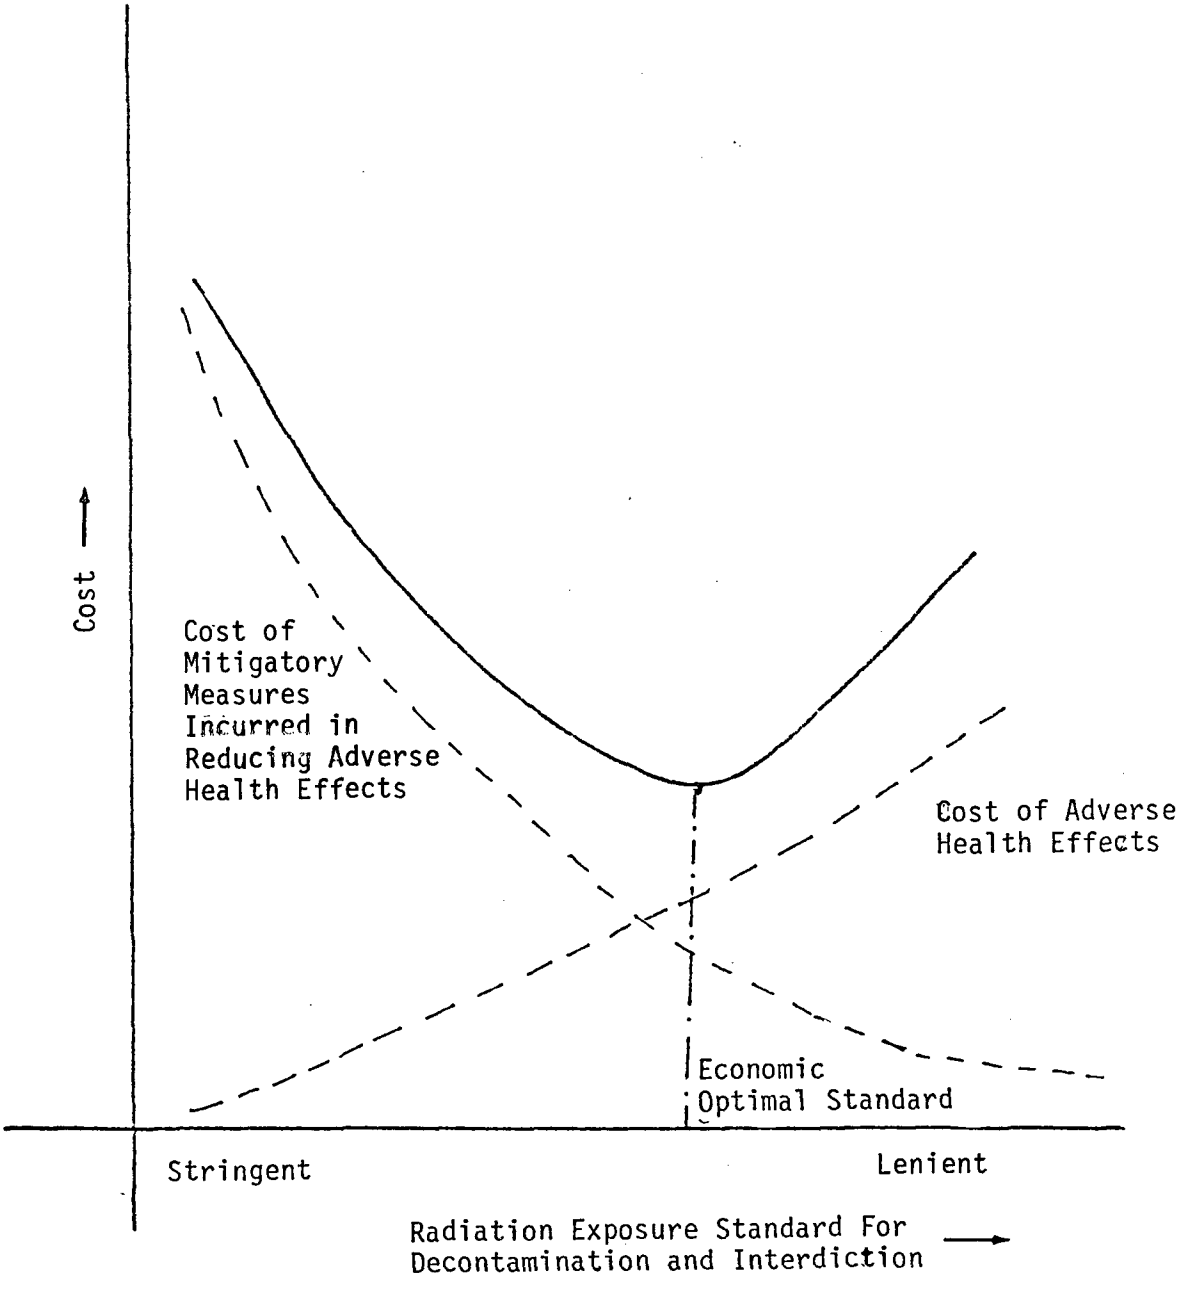

Fig. 10.2. Economic optimal radiation exposure standard for decontamination and interdiction.

Table 10.2. Value per fatality averted (1975 dollars) implied by various Societal activities (left column) and cost per 20 years of added life expectancy (right column).*

Item	$ per fatality averted	$/20 yr life expectancy
Medical screening and care		
cervical cancer	$ 25,000	$ 13,000
breast cancer	80,000	60,000
lung cancer	70,000	70,000
colorectal cancer:		
fecal blood tests	10,000	10,000
proctoscopy	30,000	30,000
multiple screening	26,000	20,000
hypertension control	75,000	75,000
kidney dialysis	200,000	440,000
mobile intensive care units	30,000	75,000
Traffic safety		
auto safety equipment–1966-70	130,000	65,000
steering column improvement	100,000	50,000
air bags (driver only)	320,000	160,000
tire inspection	400,000	200,000
rescue helicopters	65,000	33,000
passive 3-point harness	250,000	125,000
passive torso belt-knee bar	110,000	55,000
driver education	90,000	45,000
highway construc.-maint. practice	20,000	10,000
regulatory and warning signs	34,000	17,000
guardrail improvements	34,000	17,000
skid resistance	42,000	21,000
bridge rails and parapets	46,000	23,000
wrong way entry avoidance	50,000	25,000
impact absorbing roadside dev.	108,000	54,000
breakway sign, lighting posts	116,000	58,000
median barrier improvement	228,000	114,000
clear roadside recovery data	284,000	142,000
Miscellaneous non-radiation		
food for overseas relif	5,300	2,500
sulfur scrubbers in power plants	500,000	1,500,000
smoke alarms in homes	240,000	140,000
higher pay for risky jobs	260,000	150,000
coal mine safety	22,000,000	13,000,000
other mine safety	34,000,000	20,000,000
coke fume standards	4,500,000	2,500,000
Air Force pilot safety	2,000,000	1,000,000
civilian aircraft (France)	1,200,000	600,000
Radiation related activities		
radium in drinking water	2,500,000	2,500,000
medical X-ray equipment	3,600	3,600
ICRP recommendations	320,000	320,000
OMB guidelines	7,000,000	7,000,000
radwaste practice-general	10,000,000	10,000,000
radwaste practice – ^{131}I	100,000,000	100,000,000
defense high level waste	200,000,000	200,000,000
civilian high level waste		
no discounting	18,000,000	18,000,000
discounting (1%/year)	~1,000,000,000	~1,000,000,000

*Reproduced from Ref. 6.

value of life with regard to two energy options for electrical power generation. A figure of $30,000 per premature death was used for coal fired plants in a report from the National Academies of Sciences and of Engineering entitled "Air Quality and Stationary Source Emission Control." The rationale for the choice of $30,000 against $200,000 used in highway safety was justified on the grounds that deaths from coal pollutants would occur among chronically ill elderly people leading to life reduction in the range of days or weeks. On the other hand, in the ALARA (As Low As Reasonably Achievable) concept utilized by the NRC for judging control measures for routine radioactive release, a figure of $1,000 per person-rem is used. This translates to a value of 5 million dollars per premature death deferred, based on the BEIR report's conversion factor of 5,000 person-rem per statistical death.[7]

10.4 FORMS OF PROPERTY DAMAGE RISK CRITERIA AND THEIR IMPLICATIONS

In the previous section, the manner in which cost-benefit analysis could be hypothetically utilized to aid in the formulation of property damage risk criteria was discussed. Criteria based on this approach are implicit. In this section, some explicit forms of property damage risk criteria with regard to their implications are examined.

Property damage risk criteria can be expressed in the following forms:

1. As a number which specifies the expected property damage in dollars per plant year of operation that is unacceptable.

2. As a limiting curve which specifies the frequencies and the magnitudes of property damage in dollars that are unacceptable.

3. As a limiting complimentary cumulative distribution function which specifies the frequencies of equalling or exceeding specific magnitudes of property damage that are unacceptable.

The above three forms are similar to the forms of radioactive release criteria that have been considered in Chapter 7. Hence, the implications of expressing property damage risk criteria in these forms are very similar to those that were discussed with regard to the release criteria. Therefore, we will only highlight the essential implications for the sake of completeness.

The first form of criteria, as stated above, focuses only on the risk of property damage. This can be defined as the expected value of property damage in dollars per plant year of operation. In order to show compliance with this form of criterion one would need to evaluate the property damage risk of a plant and compare it with the number specified by the criterion. Based on this comparison it would be judged whether or not the evaluated risk is unacceptable. This form of criteria would not distinguish between plants which have the same property damage risk but have different mixes of high probability low consequece events and low probability-high consequence events. In addition, this form of criteria as it applies to a particular plant does not ensure that the frequency of an accident resulting in extensive property damage is lower than the frequency of an accident resulting in minor property damage. This feature is incorporated in the second form of criteria.

The second form of criteria, as stated above, focuses on the frequency of an accident as well as its resulting property damage consequence. This form of criteria is identical to Farmer's limit line which judges the risk of accidents on the basis of their frequencies and their resulting consequences in terms of iodine-131 releases to the environment. We have discussed Farmer's Limit Line in Chapter 7. By analogy, this form of criteria for property damage is called a limit line on property damage. An example of this form of criteria is shown in Fig. 10.3. In this figure, the limiting curve (limit line) which can be used to determine whether or not the frequencies and magnitudes of property damage are unacceptable is represented as a straight line.

In order to show compliance with a limit line one would first need to identify accident sequencies with a potential for property damage, and then assess their frequencies and property damage consequences. Thus, a particular accident sequence may be represented by a point in the frequency - dollar damage plane. If this point falls above the limit line, the accident sequence would be considered unacceptable thereby necessitating modifications to reduce either the frequency of the accident, or the magnitude of the property damage, or both, such that the modified accident sequence point falls below the limit line. In this manner, the limit line might serve as a useful tool to judge the unacceptability of individual accidents with regard to their property damage risk.

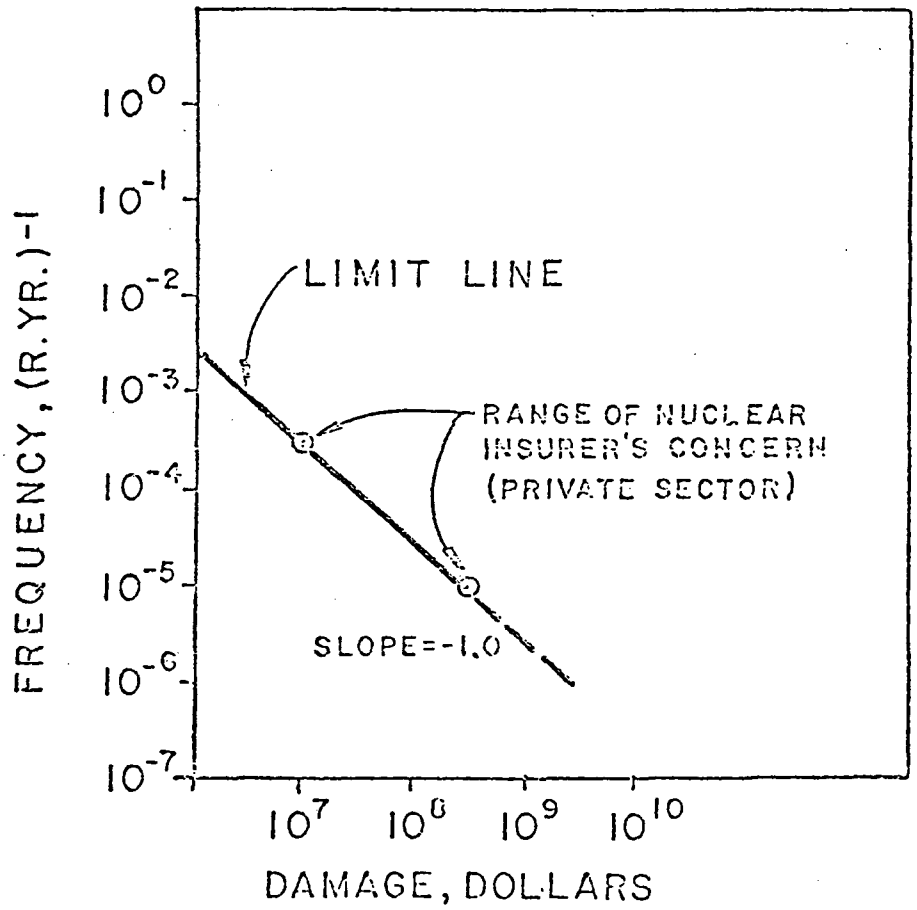

Fig. 10.3. Quantitative safety goal for public property risk proposed by Joksimovic.

*Reproduced from Ref. 8.

In the preceding paragraph, it was assumed that associated with a particular accident sequence is its frequency and property damage consequence which could be represented by a point in the frequency - dollar damage plane. If in reality such a situation prevailed, it would be easy to decide whether or not any point was above the limit line. In practice, there is a spread in the estimates of both the frequency and the property damage consequence of an accident. Therefore, an accident cannot be represented realistically by a point in the frequency - dollar damage plane. Consequently, it is not apparent how a decision can be made as to whether an accident is above or below a limit line unless the limit line is defined appropriately.

The limit line, if it is intended to serve as a decision tool to judge the unacceptability of individual accident sequences with regard to their property damage risk does not appear to constrain the total risk of property damage from all accidents. This is due to the fact that each point representing an accident can satisfy the criterion by falling below the limit line but there may be a large number of points (accidents) which make the total risk unacceptable. This limitation can be eliminated by specifying that frequencies of accidents with similar property damage consequences be summed to show compliance with a limit line.

The limit line, shown in Fig. 10.3, was proposed by Jocksimovic[8] as a preliminary safety goal for property damage outside the plant. He based the safety goal on the following considerations:

- Risks of concern to nuclear insurers range from 10 to 300 million dollars
- Insurance coverage in the absence of the Price Anderson Act
- A slope of -1.0 reflecting no risk aversion
- A low probability of 10^{-5}/per reactor year for the largest loss of 300 million dollars.

Incidentally, the upper limit of insurance coverage available from nuclear insurers (equal to $300 million), as stated above, should not be construed as the maximum insurance coverage available from insurance companies for property damage outside the plant. As stated in Section 10.2, 300 million dollars is the maximum insurance coverage available commercially for damages to property inside the plant site boundary.

As an alternative to the limit line, a criterion on property damage risk can be expressed in the form of a limiting complimentary cumulative distribution function (CCDF) of property damage. This criterion would specify the frequencies of equalling or exceeding specific levels of property damage that are unacceptable. The CCDF form of criteria has the ability to constrain the total risk of property damage from all identified accidents. In addition, it is suitable for handling the assessed spectrum of property damage consequences that may result from a particular accident. A weakness of this form of criteria, in comparison to the limit line, is its lack of discernability on an accident by accident basis as to whether or not the contribution to the total property damage risk from a particular accident sequence is inordinately high.

Criteria in the form of a limiting CCDF offers a convenient way of expressing the notion that the frequency of accidents resulting in public property damage that exceed the insurance coverage should be small. However, formulation of a criterion based on this idea is not feasible because, as stated in Section 10.2, there exists no exclusive insurance coverage for public property damage. Liability arising from public property damage in conjunction with other kinds of public liabilities are covered, as a whole, by insurance under the Price Anderson Act.[3] In this context, approaches that entail constraining the risk of public liability as opposed to those that constrain only the risk of public property damage appear to be more feasible. Consequently, a criterion in the form of a limiting CCDF can be formulated to ensure that the frequency of accidents with resultant public liabilities in excess of the insurance coverage is small.

Chapter 10 References

1. "Reactor Safety Study, An Assessmnt of Accident Risks in U.S. Commercial Nuclear Power Plants," Appendix VI, USNRC, WASH-1400, NUREG-75/014, October 1975.*

2. Clausen, M.J., Fraley, D.W. and Denning, R.S., "Improved Methods for Incorporating Risk in Decision Making," (Interim Report), PNL 3523, August 1980.

3. "Legal Consequences of Nuclear Accidents and Shutdowns," Transcript of Proceedings held in Hershey, Pennsylvania, July 27-28, 1979, Pennsylvania Law Journal, Philadelphia, 1979.

4. Kehoe, K., "The Story Behind Nuclear Insurance," Critical Mass Journal, June 1980.

5. Solomon, K.A. and Okrent, D., "Catastrophic Events Leading to De-Facto Limits on Liability," UCLA-ENG-7732, May 1977.

6. Cohen, B.L., "Society's Evaluation of Life Saving in Radiation Protection and Other Contexts," Health Physics, Vol, 38, pp. 33-51, January 1980.

7. Okrent, D., "Testimony Presented on July 25, 1979 to the Forum on Risk/Benefit Analysis in Legislative Process, Subcommittee on Science, Research and Technology Committee on Science and Technology, U.S. House of Representatives," July 1979.

8. Joksimovic, V., "Statement of Quantitative Safety Goals before the ACRS Subcommittee on Reliability and Probabilistic Assessment," General Atomic Company, July 1980.

*Available free upon written request to the Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

11. CONSIDERATIONS IN THE USE OF RISK CRITERIA

11.1 INTRODUCTION

The use of quantitative risk assessment to evaluate safety in the licensing of nuclear power plants in the United States has been suggested for many years. The publication of WASH-1400[1] in 1975 spurred interest in this concept and caused the Nuclear Reliability Subcommittee of the IEEE (SC.5) to suggest that a quantitative risk criterion should be considered in select cases as an alternative to the "Single Failure Criterion." The American Physical Society's study[2] on light water reactor safety and the Risk Assessment Review Group of the Nuclear Regulatory Commission[3] recommended in 1978 that the risk assessment approaches pioneered by WASH-1400 should come into more general use in the licensing process. In 1979 the Advisory Committee on Reactor Safeguards[4] specifically recommended to the NRC chairman that the Commission give consideration to the establishment of quantitative safety goals for nuclear power reactors. The Kemeny[5] and Rogovin[6] reports on the TMI accident also recommended that the Commission move towards the incorporation of quantitative safety into its decision making process.

The Risk Assessment Panel at the IEEE/NRC conference[7] held in early 1980 endorsed the establishment of quantitative safety goals for industry and suggested that the use of risk assessment in the area of licensing be pursued. Recently (late 1980) both the ACRS[8] and NRC's own Office of Policy Evaluation[9] have encouraged the early implementation of risk based decision making.

As discussed in this report the addition of a risk assessment approach to the licensing process will be by no means an easy process. However these difficulties do not over balance the potential benefits that can be expected from the use of probabilistic risk assessment techniques.

The sections which follow address the various problems which can affect the implementation of risk criteria. Each section addresses one level of

the hierarchical structure for risk criteria proposed in Chapter 3 of this report. In this way, Chapter 11 can be correlated with the body of the report, as shown in Fig. 11.1.

The problems may be divided in the most general way into a technological group and a socio-psychological group. The technological group includes problems related to how well the assessed risk describes the real world risk, (i.e. assessed vs. actual risk). The socio-psychological group includes problems related to how well the assessed risk is perceived even if it represents exactly the actual risk (i.e. perceived vs. actual risk). In this report we shall confine ourselves to the discussion of technological problems.

11.2 POTENTIAL PROBLEMS IN THE IMPLEMENTATION OF COMPONENT AVAILABILITY CRITERIA

As has been stated in Chapter 4, a component unavailability criterion can be established at various levels of detail. In the broadest sense, the availability could be stated in a generic form for all components within the plant, independent of use or type. In the most specific sense, it could be applied to each individual component in the plant with component specific values. In general, the demonstration of a criterion would require the implementation of existing models and/or the development of new models to represent component availability. Also required is the acquisition of component availability data. The level of detail in both model and data would vary depending on the level of the specified criteria.

11.2.1 Information Requirements

Although the level of detail in the information requirements depends upon the level of specification chosen, the information set required (i.e. the models and data) includes the consideration of the reliability of the component, the test and maintenance interval for the component, and the effectiveness of the maintenance. To derive failure rate information from the data, the number of components in a group must also be known. Thus models for the component failure rate, maintenance process, and the influence of human error must be employed in the evaluation process.

Fig. 11.1 Risk Criteria Evaluation (According to Hierarchy Proposed in Chapter 3)

Level	Risk Criteria	Information Required: Models	Information Required: Data and Descriptions
1 (Chapters 8, 9, and 10)	Risk Number (Individual & Societal Risk) • Fatalities • Illnesses • Property Damage	• Dosimetric • Ecological • Evacuation etc.	• Population • Demographic • Meteorology etc.
2 (Chapter 7)	Radioactive Release Probability • Radioactive Releases Frequency Amount	• Thermal-Hydraulics • Core Physics • Radioactivity Release etc.	• Accident Sequence • Isolation Limit • Penetration Limit etc.
3 (Chapter 6)	Accident Probability • Core Damage	• Event Trees • Systems Interaction • Safety Systems Model etc.	• Initiating Events • Frequency Data etc.
4 (Chapter 5)	System Availability • System Failure Process Safety Systems	• System Failure Modes & Effects • Fault Trees • System Human Errors etc.	• System Functional Description • System Operational Data etc.
5 (Chapter 4)	Component Availability • Component Failure • Human Error	• Component Hazard Model • Maintenance Time & Effectiveness • Human Error etc.	• Component Failure Rates • Failure Descriptions • Maintenance, Test & Repair Data etc.

Note: Arrows indicate flow of data and information.

11.2.2 Information Suitability

Independent of the level of detail of the data and models, decisions must be made as to the types of models and data that would be considered acceptable in the evaluation. For the evaluation to be properly done, model requirements and data capabilities must be compatible. For example, the failure rate model specified by the criterion might require failure rate information about a component as a function of its age. In this case, data which provide only the total number of failures of a component over some time period would be insufficient. The failure rate model might require that distinction be made between failures per demand, failures per hour in standby, and failures per operating hour. Even if the constant failure rate (with time) model is assumed valid, there may be a region in which this assumption is incorrect. In the maintenance case, the model may require distinction in data on the crew size and the skill mix, the location (in radiation area or not), organizational time, set-up time, actual repair time, and checkout time. In the case of human error, the model chosen may require detailed taxonomies (task breakdowns), and factors which are thought to influence performance (Performance Shaping Factors).

In theory, demonstration programs can run the full range from those which rely entirely upon assessments based on generic* data to those which will accept only specific** data for the component in its installed condition. In the sections which follow, several types of demonstration programs will be discussed along with their accompanying limitations. In all cases, the concept of level of detail in the definition of a component has been intentionally left open by the authors. Level of detail and associated availability of data at each level will be discussed in a separate section. However, at least some distinction between broad generic types of components (i.e. pumps, valves, diesel generators, motors, etc.) is required.

A demonstration program could be based entirely on the use of historical data collected on "similar" components in "similar" use. If a class of components is one which is known to exhibit a long mean time (or number of

*Generic data - all data on a typical component, no matter how that component is used or in what system it is placed.

**Specific data - only data on a particular component used for the same function and in similar or identical systems.

cycles) to first failure (MTTFF) when compared to the operating time (or cycles) involved (i.e. that the probability of multiple failures of the same component is small), or if the component has a rate of failure which is independent of the maintenance and repair process, or if the component is non-maintainable, then this method of demonstration could be an acceptable approach.

If we restrict applicable data collection only to those "identical" components in the "identical" end use, defined previously as specific data, and interpret this to mean that it applies to a component from the same manufacturer, manufactured to the same manufacturing specs, qualified to the same end use and installed in the same reactor design (i.e. same nuclear steam system, same generation, same architect engineer), and employed to perform the same function in the same environment, then the allowable sample for demonstration may become exceedingly small. Even using this restrictive definition does not adequately define all cases. There may be differences within each component category even from the same manufacturer due to batch or lot differences (materials or procedures), or even design changes which do not affect the procurement or manufacturers specifications. Obviously, being too specific (same lot, etc.), would seriously limit the applicability of the data.

If, in order to expand the base of available data, generic data is allowed, what will be the limits of acceptable groupings under the generic titles? As evidenced by a review of current analyses, the concensus appears to be that similar components can be grouped and then thought of as physically and functionally belonging to particular generic populations, but there seems to be no general agreement on how this generic grouping should be accomplished, nor on how to apply these grouped data in specific instances. Some data bases present rather broad generic groupings by collecting "similar" design types under very general groupings (such as MIL Std. 217[10], WASH-1400[1], the NRC PAS*/EG&G proposed safety system component failure rates from LERs[11,12,13,14], and the early versions of NPRDS[15], while others suggest far more detailed distinction requirements prior to grouping (such as IEEE Std. 500-1977[16], IIT Handbooks[17], RADC Notebook[18], NRC/ORNL/SAI[19] In-Plant Data Base). Functional and operational distinctions are sometimes not taken into account (MIL Std. 217, WASH-1400, IEEE Std. 500).

*PAS - Probabilistic Analysis Staff

Grouping data often produces interpretation problems with regard to the uncertainty bounds around the group's estimate of central tendency measures. Since any grouping necessarily implies going from a more homogenous sample to one less homogenous, uncertainty bounds generated based upon the assumption of homogeneity may become less applicable as the definition of the grouping is loosened. As this tendency toward looseness increases, the danger of misapplying the data also increases. For instance, if a mean value is used for a very broad grouping, there is a danger of either understating or overstating the appropriate value for a particular component within that grouping. If the extreme values were to be used, the result might be too conservative in some cases. The data can also be fitted to a distribution (for example, a Log-normal distribution), and the desired confidence intervals of this distribution may be used. Even with these problems, generic data can be useful in early design phases of plant licensing, but is not recommended for later phases

Also, the quality of the recorded data must be considered. Differences in reporting rules can cause errors in:

- which components must be reported
- what is considered a reportable event

Between the two extremes of the totally generic based approach to risk criteria evaluation and an approach based totally on specific data, there exists a wide spectrum of programs which could be developed. Each program could be more or less influenced by either of the two extremes. For example, the definition of "specific test data" could be expanded to include test data on the same component put to the same use in other nuclear plants, or it could be expanded to include manufacturer's test data, or it could be expanded to include any operational data collected on the same component in the same use. However, even when the definition of acceptable data is expanded, there may still not be enough information available. In this instance, the question arises of how to combine the differing types of specific data. Should each data source be treated equally and therefore should all data be "lumped" together, or should some data be considered more "applicable" than others and

therefore be given more credence (for example, Bayesian approaches). If it is decided to discriminate, how should the discrimination be implemented?

11.2.3 Human Error Data

In Chapter 4 of this report, if the component description defined in the criteria includes the direct man/machine interface of the component, this approach could allow the analyst to view the human-caused and hardware-caused failures of a component together as long as the human-caused failures occur within the component definition. In this instance, as with components, a component availability criterion can be set for the component and evaluated either on a generic data, specific data, or combined basis by considering all failures which have been recorded in the data base as a basis for demonstration of the criteria being met. At the present time, programs are underway to utilize and expand available human performance data bases and to develop and validate methods to model human performance. Nuclear industry data from the Licensee Event Reports (LERs) and training simulators are being applied to quantify Human Error Rates (HERs). Various groups are also investigating the use of non nuclear industry data in the prediction of human performance. This approach is viable as far as frequently occuring intrinsic* human errors are concerned. The "cascading and common mode failure" effect of human error, as well as the possible cascading effect of hardware failures, cannot be taken into account properly on the component level and is one of the inherent limitations of the application of a component-level risk criterion. Since this level of the hierarchy does not evaluate the influence of the error on multiple components, higher levels should be considered for common mode failures, (refer to Fig. 11.1).

11.2.4 Data Base Availability

Section 11.3.3 will discuss the fact that the level of detail given in the specification of the criteria would determine the level of detail required in the models, and the consequent level of detail in the data. It is also true that given some model and data source, the more specific the description

*Intrinsic - affecting only component or system operated on.

contained in the criteria, the greater the assurance that the evaluation would represent a true picture of the actual risk. In the most specific case, each component of the plant could have an availability number assigned to it, but this would require individual models and data for each component. Data at this level of detail, even if it were available, represent such a small number of failures that the estimates would have large uncertainty bounds associated with them.

Most data bases which have relevance to a nuclear power plant component availability criterion have been categorized by plant, plant system (function), and generic type of component. It could be argued that regardless of design similarities, the in-plant maintenance program has such an effect on the availability of a component that plant data must be combined only after careful categorization of the in-plant maintenance program and only combining data from plants with similar programs. At present, there are two data sources which may have the potential for the requirements of this approach to be fulfilled -the NPRDs, and the LER based system generated by EG&G. Both of these systems are limited by the fact that the populations only consider safety system components. Further, failure rate data derived from LER's depend on the estimates of component population. The LER and NPRD also require the definition of what is a reportable event, and what is a reportable component to be decided on an individual basis. This flexibility could allow for significant reporting differences from plant to plant. The difficulties of these systems could be lessened by a system based on actual in-plant maintenance reports on all components. Such a system is under development by ORNL/SAI in cooperation with the IEEE volunteer effort to update IEEE Std. 500.[16] However, this data base is still in the research stage of development.

An additional driver of component availability is the system function which the component is performing, the type of service, mode of operation, and the associated environment. Both the NPRDs and the EG&G data base allow for these distinctions to be drawn within their limitations. The published edition of IEEE Std. 500[16] allows distinctions to be made according to mode, and allows environment to be accounted for by utilizing "environmental factors". No allowance is made for components in different types of service,

however the next edition is intended to address this issue. The effort performed by Manning[20] in conjunction with the recent NRC reliability evaluation programs allows for mode distinctions, service type distinctions in some cases (e.g. pumps, control rods, etc.) but does not take into account environment. This base is derived from the EG&G study and therefore suffers from the same limitations involved in using LERs as a data source.

Functional distinctions also imply that the data base should be able to distinguish between failures occuring while operating, during standby periods, and as a result of a demand. The NPRD system and IEEE Std. 500[16] allow for a distinction to be made between failures per hour and failures per cycle. However, this approach groups together under the heading of cyclic failures those that occur during standby periods and are observed during an actual or test-related demand.

Finally, it could be argued that the primary driver of component availability is the generic type of component and that functional and maintenance distinctions, while important, are only secondary. This argument is supported by the fact that component groupings are made precisely to account for functional differences, and maintenance quality is something which can be controlled to some degree by specification of preventative maintenance programs. Acceptance of this argument allows many data bases to have applicability to the evaluation of component availability. Some of the sources are:

- WASH-1400 Appendix III[1]
- EEI Data (Now NERC/GADS)[21]
- UK Systems Reliability Service Data[22]
- Military Handbook 217 B (Now in a "C" revision)[10]
- IIT Notebooks[17]
- Reliability Data from In-Flight Spacecraft[23]
- GIDEP Reliability Maintainability Data Bank[24]
- RADC Notebook[18]

- AMF 66-1 (Air Force)[25]
- 3M (Navy)[26]
- IEEE Industrial Reliability Survey[27]

11.2.5 Summary of Criteria Limitations

Even if the problems of evaluation of component availability were resolved, and this resolution occurred at a level which was considered to be proper, based upon a general consensus, and even if the data base problems were resolved, and even if component availability were tracked operationally on a continuing basis, there would still be limitations on the use of a component availability criteria. Most fundamentally, these limitations arise from system interactions which cannot be accounted for in the individual component criteria. Examples of these are common mode or common cause failures due to external stimuli (seismic, fire, flood, etc.), and the multiple effects of a single failure (cascade effects). The cascade effects includes such things as, a) pump A failure causes pump B and pump C to operate at higher RPM to maintain flow, increasing their failure rate; b) pipe A bursts in such a way that fragment missiles disable operators on valves C & D; c) incorrect calibration is made on instrument A increasing the probability that one or more other instruments in the same calibration set will also be miscalibrated. These interactions cannot be covered by a component level criteria but are accounted for by all of the higher level criteria such as System Availability, Accident Probability, etc. (see Fig. 11.1).

11.3 POTENTIAL PROBLEMS IN THE IMPLEMENTATION OF SYSTEM AVAILABILITY CRITERIA

11.3.1 Evaluation Requirements

Risk criteria at the system function level could be stated in terms of the probability that a system would respond adequately to an incident which required it. The criteria should address the probability of the function responding successfully, initially as well as its continued response over the required time. Thus the evaluation process requires an accurate and complete system description. Information on the components may not be required if sufficient system level failure data are available, without this, systems level information could be synthesized from component level information and systems models which indicate the interactions among the components. Even if the component information includes intrinsic human error (see Section 11.2.3) contribution, and common cause contribution, the extrinsic error possibilities still must be considered at the systems level (e.g. multiple miscalibration during systems test, multiple failures due to component proximity). When the component level information excludes this information, allowance must be made for its inclusion in the system function model.

11.3.2 Information Requirements

The evaluation process required to demonstrate that a criterion imposed upon the system availability has been met must include the following information:

- Models of the system failure probability under both standby and operating conditions,
- models of the probability of the system successfully responding initially to a demand,
- a downtime model for repairable systems, and
- models of periodic testing and monitoring.

11.3.3 Information Suitability

If system availability is calculated and sufficient data are available at this level to support this calculation, the criteria can be demonstrated in theory on the basis of data alone. However, the system failure data base is small even when all operating experience is considered, and the application of system level data from other plants may be questionable due to differences in operating and maintenance philosophies. If the system hardware configuration is considered to be similar, differences in the operating and maintenance philosophies for the plant in question might be accounted for by using bounding calculations.

An alternative is to construct the system behavior in terms of the behavior of its associated components. In this instance a model must be developed that depicts, to the degree of resolution required, the relationship between the components as they effect the performance of the system. There are many such models but they all can be classified according to whether or not they use a deductive or an inductive approach.

Inductive approaches query the system to determine what happens to the performance of the system function if a particular component fails in a particular mode. Examples of inductive approaches are the following:[28]

- Parts Count
- Failure Mode and Effect Analysis (FMEA)
- Failure Mode Effect and Criticality Analysis (FMECA)
- Fault Hazard Analysis (FHA)
- Double Failure Matrix (DFM)

In all cases inductive methods require a great amount of investigation into the system. For those instances where only catastrophic events are of interest, or when the knowledge of the analyst concerning the operability of the system under normal and component failure mode conditions is extensive these investigations will tend to be less than productive. This condition arises because there is no apriori way of <u>not</u> considering individual failure effects which are of little system import or whose probability of occurrence is negligible. However, if degraded operation effects are of interest or if

the system design is new or unconventional the inductive approaches may be essential to the successful performance of later deductive analyses since they provide a mechanism for the analyst to more clearly understand the systems operation. These approaches may also be of considerable value for providing insight into proper deductive analysis if they are performed on the interfaces between systems. Here they can highlight system interaction effects which may be masked by the compartmentalized nature of the deductive approaches.

Deductive approaches are all based upon a process of postulating the occurrence of an event and investigating the potential causes of this event. These approaches allow for the postulation to be made in terms of a success or failure event. Fault tree analysis is a deductive failure analysis which concentrates on the occurrence of one particular undesired event. This event is the top event in a diagram from which logically constructed branches emanate. These branches terminate in blocks which enclose the immediate causes of the top event, and are connected to it via logical symbols which indicate the way in which they (as individuals or in combination) influence the occurrence of the top event. The logical connections to be utilized are chosen based upon the analyst's knowledge of the design and operation of the systems under investigation. The process continues until a branch terminates in an event which either requires no further development, or which is unable to be developed further.

When completed, the fault tree is a model which indicates graphically the various combinations of faults which will result in the occurrence of the predetermined undesired top event. It is therefore a depiction of the logical tie between basic events and the undesired event.

The fault tree approach does not, and is not intended to model all possible causes of system failures. By intent it addresses only contributory faults and only those which are considered auditable. The tree generated by the analysis, while not in itself quantitative, can be evaluated quantitatively by supplying probabilities for the basic events by the use of Boolean algebra.

Since there are many models which are available to calculate the availability of system functions, this discussion must be necessarily limited. Therefore, an attempt has been made to discuss those which are most relevant to the process of evaluating nuclear system function availability.

Computer codes are available to evaluate fault trees. A discussion of the state-of-the-art in computer code methodology is given in References 28 and 29. In the following paragraphs, a brief description of some computer codes will be presented. More complete descriptions are to be found in Reference 28.

Computer codes which have been developed to produce the minimal cut sets of a fault tree include: PREP[30], ELRAFT[31], MOCUS[32], TREEL AND MICSUP[33], ALLCUTS[34], SETS[35], and FTAP[36]. Computer codes which allow for varying basic event input data (i.e. time point for evaluation, varying failure and repair rates represented by distributions) and operate upon the same minimal cut set of the tree in question to provide a spectrum of top event probabilities (unavailabilities) include: KITT 1[30], SAMPLE[37], MOCARS[38], and FRANTIC[39].

These codes produce probabilities of system and individual component failure, quantitative rankings of contributions to system failure, and sensitivity analyses. The KITT and FRANTIC codes compute time averaged and time dependent point estimates of system failure probability. SAMPLE and MOCARS compute a distribution and error bounds for system failure probability based upon uncertainty, error, or variation in the component level characteristics.

There are also computer codes which perform direct quantitative analysis without the intervening step of producing cut sets. These codes include: ARMM[40], SAFTE[41], GO[42], NOTED[43], PATREC[44] and WAM-BAM[45]. The output of these codes is generally in the form of point estimates for the system unavailability or failure probability. GO and WAM-BAM offer the advantage of characterizing complement events and some dependencies.

A computer code which can perform both qualitative and quantitative analysis is the PL-MOD[46] code. This code modularizes the fault tree from a description of its component and gate diagram. A module is formed from

a group of components which act as a supercomponent in that the state of this group is all that is required to determine the state of the system. This code produces occurrence probabilities for the top event and all modules, it has a Monte Carlo option for computing uncertainties, and a time dependent unavailability evaluation option which can handle non repairable, repairable (revealed fault) and periodically tested components.

The last group of codes are attempts toward identifying the common cause failure contribution to system unavailability. These codes include: COMCAN[47], BACKFIRE[48], and SETS[49]. Common cause failures can dominate random hardware failures in certain instances. These codes attempt to identify the system minimal cut sets which have the potential of being triggered by a single, more basic common cause. The programs produce common cause candidates, but require data on the generic cause susceptabilities in order to produce quantitative estimates.

11.3.4 Summary of Criteria Limitations

Although there are many theoretical approaches to the evaluation of system availability, the ones in extensive use for nuclear power plants utilize the fault tree approach, supplemented by analysis of common cause failure and human error. The fault tree approach is well developed and, although the data are not yet always available, programs which are currently underway to collect operating experience should reduce the uncertainties to within reasonable bounds in the near future.

Since fault tree analysis is dependent on the choice of the top event the specification of the criteria must include a detailed specification of the system functions along with the availability requirements. The selection of system functions and the specification itself will be design dependent and this dependency may require specification on almost a plant to plant basis. Multiple specification could cause compliance difficulties. Even with similar specification the fault trees developed depend upon the analyst's approach, again producing compliance verification difficulties. These difficulties can be reduced to a degree by standardization; however, too much standardization would tend to constrain the thought process required for the production of a good analysis.

The residual theoretical limitations are in the treatment of common cause and human error. Both of these areas are represented by models and data that are in the development stages. Use of these models in evaluations can lead to large uncertainties and must be used with caution.

11.4 POTENTIAL PROBLEMS IN THE IMPLEMENTATION OF PROBABILISTIC ACCIDENT CRITERIA

11.4.1 Evaluation Requirements

Chapter 6 discussed the accident probability criteria which specify unacceptable probabilities for some defined accidents. These criteria focus on the frequencies of accident-initiating events and the unavailabilities of safety systems. If the accident probability criterion is formulated to apply to particular accident sequences, then the evaluation procedure intended to show compliance with the criterion, obviously, does not require the identification of these sequences. On the other hand, if the accident probability criterion is formulated on the basis of some classes of accidents, then the evaluation process requires that accident sequences belonging to a certain class be hypothesized.

11.4.2 Information Requirements

The development of the informational data base of events which could initiate situations which might require a response from the systems in the plant to prevent an ensuing accident can be undertaken by a review of the sources of radioactive material. One method that could be used to delineate an accident is to postulate all system responses to a given event and then to calculate expected effects on the sources of radioactive materials. In this case both frequencies of the events and the quantity and type of radioactive materials effected are needed to evaluate the overall importance of a specific accident sequence.

For each of the events postulated, an approach must be utilized which delineates the required responses of the engineered safety features of the plant to provide the safety functions required to mitigate the unchecked consequences of the initiating event. This approach must provide a capability

for modelling the interaction between systems as well as the interaction between the systems and the initiating event. The approach must also provide a means for discriminating the time sequencing of the subsequent required systems responses. A method for the estimation and incorporation of probability estimates of unsuccessful system response must also be available.

11.4.3 Information Suitability

The accident level input data and models of the accident sequence are tied directly to an understanding of the design of the plant's operating systems and its safety systems. This understanding cannot be completely specified in a formulative sense. However, techniques which organize this understanding thereby allowing an examination of the potential challenges and responses in a systematic and logical process are preferrable; their very nature tends to minimize the possibilities that neither the relevant accident initiators considered nor the relevant required system responses will not be overlooked. From the point of view of completeness, approaches are also preferred which utilize historical data to provide an indication of the possible spectrum of initiating events and responses to these events. For those cases where the frequency of the initiator is so small that there is no historical evidence but whose consequences are so large that its occurrence cannot be ignored, alternative approaches are required. All approaches entail extensions of the existing data base to the lower frequency, analytical investigations into possible precursor activities, and extrapolations of the underlying physical processes by reviewing test information at more severe levels of stress.

There are many approaches possible for the identification of accident sequences. One approach is the application of a logical systematic analysis of the successive results of an initiative. This approach has been called "Decision Tree."[50] Its original application was in the area of management decision making. The Event Tree approach, which had its origin in the Decision Tree approach, provides a systematic logical way of defining the possible accident sequences, and with failure probabilities available from lower level

analyses, data at the system level, or from estimates, allows the probability of each sequence to be determined. The WASH-1400 study utilized this approach as the principal mechanism for identifying and analyzing the accident sequences which were considered.

11.4.4 Summary of Criteria Limitations

Whenever an event tree is applied to a real world problem, the number of outcomes produced without bounding the application in some way is prohibitive. Any treatement of necessity cannot consider all the outcomes. The WASH-1400 study chose to use a method that did not allow for partial success and defined failure as complete if there was any uncertainty. Also, system failure was defined as having less than a required fraction of redundant equipment operating. This greatly reduces the branches of the event tree with the resulting error on the conservative side. But in some cases, degraded conditions (or the seriousness of the failures in terms of the function supported) are necessary for the understanding of the accident and so these conditions are not, in general, incorporated. However, this problem can be overcome by the use of an analysis, that allows degraded function to supplement the event trees.

From a practical standpoint, accident level criteria are limited by its lack of consequence estimate to discriminate between accident sequences. Suggestions have been made to utilize a qualitative measure of consequence based upon the extent to which each sequence represents a threat to, or a violation of one or more of the multiple barriers between the fuel and the public. These limitations can be accounted for by the use of higher level criteria as described in the following sections.

11.5 POTENTIAL PROBLEMS IN THE IMPLEMENTATION OF PROBABILISTIC RELEASE CRITERIA

11.5.1 Evaluation Requirements

As discussed in Chapter 7 of this report, criteria for the release of radioactivity can be stated in terms of:

1. Curve of unacceptable probability vs. the amount released.
2. Some measure of the unacceptability curve such as its expected value.

The amount of radioactivity released can be expressed in terms of the total release of radioactivity, or the amount of individual isotopes released. Regardless of the isotopes selected, demonstration of compliance with the criterion requires the identification of events which could lead to a release in order to determine their probabilities and consequent release magnitudes. This evaluation requires a review of potential accident sequences which could lead to such a release. For those sequences which involve possible core damage, information about accident sequence conditions must be obtained and their resultant impact on the core must be evaluated. The following must be addressed: core thermal-hydraulics, the chemical and physical characteristics of the core in terms of its response (i.e. liquefaction, melt, cladding defects) to the resultant thermal hydraulic conditions, and the release of radioactivity as a result of this response. Also, the transport of released radionuclides from the core and into the containment, mechanisms which can cause reduction of the amount actually in containment as opposed to the radioactivity released from the primary envelope, and the capability of the containment to mitigate environmental release must also be ascertained. For those sequences which do not involve significant impact on the core but could lead to releases, an evaluation consistent with the release level is required.

11.5.2 Information Requirements

Demonstration that the release criteria have been met for those accident sequences which could involve core damage requires the utilization of models to supplement the data available on core behavior as discussed in the previous section:

- Efficient thermal hydraulics models
- Chemical and physical process models the for core under accident conditions
- Release of radionuclides from the fuel
- Fission product transport and reduction within the containment
- Containment integrity models

In addition to the above models, the informational data base which is required includes data on:

- Physical design limits of containment
- Isolation limits of containment, (design leak-rate)
- Accident initial conditions

In addition to the data outlined above, the data discussed in the previous sections on system and component availability are also required. These include component and system failure models and data, with emphasis on accident initiators and sequencies. As shown in Fig. 11.1, each successive stage in our hierarchy of criteria depends on and makes use of the data developed for the preceeding stage.

11.5.3 Information Suitability

In addition to simulating the degraded core thermal hydraulics, the models must have the capability to provide estimates of fission products which are expected to be released from the core under accident conditions. This fission product source term model must include estimates of the gap release, the meltdown release, the vaporization release, and the oxidation release components for the various fission products in the core.

For those release components that contribute to the effluent released from the primary envelope through escape of primary system coolant (i.e. the gap and meltdown releases), models must be constructed which account for the mechanics by which the coolant escapes and the interaction of the coolant with the fission products. This model must incorporate the bulk fluid flow during the course of an accident insofar as it acts as a driver to carry fission products (vapors and aerosols) out of the primary system, the deposition and plate out which would occur, and the absorption by water within containment.

Information should include containment pressures and temperatures, radionuclides within the containment, and the leak rates to the atmosphere or puff releases. The models must be sufficiently flexible so that realistic differentiation due to containment and plant design differences are allowed and they should take into account variable initial conditions pertaining to an initiating event.

Limitations of Existing Information

For the case of core degradation, there is no experience available to provide the required information base. However, thermal hydraulic models are available for the prediction of core response. These models are in some cases supported by experimental data, but in many cases the experimental basis is quite limited. Uncertainties can be taken into account assuming the models are correct, if the input parameters are varied sufficiently. However, this "estimated uncertainty" does not account for the assumptions in the model. These must be validated by additional experimental programs. Additionally, there is the problem of lengthy computer running times that could limit this approach. In view of the lack of a firmly established basis for some of the underlying models, the results should be carefully applied. WASH-1400 treated thermal hydraulics with a set of models directed at specific LWR designs. These models also concentrated on large break accidents. The MARCH[51] code addresses the Meltdown Accident Response CHaracteristics so as to include a broad range of LWR designs, extend the applicable range to transients and small break accidents, and provide for an integrated and consistent treatment of the entire accident time period. The MARCH code also attempts to incorporate experimental data that have been developed in the interim. However, recent events indicate that further analytical and experimental work is necessary to realistically model accidents, especially for the cases that involve partial core melt.

The CORRAL (Containment Of Radionuclides Released After LOCA) code[1] was used in WASH-1400 to compute the transport and deposition of radioactive material within the containment and the release of radioactivity to the atmosphere. A revised and generalized version of CORRAL known as CORRAL II[52] is now available. Here also, the uncertainty of the results can be estimated in terms of the response to ranges of input variables.

In summary, the MARCH/CORRAL combined code addresses the very complex problems of analyzing the core melt and consequent radionuclide transport and deposition phenomena, and while its experimental basis is in some cases weak, it at least represents a basis of relative comparison of results and the beginning of future develomental work in this area.

For accident sequences which do not significantly affect the core, the above codes may not be applicable. However, these events have frequencies of occurrence high enough so that there exists a body of historical data available to evaluate the release probabilities. One approach which has been applied to the assessment of a PWR Class 3-8[53] accident risk utilizes the LER data to provide an actuarial model for the assessment of some of these accidents where the data applies, and extends this evaluation via engineering analysis and judgement to lower probability and higher consequence events. The major uncertainty in this approach is due to the many and varied sequences possible for these releases and the uncertainty involved in extending the analysis.

11.5.4 Summary of Criteria Limitations

The residual limitations associated with the demonstration of a release criterion may result from the use of simplified models and conservative assumptions in the evaluation to describe and estimate the release of radioactivity from the fuel and thence to the environment under accident conditions[54]. In the past, it was assumed that iodine escaping from the ruptured fuel rods due to accidents is in the elemental form. Experimental evidence presented in Reference 55, suggests that iodine as a metal iodide escapes from the fuel. This finding would lead to a reduction in the estimate of the environmental release of iodine in comparison with estimates that have been made in the past.

11.6 PROBLEM IN THE IMPLEMENTATION OF RISK NUMBER CRITERIA

As stated in Chapters 8, 9 and 10, the risk number criteria can be formulated to control individual plant risk or total risk from all plants. The techniques available must produce credible estimates for the following risk number criteria:

- Individual and/or Societal Risk Criterion for adverse health consequences of interest such as early and latent fatalities (see Chapters 8 and 9).
- Property damage risk criterion (see Chapter 10) which applies to public property damage (i.e. costs associated with evacuation, relocation, decontamination and interdiction) or it may include plant property damage (i.e. costs associated with cleanup, replacement and requalification of equipment, replacement power, etc.).

The information required to allow risk numbers of interest to be calculated include:

- Radionuclides vs. various exposure pathways
- Dose vs. distance from the release site
- Dose vs. health effects
- Evacuation effectiveness information
- Shielding effectiveness information for the residual population
- Information on the economic consequences of a radioactive release

The development of the above information is carried out by exercising models of individual aspects of the consequences of a radioactive release. The collection of these individual models is often referred to as the "accident consequence model" or simply the consequence model. These models predict the results from individual radioactive releases (as discussed in 11.5). These predictions depend upon:

- How the radioactivity is dispersed in the environment
- The number of people exposed and the extent of property damage
- The effects of radiation on people

Due to the difficulties in modelling the three subjects mentioned above, coupled with the uncertainties discussed previously in this chapter, these calculations are likely to have the largest uncertainties of all and will be further discussed in the following paragraphs.

In the case of input data, even if the data are collected in the exact form needed for input to a selected model, there is an underlying uncertainty in the collection process. Here data quality is an issue. Often practicality requires use of gross measures derived from data rather than the data itself. Errors can result when the consequence calculations must utilize models of the local meteorology and its effect on the atmospheric dispersion of the radioactive effluent, models of the local ecology and pathways to man. In addition, population evacuation models can lead to large uncertainties.

11.6.1 Meteorological Models (Atmospheric Dispersion)

In assessing the consequence of a postulated reactor accident, the major concern (aside from water contamination) is the possibility of an airborne release of radionuclides. In this case an atmospheric diffusion model is used to estimate the subsequent transport, diffusion and removal processes of the aerosols. Most such models, long used in air pollution studies, are based upon the assumption that the concentration profile has the shape of a Gaussian distribution. Several properties of the aerosols can be tracked, such as chemical composition and size, although computational time or lack of knowledge generally requires compromising on the numbers of aerosols and their properties to be followed. For example, the WASH-1400 study selected 54 radionuclides out of several hundred and this number was further reduced for the consequence calcuations.

In dispersion models, the initial release from containment is corrected for plume rise[56] based on the thermal energy release rate, wind speed and atmospheric stability; the turbulence caued by the presence of the building[57] and cloud depletion due to decay and dry and wet deposition depending upon the initial meterology. The duration of release (0.5 to 10 hours in WASH-1400) accounts for horizontal dispersion due to wind meander.

The effluent concentration now depends upon meterological parameters such as the time and space variability of the wind, atmospheric stability, mixing depth and precipitation, and aerosol parameters such as decay, fall velocity and deposition. Except for decay these factors can be quite variable. In Gaussian models the standard deviations of the crosswind and vertical distributions of the contaminant, σ_y and σ_z respectively, are determined by experimental observations for several atmospheric stability categories. In Regulatory Guide 1.23[58] the stability categories are divided into six classes ranging from extremely unstable to very stable (A to F) based upon temperature lapse rate data. Typical recommended formulas for σ_y and σ_z are cited in Reference 59. However, vertical diffusion is limited by the mixing height whereafter only horizontal diffusion remains effective. This height depends upon the underlying stability and the season[60].

The depletion of the plume is also computed in diffusion models. Radioactive decay is adjusted for travel time by a well understood model that also includes the growth of radioactive daughters. Dry and wet deposition processes are less well understood because of the chemical composition, surface type, atmospheric stability and many other factors. In the WASH-1400 model dry deposition was assumed to occur at all times with a constant deposition velocity of 10^{-2}m/sec for all nuclides except the noble gases. The wet deposition process is a simple exponential depletion with two constants, one for stable conditions and a larger constant for convective storms for both particles and gases except noble gases.

In the WASH-1400 model several assumptions are made about the plume expansion. For example, hourly changes in wind speed, stability class and mixing depth at the release site are assumed to occur simultaneously at all downwind locations of the plume; the wind direction does not change and the effect of vertical wind shear is not considered; and the crosswind Gaussian shape is replaced with a uniform square wave or "top hat" function which has an amplitude within 20% of the Gaussian peak. Also any remaining radioactive material (except noble gases) is deposited after 500 miles uniformly to 2000 miles. These assumptions are justified by an overall consideration of the state-of-the-art knowledge of the factors used in the consequence model. Atmospheric dispersion models can be much more sophisticated depending upon the meteorological data and computer time available.

11.6.2 Ecological Models

In addition to the health consequences of land contamination, the effect on the local ecology must be considered. Contamination with radionuclides can make land unsuitable for agriculture even if it is considered suitable for human occupation[2]. Time and distance thresholds have a great effect upon the estimates of property damage. The application of absolute thresholds while significantly simplifying the calculations can also lead to significant over or underestimate of the effects.

11.6.3 Dosimetric Models

Using the above described models, estimates can be made of both the atmospheric and ground surface concentration for each radionuclide as a

function of space and time following a given release. Exposure to radiation can be classified into early and chronic exposure depending on the time over which exposure occurs. The former dominate the direct exposure to the effluent release cloud due to its radionuclide inventory and includes the external dose and the internal dose. In addition the release may result in ground contamination with consequent long term exposure due to long lived isotopes (Cs 137, Sr 90). The long term exposure modes are irradiation from radioactive material deposited on the ground, inhalation of resuspended radioactive material from ground deposits, and ingestion of contaminated food and water.

11.6.4 Epidemilogical or Health Effects Models

Once the dose from releases is assessed the health effects that might be associated with a hypothetical release must be calculated. The health effects can be generally divided into three categories:

- early and continuing somatic effects
- late somatic effects
- genetic effects

The early somatic affects include illness and early fatalities which would result from large doses, and these exhibit themselves for a period of a few days to several months after exposure. Individuals who recover from early illness, and those who may have received doses too small to produce illness may still be vulnerable to late somatic effects. These effects include latent cancer fatalities, and morbidities and the development of benign nodules in the thyroid. These varying health effects take place between 2 and 30 years after exposure. An exposure to radiation may increase the frequency of mutations in cells of the exposed individual, and this may result in genetic disorders.

11.6.5 Demographic and Econometric Models

In order to determine the societal impact, individual health effects should be considered for the exposed population. In addition, the property damage effects of a release must be evaluated in terms of the contaminated land area.

Population data can be obtained from census information, but must be correlated to the specific site in terms of the density variation around the site, and the change in population between the time of the estimate and the time of the census. The distribution of population around a plant site must be estimated. In order to reduce calculational requirements the calculation is cutoff at some distance from the plant where the concentration becomes relatively small.

The property information such as land use fractions, and land use characterizations can be obtained from the U.S. Statistical Abstracts and the County and City data book*. The latter source gives an estimate of the value of property in active agricultural use. The value for nonagricultural land and property other than land is not available from this source and must be estimated. In WASH-1400 costs associated with non-agricultural assets were assumed to be proportional to the population. The calculation of relocation duration and the costs due to contaminated land are more or less dependent upon the success of post accident decontamination.

11.6.6 Mitigation Effectiveness Models

The above models produce outputs in terms of deposited radionuclides given a radioactive release, and also give the resultant health effects and property damage. Both naturally occurring events and human efforts could mitigate the exposure from a given release. These actions or efforts include evacuation, terrain self-shielding, shielding from buildings, and decontamination. In order to estimate societal risk, these considerations must be included. The mitigating measures can be divided into immediate and long term.

The possible immediate actions taken (within hours of the accident) include evacuation, sheltering, and potassium iodide pills. Evacuation effectiveness models depend upon estimates of the warning time available. The time of the action depends upon the local meteorological conditions and the evacuation effectiveness. The EPA 1974[61] study is the most often quoted recent work on evacuation. The WASH-1400 study fits these data to a log normal model of evacuation speed to estimate effectiveness. These models assumed that

*Annual publications of the U.S. Dept. of Commerce

there would be a time period of warning prior to the release. Due to the potential of indecision by government officials in issuing the evacuation order the assumption of alert time in the calculation can be questioned.

For the residual population, not evacuated, remaining indoors can reduce the potential inhalation dose. The amount of reduction depends on, among other factors, the ventilation rate of the building. In the Reactor Safety Study (WASH-1400), the reduction in inhalation dose was not calculated because the study concluded that little reduction in the dose, averaged over a large population, was expected since the public would be unfamiliar with the protective measures and may not take the appropriate actions that would lead to a reduction in the dose.

The shielding afforded by buildings was included in WASH-1400, based upon estimates of the shielding factors provided by representative structures and the population expected to be indoors.

However, these shielding models involve studying the exposure time, building use, and building type distribution of the individual sites before the site related effects can be evaluated.

There are principally two methods of mitigating the long-term effects of radiation exposure due to effluent release, interdiction and decontamination. Interdiction is the denial of land and its improvements for normal intended use. In the case of radiation, the use of land can be prohibited until the radiation dose is decreased below some specified criterion. Interdiction of exposed land can be in the following four areas:

- Total denial of land and assests for a long period
- Limited denial of land for a few years
- Denial of crops grown on the land
- Denial of milk produced

The decisions in the first two areas are based upon external radiation doses, and the second two on the doses received from the ingestion of contaminated food. The most restrictive contamination criteria is on milk, and the largest interdicted area is associated with milk impoundment.

The second mitigation method is decontamination. Decontamination includes the removal of radionuclides by physical removal, stabilization in place, and environment management. The effectiveness of decontamination depends upon the surface type, the level of exposure, and the procedures used. These are expressed in terms of decontamination factors. Decontamination factors are established based upon experience with specific surfaces and the effectiveness of available measures on these surfaces. The factors are then applied to the distribution of surfaces present in the exposed area. Site specific information is desirable when estimating decontamination factors, however this might not be readily available. Since subsequent to an accident the contaminated land and property which cannot be decontaminated to an acceptable level must be interdicted, the uncertainty of both the interdiction and decontamination models depends upon the uncertainty in the site information used.

11.6.7 Summary of Criteria Limitations

In summary, the uncertainties in the overall consequence calculation which must be taken into account to determine if an assessment is valid. These uncertainties may present a problem when implementing a safety goal. The calculations depend on the physical characteristics of the release, the evacuation speed, population density and distribution, and mitigation efforts.

The uncertainties in the evacuation models and cost can be reduced by comprehensive site specific information.

In the area of high consequence events it appears that there are less significant problems in the implementation of a societal risk criterion as long as estimates within an order of magnitude are considered acceptable. However, in the low consequence area the associated large uncertainties could cause problems in evaluating health consequences. These consequences result from much more probable and realistic events and therefore should be considered. However, the uncertainty in the low release regime is one of the primary weaknesses in the exclusive use of a societal risk goal.

11.7 APPLICATION OF A RISK CRITERION TO THE NUCLEAR POWER PLANT LICENSING DECISION MAKING PROCESS

11.7.1 Introduction

The previous sections in this chapter discussed the manner by which evaluations can be carried out to demonstrate compliance with different types of criteria in our hierarchy. At each level, the strengths and weaknesses of the evaluation were discussed in terms of both the risk elements whether addressed or not addressed, and in terms of the underlying uncertainties in the evaluation process. Another important aspect of the implementation of a risk criterion which must be addressed is the manner in which it might be applied. In order to investigate the applications of a criterion, it is assumed that risk assessments have been performed by the licensees and these assessments produce an estimate of the appropriate individual plant risk measures that are consistant with the level of the specified criteria. The question of application of the criterion now involves a process of review of the completed assessments, and decision making based upon the results of this review.

11.7.2 Treatment of Uncertainties in Risk Assessment in Relation to Criterion Compliance and Criterion Detail

In this section, the question of the measures selected to indicate compliance and their implications will be discussed. In all cases in the discussion, uncertainty is meant to be the assessed uncertainty presented in the licensee's risk evaluation, and the limits of the assessed uncertainty are given by a certain confidence bound (say, for example, 90%). The confidence bounds, although double sided, are only discussed in terms of one side (toward unacceptability) since the criterion is a one-sided criterion (an unacceptability criterion). For simplicity, the implications of the discretionary region are not discussed.

First consider an unacceptability criterion established on the basis of the central measure of the risk evaluated by the licensees. In this instance, three representative cases could result. These cases are portrayed in Figure 11.2. In Case A, the estimated central measure could produce a risk value

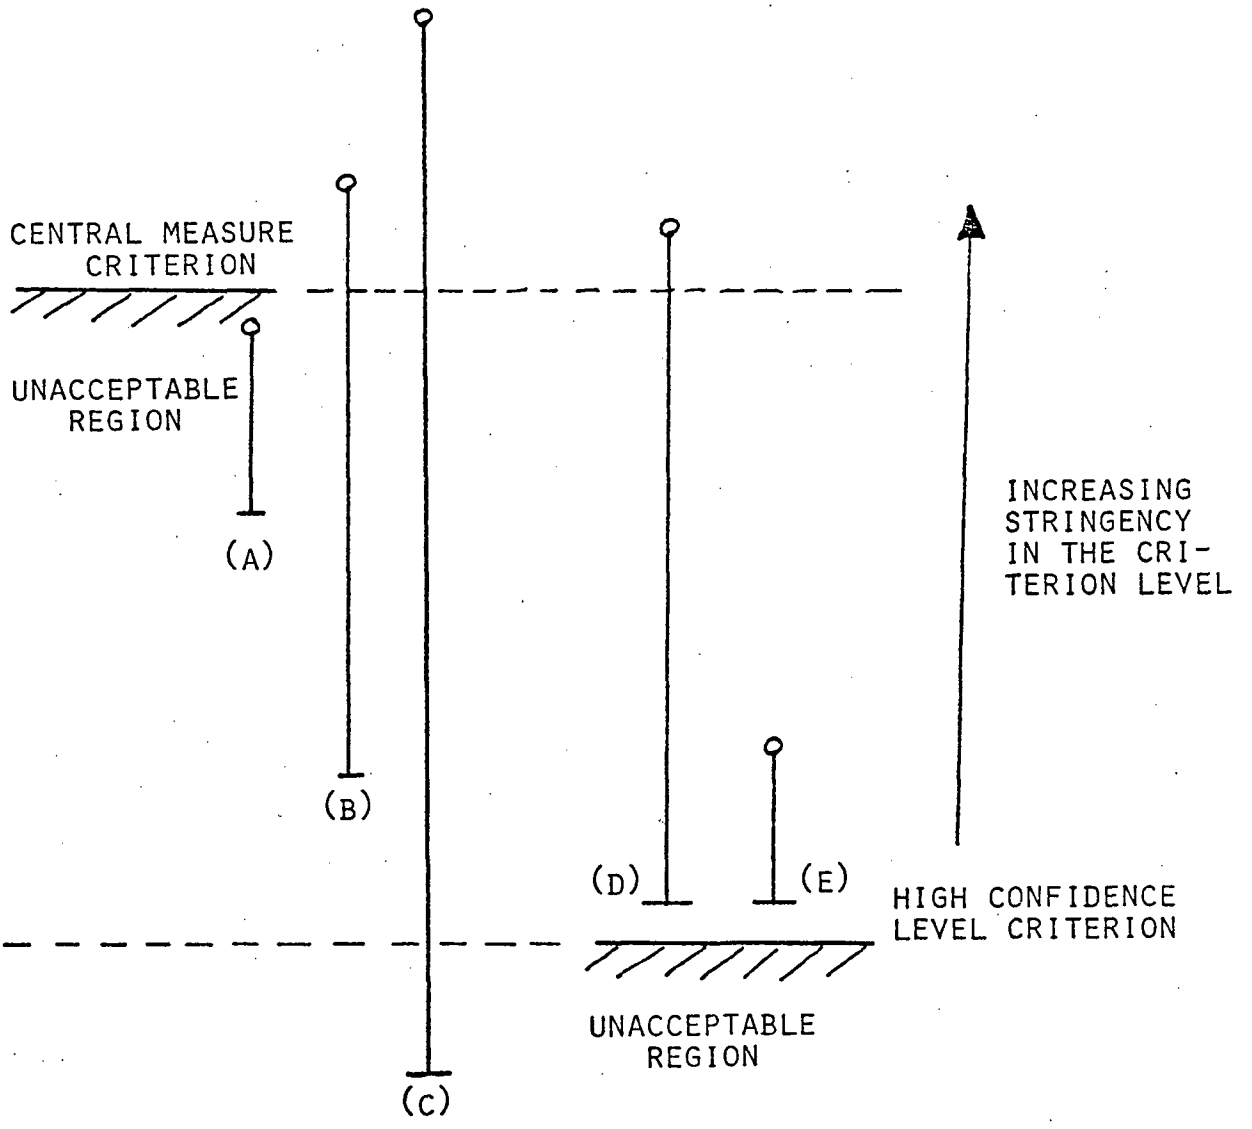

(A), (B), (C), (D), AND (E) ARE ASSESSMENT RESULTS

Fig. 11.2 Treatment of Uncertainties in Risk Assessment in Relation to Criterion Compliance and Criterion Detail.

greater than the central measure criterion, but this central measure is tightly bounded in that the 90% confidence limit is very close to the central measure value. In Case B, the estimated value is less than the criterion, but the 90% confidence bound is well beyond the central measure and well into the unacceptable region, and higher than the 90% value of Case A. Using only a central tendency measure would require that A be considered unacceptable, and B not be considered unacceptable, even though the upper bound of Case B is much higher than the upper bound of Case A. Case C suggests an even more extreme instance where the uncertainty is so large that the 90% confidence bound falls even beyond that of Case B in the unacceptability region, but the estimated central measure risk value is far less than even that of B. In this case the estimate essentially tells us to place a minimum of belief in the central measure value, but the criterion indicates that the plant is not unacceptable.

If an unacceptability criterion is established on a high confidence bound of the evaluated risk estimates (again say 90%) it would be set at higher risk value than the central measure value (see Figure 11.2). In this instance, Case D and Case E, both have the same 90% confidence bound, but the central measure of E is very close to the 90% bound (and well into the unacceptability region established by the central measure criterion) and central measure of D is much smaller than the 90% bound and out of the unacceptability region as defined by the central measure criterion. The use of a high confidence level bound would treat both of these cases as acceptable, even though E represents a case where the central measure is much higher than that of Case D. The establishment of multiple criteria could resolve some of these problems. For example, if the criterion were specified such that the assessed central measure value must be less than X, and the 90% confidence value must also be less than Y (where $Y > X$). This would make Cases A, C, and E unacceptable, whereas Cases B and D would not be unacceptable.

11.8 CONCLUSION

In Chapter 11 the authors have attempted to discuss the obstacles that could be found when applying probabilistic safety criteria to the licensing of

commercial nuclear power plants. These obstacles do not represent insurmountable barriers to the use of the techniques, but instead identify the presently existing weaknesses in the governing models. In every engineering endeavor, various problems must be identified before assumptions can be made and impediments removed. Knowing what is required enables a plan of action to be developed to realistically achieve the goal; and with this in mind, we have set out to list the possible direction that will lead to the desired results.

Chapters 4 through 10 of this document address the direction whereas, Chapter 11 is an attempt at the identification process and reviews the present limitations of each goal. Probabilistic risk assessment techniques are powerful tools for the engineer that must also be used within their limitations. These techniques, along with the presently used deterministic methods, should prove to optimize the technical evaluations needed in the existing licensing process of commercial nuclear power plants.

Chapter 11 References

1. "Reactor Safety Study - An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants," USNRC, WASH-1400, NUREG-75/014, Washington, D.C., 1975. *

2. "Report to the American Physical Society by the Study Group on Light-Water Reactor Safety", Reviews of Modern Physics, Vol. 7, Suppl. No. 1, Summer 1975

3. Lewis, H.W., et al, "Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission, NUREG/CR-0400, September 1978. **

4. Carbon, Max W., Chairman, Advisory Committee on Reactor Safeguards, letter to Joseph M. Hendrie, Chairman of the Nuclear Regulatory Commission, May 16, 1979.

5. Kemeny, J.G., et al, "Report of the Presidents Commission on the Accident at Three Mile Island", October 1979.

6. Rogovin, M., et al, "Three Mile Island, A Report to the Commissioners and to the Public", Nuclear Regulatory Commission Special Inquiry Group, NUREG/CR-1250, January 1980.***

7. IEEE-NRC Conference on Advanced Electrotechnology Applications to Nuclear Plants, IEEE Cat. No. TH0073-1, January 15-17, 1980, Washington, D.C.

8. An Approach to Quantitative Safety Goals for Nuclear Power Plants, ACRS, USNRC, NUREG-0739, October 1980.***

9. "Plan for Developing a Safety Goal", Office of Policy Evaluation, Office of the General Counsel, USNRC, NUREG-0735, October 1980.***

10. MIL - Handbook 217C, "Reliability Stress and Failure Rate Data", Department of Defense.

11. Hubble, W.H, Miller, C.F., "Data Summaries of Licensee Event Reports of Valves in U.S. Commercial Nuclear Power Plants - January 1, 1976 to December 31, 1978," NUREG/CR-1363, EGG-EA-5125, May 1980.***

12. Sullivan, W.H., Poloski, J.P., "Data Summaries of Licensee Event Reports of Control Rods and Drive Mechanisms at U.S. Commercial Nuclear Power Plants - January 1, 1972 to April 30, 1978," NUREG/CR-1205, EGG-EA-5044, January 1980.***

13. Hubble, W.H., Miller, C.H., "Data Summaries of Licensee Event Reports of Control Rods and Drive Mechanisms at U.S. Commercial Nuclear Power Plants - January 1, 1972 to April 30, 1978," NUREG/CR-1331, EGG-EA--5079, February 1980.***

14. Poloski, J.P., Sullivan, W.H., "Data Summaries of Licensee Event Reports of Diesel Generators at U.S. Commercial Nuclear Power Plants - January 1, 1976 to December 31, 1978," NUREG/CR-1362, EGG-EA-5092, March 1980.***

15. "Nuclear Plant Reliability Data System (NPRDS)", 1979 Annual Report of Cumulative System and Component Reliability, Joint Publication of ANS 58.20, APPA, EEI, TVA, and USNRC, NUREG/CR-1635, USNRC, Washington, D.C., September 1980.***

16. "IEEE Guide to the Collection and Presentation of Electrical, Electronic, and Sensing Component Reliability Data for Nuclear Power Generating Stations", Institute of Electrical and Electronics Engineers (IEEE), IEEE Std. 500-1977, 1977.

17. Feduccia, A.J., "Reliability Data Books", Reliability Analysis Center RADC/RBRAC, Griffiss AFB, NY: (1) MDR-1(4) Digital Detailed Data, (2) MDR-2(6) Linear/Interface Data, (3) MDR-7 Memory/LSI Data, (4) MDR-8 Digital Failure Data, (5) MDR-9 Hybrid Circuit Data, (6) DSR-2 Transistor/Diode Data, and (7) NPRD-1 Nonelectronic Parts Reliability Data.

18. Cottrell, D.F., et al, U.S. Air, Rome Air Development Center RADC Nonelectronic Reliability Notebook, RADC-TR69-458, 1970.

19. Penland, J.R. et al, Interim Report - In-Plant Reliability Data Base Development, ORNL/SAI, Report No. 62B-13819C/62X-06, Oak Ridge, TN, November 1980.

20. Manning, F., "Component Failure Rate Analysis", USNRC (private communication).

21. Neibo, R., National Electric Reliability Council/Generating Availability Data System (NERC/GADS), Princeton, NJ, 1980.

22. Systems Reliability Service Data Bank, U.K. Atomic Energy Authority, Culcheth, Warrington, WA34NE, England.

23. Planning Research Corporation, Reliability Data from In-flight Spacecraft, Los Angeles, CA.

24. Government Industry Data Exchange Program (GIDEP), FARADA Failure Rate Data Manual Statistically Analyzed, U.S. Naval Fleet Missile Systems Analysis and Evaluation Group, Corona, CA, February 1971.

25. U.S. Air Force, Wright Patterson Air Force Base, Maintenance Experience Data (AMF 66-1) H.Q. A.F. LOG, Command, Wright Patterson Air Force Base, Ohio.

26. U.S. Navy, Maintenance Material Management Data, (3M), Naval Maintenance Command, Mechanicsburg, Penna.

27. Institute of Electrical and Electronics Engineers (IEEE), "IEEE Industrial Survey - Report on Reliability Survey of Industrial Plants Part I", IEEE Transactions on Industry and Applications, Vol. IA-10, IEEE, NY, March -April 1974.

28. Vesely, W.E. et al, "Fault Tree Handbook, USNRC, NUREG-0492, January 1981.***

29. Fussel, J.B. and Arendt, J.S., "System Reliability Engineering Methodology: A Discussion of the State of the Art", Nuclear Safety, Vol. 20, No. 5, September-October 1979.

30. Vesely, W.E. and Narum, R.E., "PREP and KITT Computer Codes for the Automatic Evaluation of a Fault Tree", Idaho Nuclear Corporation, IN-1349, August 1970.

31. Semandres, S.N., "ELRAFT - A Computer Program for the Efficient Logic Production Analysis of Fault Trees", IEEE Trans. on Nuclear Science, Vol. NS-18, No. 1, pp 481-487, February 1971.

32. Fussel, J.B., Henry, E.B. and Marshall, N.H., "MOCUS - A Computer Program to Obtain Minimal Sets from Fault Trees, Aerojet Nuclear Company, USAEC Report ANCR-1156, August 1974.

33. Pande, P.K., Spector, M.E. and Chaterjee, P., "Computerized Fault Tree Analysis: TREEL and MICSUP", Operations Research Center, University of California, Berkeley, ORC 75-3, April 1975.

34. VanSlyke, W.J. and Griffing, "ALLCUTS - A Fast, Comprehensive Fault Tree Analysis Code", Atlantic Richfield Hanford Company, Richland, Washington, ARH-ST-112, July 1975

35. Worrel, R. and Stack, D.W., "A SETS User's Manual for the Fault Tree Analys, Sandia Lab., Albuquerque, New Mexico, NUREG/CR-0465, SAND 77-2051, November 1978.**

36. Willie, R.R., "Computer Aided Fault Tree Analysis", Operations Research Center, University of California, Berkeley, ORC-78-14, August 1978.

37. "Reactor Safety Study - An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants", Appendix II, USNRC, WASH1400, NUREG-75/014, October 1975.*

38. Matthews, S.D., "MOCARS: A Monte Carlo Simulation Code for Determining the Distribution and Simulation Limits", EG&G Idaho, Inc., TREE-1138, July 1977.

39. Vesely, W.E., and Goldberg, F.F., "FRANTIC - A Computer Code for Time Dependent Unavailability Analysis", USNRC, NUREG-0193, October 1977.**

40. McKnight, C.W., et al, "Automatic Reliability Mathematical Model", North American Aviation, Inc., Downey, California, NA66-838, 1966.

41. Garrick, B.J., "Principles of Unified Systems Safety Analysis", Nuclear Engr. and Design, Vol. 13, No. 2, pp 245-321, August 1970.

42. Gately, W.Y., Williams, R.L., "GO Methodology - System Reliability Assessment and Computer Code Manual", Kaman Sciences Corp., Colorado Springs, Colorado., EPRI-NP-766, May 1978.

43. Woodcock, E.R., "The Calculation of Reliability of Systems: The Program NOTED", UKAEA Authority, Health and Safety Branch, Risley, Warrington, Lancashire, England, AHSB(S) R153, 1971.

44. Koen, B.V., et al, "The State of the Art of PATREC: A Computer Code for the Evaluation of Reliability and Availability of Complex Systems", presented at the National Reliability Conference, Nottingham, England, September 1977

45. Leverenz, F.L., Kirch, H., "User's Gude for the WAM-BAM Computer Code", Science Application, Inc., Palo Alto, California, EPRI-217-2-5, January 1976.

46. Olmos, J. and Wolf, L., "A Modular Approach to Fault Tree and Reliability Analysis", Department of Nuclear Engineering, MIT, MITNE-209, August 1977.

47. Burdick, G.R., Marshall, N.H. and Wilson, J.R., "COMCAN-A Computer Program for Common Cause Analysis", Aerojet Nuclear Company, ANCR-1314, May 1976.

48. Cate, C.L. and Fussel, J.B., "BACFIRE - A Computer Program for Common Cause Failure Analysis, Nuclear Engineering Dept., Univ. of Tennessee, Knoxville, NERS-77-01, May 1977.

49. Worrel, R.B. and Stack, D.W., "Common Cause Analysis Using SETS", Sandia Laboratories, Albuquerque, New Mexico, SAND 77-1832, 1977.

50. Raiffa, H., "Decision Analysis, Introductory Lectures on Choices under Uncertainty", Addison-Wesley Publishing Company, Second Printing, July 1970.

51. Wooton, R.O. and Avci, H.I., "MARCH (Meltdown Accident Response Characteristics) Code Description and User's Manual", USNRC, NUREG/CR-1711, BMI-2064, October 1980.***

52. Burian, R.J. and Cybulskis, P., "CORRAL II User's Manual", Battelle Columbus Laboratories, Columbus, Ohio, January 1977.

53. Hall, R.E., et al "A Risk Assessment of a Pressurized Water Reactor for Class 3-8 Accidents", NUREG/CR-0603, BNL-NUREG-50950, October 1979.***

54. Levenson, M. and Rahn, F., "Natural Limits on the Dispersal of Radioactivity in Nuclear Accidents", Electrical Power Research Institute (EPRI), November 1980, (Available from Nuclear Power Division, EPRI, 3412 Hillview Avenue, Palo Alto, California 94304).

55. Campbell, D.O., Malinauskas, A.P., and Stratton, W.R., "The Chemical Behavior of Fission Product Iodine in Light Water Reactor Accidents", (to be published in the May 1981 issue of Nuclear Technology).

56. Briggs, G.A., 1969, "Plume Rise," U.S. Atomic Energy Commission, Critical Review Series.

57. Slade, D.H. (Ed.), 1968, "Meterology and Atomic Energy 1968," U.S. Atomic Energy Commission, Oak Ridge, Tennessee, TID 24190.

58. U.S. Nuclear Regulatory Commission, Regulatory Guide 1.23, "Onsite Meterological Programs".

59. Gifford, F.A., 1976, "Turbulent Diffusion Typing Schemes: A Review", Nuclear Safety, 1976, Vol. 1976, Vol. 17, No. 1, pp 68-86.

60. Holtzworth, G.C., 1972, "Mixing Heights, Wind Speeds, and Potential for Urban Air Pollution Throughout the United States," Publ. No. AP-101, USEPA, Office of Air Programs, Research Triangle Park, N.C.

61. Hans, J.M., Jr., and Sell, T.C., "Evacuation Risks - an Evaluation", USEPA, EPA-520/6-74-002, 1974.

*Available free upon written request to the Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, DC 20555.
**Available for purchase from the National Technical Information Service, Springfield, VA 22161.
***Available for purchase from the NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission, Washington, DC 20555, and/or the National Technical Information Service, Springfield, VA 22161.

12. MANAGING A QUANTITATIVE RISK CRITERIA

12.1 INTRODUCTION

In the areas of nuclear power plant design, construction, and operation, there exists a large body of technical standards, specifications and guidelines. These include, but are not limited to the ASME codes, and the IEEE and ANS standards. These standards and specifications, while specifically applicable to nuclear power plants, are patterned after, and indeed some are derived from, the existing codes and standards for fossil fuel power plants. In general these standards represent the collective efforts of the profession and are promulgated or published by the engineering societies interested in the particular discipline they represent. (i.e., ASME for mechanical equipment, IEEE for electrical, etc.).

Using these codes and standards, the various regulatory agencies which have cognizance are able to implement their use in licensing procedures, both for construction and operation of the particular installations under their jurisdiction. In addition to these standards, most regulatory agencies have added many more regulations which must be met before the applicant is allowed to build and operate a power plant, be it a nuclear reactor, where the NRC has primary jurisdiction, or a fossil fuel power plant where the states have primary jurisdiction, but the EPA has requirements which must also be satisfied.

These standards and the manner in which they are implemented are not static phenomena and are continually evolving. Both aspects are constantly under review and subjected to revision.

In summary, in the areas of power plant design, construction, and operation, there exists both a large body of technical standards, and mechanisms which can be promulgated, altered, improved, and implemented. In the area of risk assessment, the standards do not exist in a usable form nor is there presently a mechanism for the implementation of these standards. Therefore, while it is possible to promulgate risk criteria as discussed in the preceding chapters of this report, there is at present the need to review and establish agreed upon mechanisms for their use in the decision process of licensing nuclear power plants.

Implementing the use of risk criteria in the NRC decision process should be accomplished with several realities in mind. First, the nature of risk criteria is that these generally impose a requirement of what is unacceptable in terms of risk without reviewing the benefits from the continued operation of a plant; secondly, the economic impact of the decisions that may result from the use of risk criteria are potentially severe. Finally, although the risk criteria themselves may be clearly defined and exact, the risk analyses which will be used to judge acceptability or non acceptability of the risk have the potential, as in all analyses, to include biased judgement. Biases could lead to potential conflicts which could affect risk estimates in both economic and societal ways. Recognition of the existance of these factors is the first step toward mitigating the foreseen conflicts. In fact, recognition of these potential conflicts has occurred, and several organizational techniques have been identified by the community as a means of reconciling the problem of managing the use of risk criteria. The objectives of this chapter are to first examine the characteristics of risk analysis techniques that make them potentially susceptible to implicitly biased assumptions and judgements, and then to discuss some of the organizational techniques that have been proposed to mitigate this problem. The objective here is not to recommend implementation of one or another of the organizational techniques; rather, an attempt is made to present and address both sides of the issues.

12.2 EXISTING PROBLEMS IN THE FIELD OF PROBABILITIES RISK ASSESSMENT

At the present time, no set of techniques has been universally accepted as the methodology by which risk is to be estimated. Approaches to estimating risk range from estimates of upper bounds on core melt accident frequencies based on the number of years of operating experience, to Baysian techniques that account for, among other things, subjective expert opinion, to the detailed decision tree (fault tree/event tree) based methodology that was used in WASH-1400.[1] Notwithstanding the variety of possible approaches, one can map out the general characteristics of a methodology acceptable to at least the NRC and a large portion of the risk analysis community. In the past, this has been the decision tree approach where possible sequences

leading to adverse effects on the public are first identified using decision trees, and then the frequency of each sequence is estimated using probabilistic analysis. Consequences to the public are estimated based on the characteristics of each sequence of failures. Thus, the frequency of each sequence from a spectrum of possible consequences results from this analytical approach. There appears to be a general consensus that realistic as opposed to conservative or upper bound estimates of both consequences and the frequency of consequences should be attempted, and the acceptable risk criteria be based on this premise.

As discussed earlier in this report there is always some degree of uncertainty in the estimates of consequences and frequencies of consequences which can be orders of magnitude in range. Therefore, a concomitant problem in estimating the frequency of accident sequences is to develop some type of bounds for the frequency. Assuming these bounds can be developed, the output of a risk analysis is a statement of consequences, frequency of consequences, and uncertainty bounds of this frequency. Thus risk criteria, as a minimum, should address the consequences, and both the frequency of accident sequences and the uncertainty bounds on that frequency.

There is some concurrence on methods for estimating the frequency of accident sequences, but numerous assumptions and judgement are required in a practical risk assessment. There are only very vague rules or measures of sufficiency or completeness of a risk analysis. Therefore, considerable leeway is given to the risk analyst who must identify accident sequences and estimate frequency and bounds on frequency of these sequences. To some extent, the risk analyst chooses from possible assumptions according to his or her experience. For instance, the analyst may choose to assume either that certain hardware faults should be characterized by coupled (common cause) failures, or that these should be considered to fail independently of one another. Different analysts, representing perhaps different institutional interests, cannot be expected to make the same choices and assumptions since few generally accepted guidelines are available.

A similar situation exists for the evaluation of human faults in a risk analysis. As stated in NUREG/CR1278[2] the assessment of human errors is

greatly facilitated if the assessment is done by a human factors expert. Human factors analysts capable of contributing to the assessment of human errors are often not available to a risk analysis team. Coupling between human faults of a similar nature (common mode human errors) is another area where both data and methodology are sparse. In certain situations recovery from faults (both human and hardware) may be assigned a probability; alternatively the analyst may decide conservatively that no credit for possible recovery should be given. Again, the choices and assumptions are left to the risk analyst, and different analysts may make different choices and assumptions.

In some cases, it is not clear whether or not certain system failures should enter into the risk analysis. Consider, for instance, a safety grade component cooling system whose function is to supply cooling water to several safety related (normally not operating) pumps in several safety systems. The question that arises in a risk analysis is, does failure of the component cooling system result in failure of the associated safety grade pumps? No single simple answer may be correct for this question; and the correct answer may be different for different pumps and for different reactors and may depend on pump design. Information for any given pump that would provide the correct answer may not be available to the analyst, or may not be available at all. Two gross assumptions are possible: that the failure of component cooling will result in pump failure almost immediately, and that the pump will not fail in the time period in which it is required. The real answer may be somewhere in between, such as the pump will operate for some length of time before failing. If this is the case, component cooling may be recovered before the pump fails. But the probability of recovering component cooling will depend to some extent on the mechanism by which component cooling failed. Therefore, one should convolve two probability density functions (pdfs): one describing the probability of pump failure without component cooling versus time, the other describing recovery of component cooling versus time, to obtain the probability of pump failure from loss of component cooling versus time. (Convolutions of this type has been termed "partial failure analysis".) One could even go so far as to specify different accident categories depending on whether the pump failed early or late, if this was appropriate. However, specifying the above pdfs presents another problem to the analyst, since data may not be available to obtain reasonable estimates of the pdfs. Judgement is required,

assumptions are necessary, and the potential for institutional bias is apparent. The two "simplifying assumptions" that failure of component cooling does or does not fail the pump and in the time period in which it is required to operate, lead to two results which may greatly influence the risk analysis. In the first case, failure of component cooling may be a dominant contributor to risk, since the component cooling system services several safety systems. In the second case, failure of component cooling does not even enter into the analysis. The available systems data may not be sufficient to indicate which assumption is more nearly correct. If the conservative assumption is made, the analysis may overestimate the risk, while the frequency of accident sequences may be underestimated if the non-conservative assumption is made. One could, of course, attempt to obtain data on failure of specific pumps versus time after loss of component cooling. This requires both time and resources. Obtaining information on recovery of component cooling, in an accident situation, where the system may have failed for any number of different reasons, is a more difficult problem and fraught with the necessity of making engineering judgements and assumptions. The result is that there is presently no escaping the requirement for engineering judgement and assumptions in a risk analysis, and some of these judgements may have a large impact on the estimated frequency of accident sequences.

It is possible to estimate the frequency of accident sequences with first conservative and then non-conservative assumptions to determine the effect if any on the estimated risk. If a large impact on the estimated risk is found, the problem would be to delineate which assumptions are more "realistic In addition, an in-depth risk analysis may present literally hundreds of opportunities for assumptions and judgements on the part of the risk analyst. It may not be practical to always employ the "conservative/nonconservative" bounding technique.

Similar requirements for judgement and assumptions exist in attempting to obtain bounds for the frequencies of accident sequences. Judgement must be employed in establishing realistic bounds on individual components of the analysis, such as on failure rates, common cause or common mode coupling coefficients, etc. If distributions are to be assigned to the individual components, the choice of distributions is usually accomplished through assumptions. The size of the tails of the assigned distributions could affect the

bounds on the frequencies of accident sequences and, therefore, affect the decisions resulting from using decision criteria. In most cases, insufficient data are available to suggest a choice of distribtuions for individual components of a risk analysis, so the choice of distributions is left to the risk analyst. Different analysts may choose different distributions, and this will affect the use of any set of decision criteria that employs the estimate of uncertainty of the frequency of accident sequences as part of the criteria.

Another way of calculating bounds on the estimated frequencies of hypothesized accident sequences is through a sensitivity analysis. This precludes the necessity of judgemental choices of distributions, but judgement must still be exercised in the choice of the sensitivity bounds for each element of the analysis. If the sensitivity analysis is to be used to suggest bounds on the frequencies of accident sequences, the choice of bounds for each element of the analysis will, of course, determine the bounds on the accident frequencies.

The above examples suggest that acceptable risk criteria should not be implemented simply by subjecting risk analysis results to the criteria to determine whether or not the criteria are met. The decision maker must be in a position not only to apply the criteria, but also to judge the analysis. This is true even if a bounding analysis has been conducted and the risk criteria used account for both the frequency of consequences and the uncertainty in those frequencies. The requirement for judgement of the analyses and the ensuing conflicts between institutional interests that may arise have been recognized in the industry and several implementing techniques have been suggested to mitigate these expected conflicts. These techniques include, but are not limited to:

- Establishment of a "Science Court' to resolve institutional conflicts
- Certification of risk analysts
- Certification of risk analyses studies

The reader should not interpret the following discussion of the above three illustrative examples as endorsement by the authors. The intended purpose of the remaining sections of this chapter is to present the positive

as well as negative aspects of each of the selected examples. In addition, the authors do not suggest that the science court, analyst certification or study certifications are the only potential solutions to the decision makers. At present, there are various methods used in the deterministic licensing review. However, a detailed review of all techniques is outside the scope of this report. The three examples were selected since they are presently not in general use when reviewing engineering calculations in licensing nuclear power plants for operation, but have recently been referred to by name numerous times.

Before discussing these techniques of implementation, we should outline the standards and data required before any system or technique for implementation can be attempted.

12.3 REQUIREMENTS FOR ANY MANAGEMENT SCHEME

A common requirement for any technique or system used for the implementation of criteria is agreement or concensus by the professionals involved on a set of guidelines or methodology to be used in performing risk analyses.

The different organizational techniques each require an accepted risk methodology, but the depth to which this methodology must be defined is different for the three techniques mentioned above. To certify analysts, acceptance on the details of the methodology and application would be required, since it would be necessary to test an applicant's in-depth knowledge and skills in the area. Both the Science Court and certification team for risk analyses would require less detail of definition of the methodology, since these groups would be charged with employing scientific judgement concerning the adequacy of specific analyses.

The question arises: what is to be gained by simply publishing comprehensive guidelines for acceptable methodology and data for the conduct of risk analyses. Presumably, these guidelines could be published by the NRC or a neutral body such as one of the technical associations or journals. This would require a decision by the NRC to accept these guidelines as is done presently with professional codes and standards, and judge acceptance of the

risk analyses performed by the utilities on conformance to the guidelines. Updates to the guidelines could be made as new methodology and data became available.

The positive effect of uniformity of analyses would be facilitated by such an organizational procedure. The guidelines may help in resolving conflicts about analyses, but no mechanism is specified for the final resolution of conflicts. Following the guidelines would presumably require the risk analyst to consider all of the portions of the analysis defined as the risk methodology. The guidelines would probably do nothing to mitigate institutional or personal biases that can impact a risk analysis at many different places and in many different ways, since these involve assumptions and judgements that are very difficult to treat a priori in a set of guidelines. The requirement for assumptions and judgements appear on a case-by-case basis, and can be expected to be different for different risk analyses, hence the need for one or more of the techniques described in Section 12 below.

In addition to these guidelines, a standard Data Set would be desirable, for use when actuarial or test data is not available for such things as component failure rates or human error rates. A system which encourages the continual update of this data set is also required. Establishing this data set which again represents the concensus of professionals in the field of risk assessments would help to avoid arguments in the later stages of the review process for a risk analysis.

12.4 PROPOSED MANAGEMENT TECHNIQUES FOR THE IMPLEMENTATION OF RISK CRITERIA

If after designing a piece of hardware using existing standards and specifications, if there is still a question as to whether it can perform as desired, or that the design safety factors are sufficient, a test can usually be devised to settle the question. However, in the case of a risk analysis, questions of adequacy cannot be easily demonstrated by testing. For this reason, it has been suggested that additional means beyond accepted standards and specifications (The guidelines mentioned in Section 12.3 above) may be required for the implementation of risk criteria. These are discussed in the following paragraphs.

12.4.1 Science Court

Establishment of a Science Court that would act as a dispute resolving body for scientific issues, either as a simple arbitration panel or as a formal court, has been suggested as a means of resolving possible conflicts resulting from implementation of risk criteria. The Science Court would presumably be composed of a panel of leaders in the scientific field, who either have no conflicting institutional affiliations, or represent the various viewpoints in the conflict. The function of the Science Court would be to judge the accuracy and completeness of a disputed risk analysis, and the validity of application of risk criteria. Thus in contrast to the NRC, whose primary concern must be the public safety, the Science Court would be charged with balancing both the risks and benefits from decisions affecting nuclear power or elements of that technology. The scope of the Science Court is visualized as spanning the perspectives of the NRC and of the nuclear industry.

It may be speculated that the proposed Science Court would operate on a technological level somewhat like the U.S. Supreme Court operating on a legal level, and therefore it is tempting to draw comparisons between the two. One such comparison involves the basis by which conflicts are to be resolved. The U.S. Supreme Court basically uses the U.S. Constitution the codified law of the land and previous judgements as a basis of reference for resolving conflicts. An analogous basis of reference for the Science Court could be a set of rules or guidelines for the conduct of an "acceptable" risk analysis. These guidelines might take the form of "standard practices" that define methodology and data base utilization, fault types that must be included, some measures of completeness or comprehensiveness of the analysis; in other words, a standardized, codified reference methodology.

Another point of comparison may be made regarding the adversary process and burden of proof. The three major branches of U.S. jurisprudence are criminal law, civil law, and equity. Under criminal law as practiced in the U.S., the burden of proof is on the state, who is the claimant in the case (innocent until proven guilty). Under civil and equity law, the burden of proof is generally on the plaintiff, or claimant, but may shift back and forth between plantiff and defendant during the course of the trial. The rule appears to be that the burden of proof is on the party with the greatest access to pertinent documents and witnesses.

In the scientific community the burden of proof has traditionally been on the party proposing a new theory, or analysis, or methodology. In fact, the scientific community takes great comfort in this arrangement, claiming that the degree of scrutiny by peer review and repeatable experimental verification will eventually lay bare unfounded claims for a theory or analysis. However, it is difficult to see how a risk analysis can be experimentally verified. Therefore, we are in a position of substituting the judgement of the Science Court for the judgement of the analysts who performed the risk analysis, and the decision makers who applied a risk criteria. Who should have the burden of proof in this case? The utilities will generally have better access to the types of data necessary to perform the risk analysis, but the NRC, charged with the public safety, should have a better hold on what risk is acceptable to the public; both items would ostensibly be under judgement by the Science court.

It is necessary to speculate on ways that the proposed Science Court may be organized and the rules by which it might operate, since a variety of options are possible. One organizational scheme might see a Science Court composed entirely of impartial leaders in the field of risk analysis or related disciplines, presuming impartial leaders can be found. Are the members of the Science Court to be elected or appointed? By whom? Certainly the answers to these questions will to some extent determine the degree of impartiality of the court members, since even disinterested professionals in the area will have some predisposition toward one viewpoint or another.

Another option might see a Science Court composed of equal numbers of members from the contending sides (and there may be three sides here: the utilities, the NRC and the intervenors) with an impartial tie-breaker. Such a court might have a difficult time reaching a true concensus since the viewpoints of the competing institutional interests are built into the court. This organizational structure may result in the tie-breaker making a majority of the decisions de-facto. However, it is pointed out that the EPA employs such a "task force" system for resolving disputes between itself and local governing bodies, apparently with some success.

Establishment of a Science Court may accomplish the positive result of putting the risks from nuclear power plants into perspective with the benefits obtained from nuclear power, if only in a heuristic, judgemental way. Further, it is a technique for resolving conflicts among differing institutional

interests. It is also a way of gaining acceptance of some set of risk criteria, since the judgements of the court would essentially serve to define estimated nuclear risks that are acceptable and unacceptable. However, it is difficult to visualize how such a court would operate without an agreed-upon body of methodological and data techniques. Are the court members to pass on each assumption and judgement in a risk analysis? Certainly they must be in a position to agree or disagree with the assumptions that are required to evaluate the dominant accident sequences. But conservative and non-conservative assumptions might influence which accident sequences are dominant in any given analysis. The work load on the members of such a court may be awesome if each risk analysis that comes before the court must be reviewed as to technical acceptance. The time required to resolve and rule on each issue could also be unacceptable. Finally, the question of the legality of decisions handed down by the court must be raised. This will depend on the charter for the court. Presently NRC is legally responsible for regulating the nuclear industry. Is the court to be an arm of NRC, or will it have a separate identity? Will the court review all energy options or just the nuclear part? Will there be an appeal process through other legal mechanisms on decisions handed down by the court? These questions involving the legal relationships among the court and other institutions must be addressed and options formulated.

12.4.2 Certification of Risk Analysts

A procedure for certifying risk analysts in much the same way that professional engineers are certified, has been suggested as a way of assuring a minumum level of competence for risk analysts. The hope is that an improvement in the quality of risk analyses would result. Presumably no risk analysis would be accepted as an authoritative statement of the risk associated with a nuclear power plant unless the analysis was done or certified by certified risk analysts.

This organizational technique would require that a certification procedure be established. Presumably the certification procedure would require thorough testing of the analyst being certified. A demonstrated knowledge on the part of the analyst of accepted risk analysis techniques would be required before certification could occur including: generation of decision trees;

success criteria evaluation; quantification of decision trees including data base development and utilization, hardware and human error quantification; common mode and common cause analysis, and test and maintenance contributions; Boolean algebra manipulations required for event tree evaluation and systems interactions; source term estimation; and consequence estimation. It may be unrealistic to expect a single individual to be certified in each of these areas. A large number of disciplines are required to perform a complete risk analysis. However, it would be desirable for a certified risk analyst to be at least familiar with all aspects of a risk analysis. It may be more realistic to limit certification to one of several areas, as with the present professional engineer system. This group of certified engineers could then review and "approve" analyses of many others. One possible breakdown, roughly on the basis of discipline and tradition, would be probabilistic safety analysis including decision tree generation and quantification and data base analysis, source term estimation (including isotopes, core melt environments, transport, plate-out, etc.), and consequence modeling including plume modeling, evacuation, fallout, health effects, pathways, immediate and long term deaths, etc. It is interesting to note that the first of these categories (decision trees and data base) is required for analyses upon which the core melt risk criteria would be applied, while the others are not. Thus, certification of analysts in the decision tree and data base category would impact all of the decisions to be made on the basis of acceptable risk criteria. But it is precisely in the area of decision trees and data base that many of the major areas of analysis uncertainty, questions of completeness of analysis, and disagreements on methodology and data are expected to lie. This presents both opportunities and problems for the risk analyst certification process. The ability to certify risk analysts implies an ability to concur on a set of techniques upon which certification will be based. And the concurrence must be in some depth since meaningful tests are to be administered. Since the methodology is still evolving, it may be difficult to reach concurrence on a set of standard risk analysis practices and who would be involved in the concurrence. Also, since the methodology is evolving, a major problem of updating certification of analysis is evident. This may take the form of recertification tests at regular intervals as new techniques or data become available.

Certification of risk analysts, particularly in the area of decision tree analysis and quantification, would accomplish some objectives with respect to managing the implementation of risk criteria:

- Analyst education: The risk analyst would be required to learn the body of techniques comprising the risk analysis methodology in order to be certified. The certification and testing process would provide the structure and discipline necessary for this process (presumably). This, however, does not imply that the certification procedure provides the best format and initiative for training risk analysts.

- Standardized risk analysis techniques: Since the certification procedure will require tests on the application of risk analysis methodology and data, will have to be reached on a set of standardized risk analysis techniques. This is not to imply, however, that the standardized techniques will or should remain stagnant. They should not. As new techniques and data become available, and their usefulness demonstrated they should be included in the body of standardized techniques. The maintenance of this updating system could prove quite difficult.

- Uniformity of analyses: The implementation of risk criteria would be greatly facilitated if it could be assured that all risk analyses to be subjected to the criteria were uniformly performed with respect to methodology and data used, and assumptions made. Certified analysts could be expected to be more uniform in the application of the methodology and data. However, this does not imply that differences in analyses due perhaps to the availability of different data or to implicit personal or institutional biases can be eliminated. It is simply a step in the direction of greater uniformity in risk analyses. It also does not assure that the analyses will be accurate or realistic.

- Resolution of conflicts: Since the potential economic impact of implementation of risk criteria is so large, conflicts between the NRC and the nuclear industry are probably inevitable. Certification of risk analysis personnel may serve to diffuse some of these conflicts. The fact that both sides will be arguing from a common basis of "accepted" techniques and data should at least focus the discussion to detailed paticulars (e.g. assumptions), and away from arguments about whether or not the methodology is correct or complete.

Not all of the implications of establishing a certification procedure for risk analysts are positive. Implementation of this organizational technique may also lead to the following:

- The process of certifying risk analysts could stifle initiatives for developing and implementing new methodology and data. If the viewpoint is taken that new methodology or data will simply contaminate a functioning system, (i.e as in the codified area a reluctance to

change code practices) the introduction of new techniques may be very difficult. That is, the introduction of new techniques may be delayed if there is a reluctance to accept the new techniques (on the basis that they are "non-standard") on the part of the certified analysts.

- A risk analyst certification process may provide disincentives for qualified personnel attempting to enter the field. At the present, there appears to be a shortage of personnel capable of contributing to risk analyses. Disincentives would of course worsen this situation. The impact of certification on this shortage would depend on how the certification procedure was intended to operate. If the only function of the certified risk analyst was to review and stamp approval of risk analyses, then fewer certified analysts would be required than if the entire analysis were to be done by certified analysts. However, using the certified analyst as a reviewer only would likely result in fewer accrued benefits from the certification process than if the certified analyst was actively engaged in the analysis.

- The certification process would have to handle a geat deal of applicants in a short time interval so that the industry would have the ability to respond to a risk criteria in a timely manner. This would imply a large, efficient organization would be required.

To summarize, the positive aspects that could be expected to result from establishing a risk analyst certification procedure are in terms of: codifying a set of standard practices for risk analysis; education of risk analysts in this standard methodology; and increased uniformity of analyses and the opportunity to resolve conflicts more easily than will result from the first two benefits. On the other hand, the process of certification could delay innovation and stifle further development of the risk analysis field, and may provide barriers to qualified analysts who would otherwise enter the field. The positive impacts would facilitate the implementation of acceptable risk criteria. The negative impacts may be detrimental to the risk analyses field as a whole.

12.4.3 Certification of Risk Analysis Studies

Certification of risk analysis studies has been suggested as an alternative organizational technique to certification of analysts for improving the quality and uniformity of risk analyses. The purpose of a certification team would be to review risk analyses to assure a minumum level of study competence and conformance to accepted risk analysis principles and techniques. In contrast to the proposed Science Court, the purview of the certification team

would be only on the risk analysis study; not on the validity of the application of acceptable risk criteria or on the tradeoffs between risks and benefits from the implications of applying acceptable risk criteria. Presumably the certification team could be composed of a panel of impartial experts, or by a panel consisting of representatives from each institutional interest, e.g. NRC, the industry and intervener groups.

Certification of risk analyses would potentially accomplish many of the same positive objectives that were enumerated for certification of risk analysts (except analyst education) without some of the negative aspects noted there. The study certification process would not be expected to delay implementation of new techniques or data, nor provide disincentives for analysts entering the risk analysis field. It would still be required to obtain concurrence on a standardized body of acceptable risk analysis practices. However, it is expected that the details of the accepted methodology could be more loosely defined, since each study is to be reviewed on its own merits.

Questions still remain as to the implementation, operation and authority of a certification team. How would the certification team interact with the various institutional interest? What authority does the certifying team have to influence the risk analysis? Could the certifying team refuse to certify an objective risk analysis simply because of institutional biases or because the study does not completely conform to the accepted methodology?

The authority of a certification team may be anywhere on a spectrum ranging from peer review, with no authority to dictate changes, through absolute authority to not certify a study until it is totally acceptable in all aspects to the certification team. To some extent this will be left up to the certification team style of operation. The process of giving a certification team the power to certify or to not certify implicitly grants that team a wide range of options for exercising authority. It is pointed out that the first case, peer review, is similar to the peer review presently practiced by NRC. The second case essentially substitutes the certification team as a super analysis-management function. A case in between may be where the certification team refers the problem to a Science Court if unreconcilable differences exist between the certification team and analysis team.

Rules of operation for the certification team will also require resolution. Do the rules require unanimous concensus or a simple majority vote? If the latter, the make-up of the team in terms of institutional interests becomes very important. A tie-breaking mechanism would be required. If the former, the various institutional interests represented on the certification team may result in many no-decisions or "hung juries."

To summarize, institution of a risk analysis certification team would accomplish many of the objectives of the Science Court and certification of risk analysis personnel organizational techniques, but does not completely supplant either. The authority and rules of operation of a certification team would have to be clearly defined to avoid instituting an organizational technique that was not intended.

12.5 SUMMARY

The promulgation of risk criteria will require some additional regulatory techniques for proper implementation. The standards and specifications or guidelines alluded to in Section 12.3 are more readily accepted by the public and industry if they are first established by concensus of the professionals in the fields they cover. While it is possible for regulatory agencies to establish these guidelines, it is probably more acceptable if a regulatory agency were to encourage their creation by the professional societies.

Once these guidelines have been promulgated, and accepted by the regulatory agencies involved, the additional techniques discussed in Section 12.4 will become useful in the required review process. In this area, the regulatory agencies must take the initiative and choose the one or more methods required. Of the three techniques discussed in Section 12.4, the Science Court is probably the most direct and quickest to become operational. Certainly, if risk criteria were to be put into effect within one or two years, the others, requiring certification of individuals or groups would not be effective in time.

Chapter 12 References

1. Reactor Safety Study (RSS), "An Assessment of Accident Risks in U.S. Commerical Nuclear Power Plants," WASH-1400, NUREG 75/014, 1975.

2. Swain, A.D., Guttman, H.E., "Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications," NUREG/CR-1278, April 1980.

APPENDIX A

ASSESSMENT OF COMMERICAL NUCLEAR POWER EXPERIENCE AND PROJECTION UP TO THE YEAR 2000

The number of light water reactor-years of operating experience that have been accumulated from 1960 to the end of the year 1979, and the number of light water reactor-years that are projected to accumulate from the beginning of 1980 up to the year 2000 are estimated in this appendix.

The estimated cumulative reactor-years of experience were computed based on the starting dates of commercial operation of all light water reactors (LWRs) obtained from Reference (1). A total of 71[1] reactor units had started commercial operation before the end of the year 1979. Of these, 68[1] were LWRs. However, extended shutdowns of three reactors (Indian Point 1, Three Mile Island 2 and Humbolt Bay 3) were still in effect at the end of the year 1979. Therefore, 65 out of a total of 68 LWR reactor units were considered operational at the end of 1979. The estimate of cumulative reactor-years of experience took into account the extended shutdown periods for the above-mentioned three reactors, but did not consider the periodic shutdowns for maintenance and/or refueling for the other 65 reactor units. The estimated cumulative reactor-years of experience from 1960 to the end of year 1979 are shown in Figure A.1. From this figure, it is seen that about 440 reactor-years of LWR experience were accumulated at the end of the year 1979. The contributions by PWRs and BWRs to the total experience of 440 reactor-years were 246 and 194 reactor-years respectively. The manner in which the projected number of reactor-years will accumulate in the future is described next.

Reactor-years projected to accumulate from 1979 up to the year 2000 were based on Projection C forecasts of nuclear power prepared by the Energy Information Administration (EIA) of the United States Department of Energy[2]. Projection C forecasts correspond to an energy future characterized as medium demand and medium supply. There are three estimates of Projection C forecasts, low, medium and high. These three estimates correspond to the current outlook on the variability of oil prices. A summary of Projection C forecasts is shown in Table A.1.

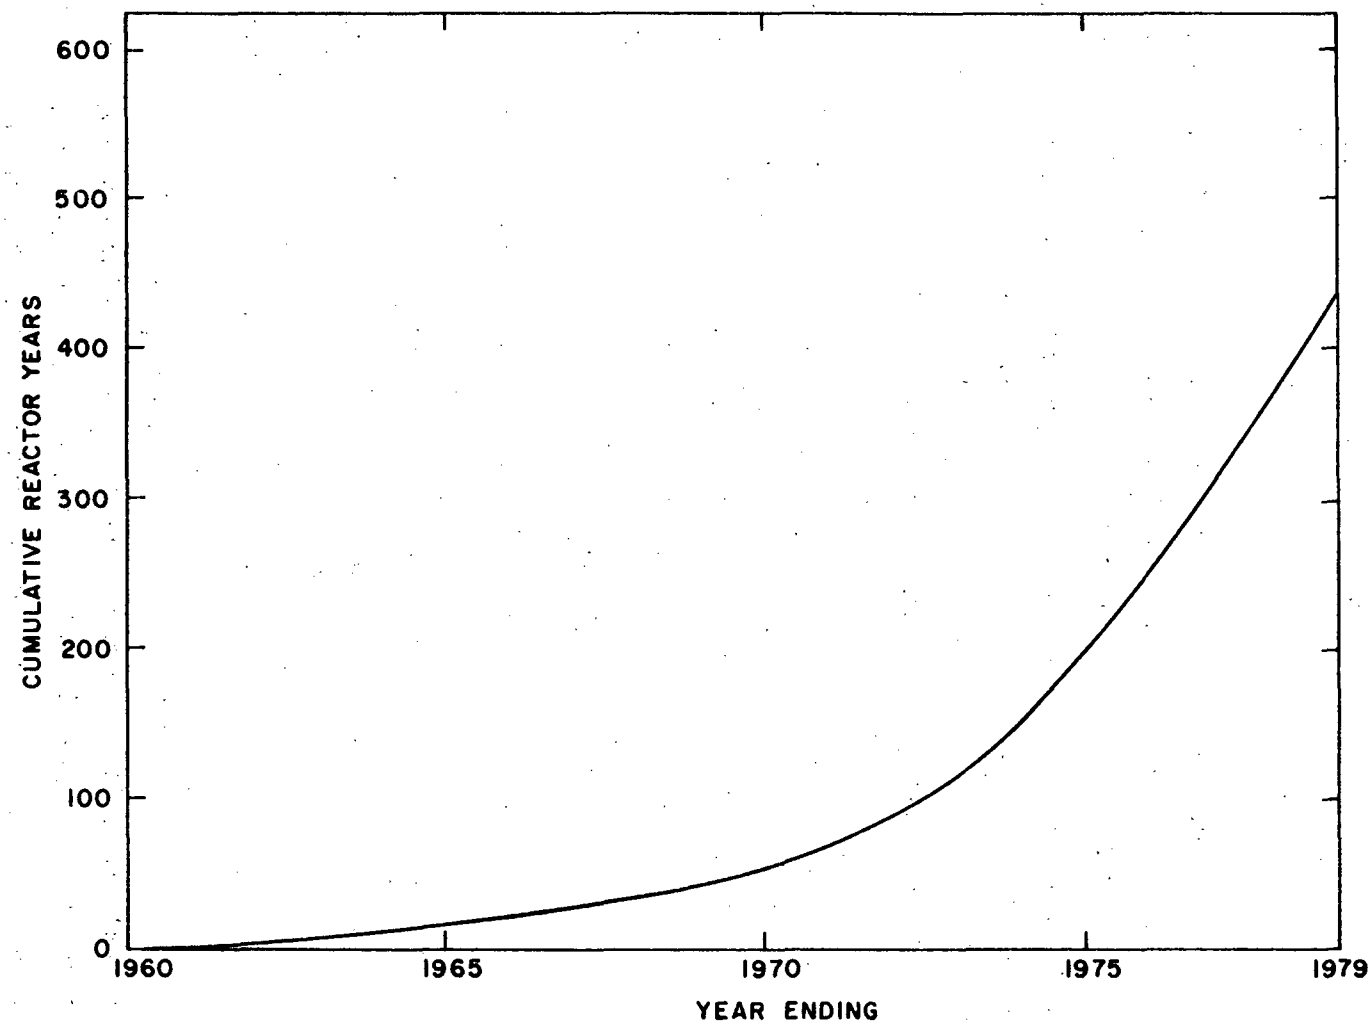

Fig. A.1 Estimated light water reactor-years of experience up to the end of the year 1979.

TABLE A.1

Updated Domestic Nuclear Power Forecasts[2] (GWe)

Year Milestone	Updated Forecast October 1979 Low	Medium	High
1985	95	106	113
1990	129	140	155
1995	156	179	196

Growth rates of nuclear power corresponding to the above forecasts were used to estimate the projected installed nuclear capacity from 1980 to 1995. The July 1979 forecasts of EIA[3] were used to infer growth rates from 1995 to the year 2000 which were then used to assess the projected installed capacity in the same period. These projections are shown in Figure A.2. The number of projected reactor-years that is forecasted to accumulate in the period 1979 to 2000 was obtained by subtracting the installed nuclear capacity at the end 1979 (51,166 MWe) from the projected installed capacity. In addition, an assumption was made that all future reactors would be LWRs with an installed capacity of 1GWe. The esimates of reactor-years projected to accumulate in the period 1980 to 2000 are plotted in Figure 6.2 (see Chapter 6). A summary of these estimates is presented in Table A.2.

TABLE A.2

Projected Cumulative Light Water Reactor-Years of Experience

Beginning of Year	Forecast of Cumulative Reactor-Years Low	Medium	High
1985	856	879	894
1990	1451	1528	1590
1995	2205	2356	2495
2000	3117	3392	3645

The above table considers some growth in nuclear power. Hypothetically, if one assumes that there is a moratorium on the operation of any new reactors from 1979 onwards, then as stated earlier, 65 light water reactor-years would be added every year in the future to the 440 reactor-years of experience

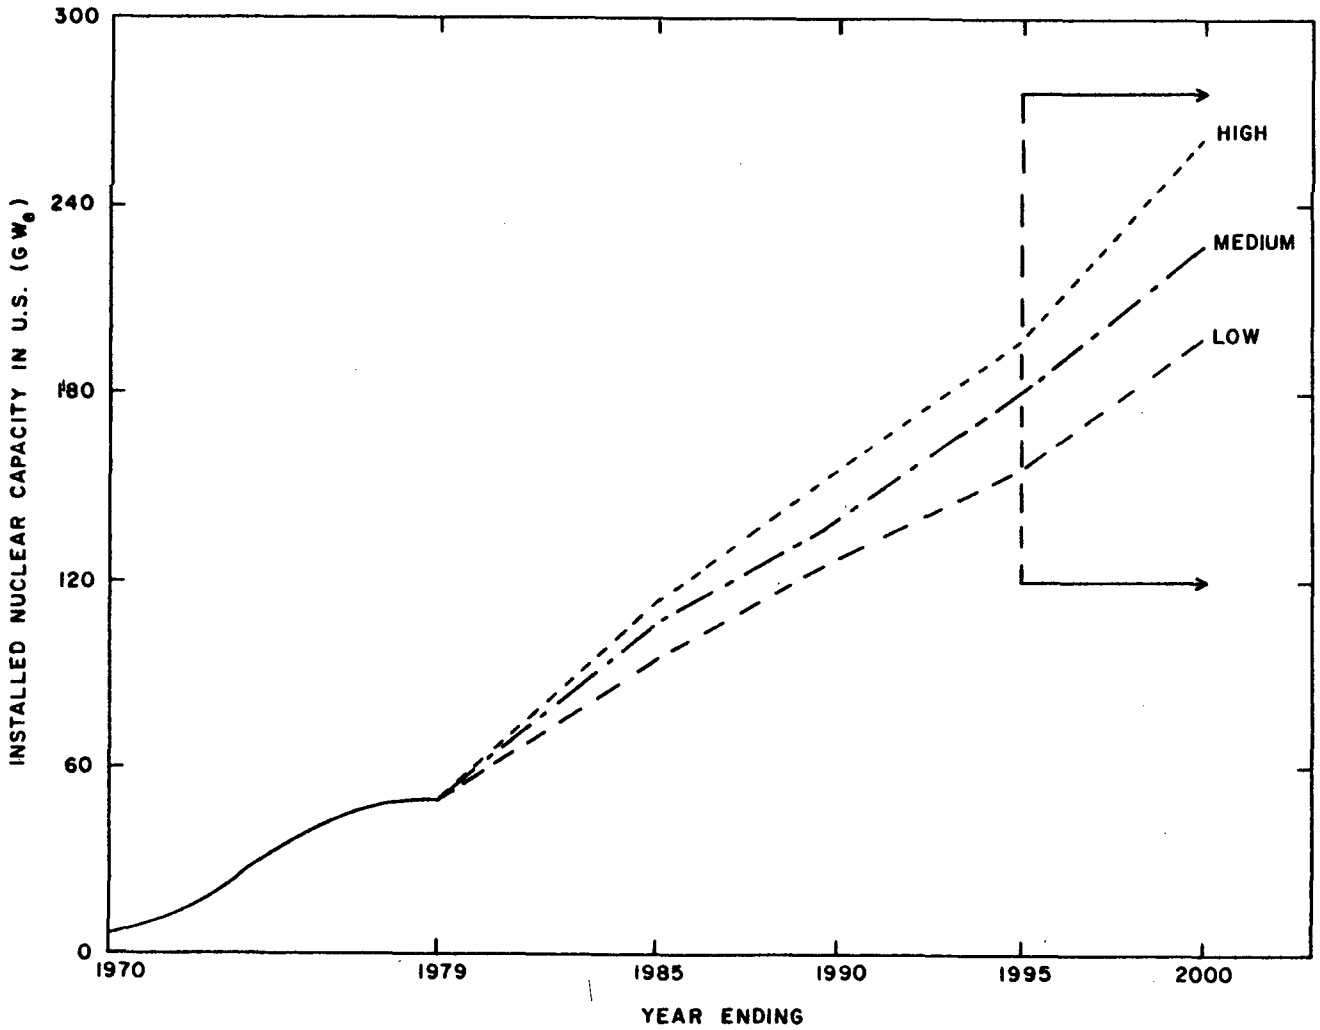

Fig. A.2 Installed nuclear capacity up to the year 1979 and projected nuclear capacity from 1980 to the year 2000.

accumulated up to the year 1979. If it is assumed that these 65 LWRs would continue to operate in the future, then in the next 10 and 20 years, the projected cumulative reactor-years may be estimated as 1090 and 1740 respectively.

Appendix A References

1. "World List of Nuclear Power Plants, Operable, Under Construction, or on Order (30 MWe and Over) as of December 31, 1979," Nuclear News, February 1980.

2. Clark, R.E., "Commercial Nuclear and Uranium Market Forecasts for the United States and the World Outside Communist Areas," U.S. Department of Energy, Energy Information Administration (EIA), DOE/EIA-0184/24, January 1980.

3. "Annual Report to Congress, 1978, Volume Three: Forecasts," U.S. Department of Energy, Energy Information Administration, DOE/EIA-0173/3, July 1979.

APPENDIX B

VARIABILITY IN THE WEIGHTED SOCIETAL RISK DUE TO DESIGN AND SITE DIFFERENCES

An attempt is made in this appendix to show the variability in the weighted societal risk due to design and site differences. Table B.1 shows the variation in the weighted societal risks of early fatalities, latent fatalities, and property damage (both in-plant and out-of-plant) by hypothesizing that the same plant is located at different sites or that different plants are located at the same site. The first consideration shows the variation in the weighted societal risks of different types of consequences as they relate to different sites, while the second shows the variation in the same measures as they relate to different designs.

TABLE B.1

Variation in the Weighted Societal Risks Due to Site and Design Differences

	Weighted Societal Risk ($/plant year)			
	Early Fatalities	Latent Fatalities	Property Damage In-Plant	Property Damage Out-of-Plant(c)
Same Plant Located at Different Sites (a)				
SITE:				
Diablo Canyon	16	1620	6×10^4	1290
Palisades	290	2430		2670
Fermi	920	3240	Same	4780
Limerick	3500	4230	as	6980
Zion	4700	3870	above	6030
Indian Point	6100	4860		9550
Different Plants Located at the Same Site(b)				
PLANT:				
Indian Point	630	396	3×10^4	700
Sequoyah Ice Condenser	2700	10800	4×10^4_4	14800
Peach Bottom	17000	9900	3×10^4	13500

(a) Surry Plant
(b) Indian Point Site
(c) In terms of 1974 dollars

The weighted risks as shown in Table B.1 were obtained by weighting the societal risks as assessed in Reference 1 with the same weighting factors that are shown in Table 10.1. The weighted risk of latent fatalities was obtained by following the same procedure that was used in Reference 2 for calculating the results of Table 10.1 (i.e. the societal risk of latent fatalities per year is multiplied by an assumed plateau period of 30 years). The plant property damage risk was obtained by multiplying the frequency of core melt accidents[1] by its assumed cost ($1 x 10^9).

From Table B.1 it may be observed that the risks of plant property damage in all cases are higher than the weighted risks of other types of consequences. On the other hand, the sum of the weighted risks of early and latent fatalities in all cases are comparable to the out-of-plant property damage risks.

Appendix B References

1. Bernero, R.M., Blond, R.M., Pritchard, W.C., Taylor, M.A., Eysmontt, G., and Sege, G., "Task Force Report on Interim Operation of Indian Point, USNRC, NUREG-0715, August 1980. *

2. Clausen, M.J., Fraley, D.W., and Denning, R.S., "Improved Methods for Incorporating Risk in Decision Making," (Interim Report), PNL 3523, August 1980.

*Available for purchase from the NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission, Washington, DC 20555, and/or the National Technical Information Service, Springfield, VA 22161.

ACKNOWLEDGEMENTS

The authors wish to express their thanks to Dr. W. Vesely, Jr. and Mr. S.R. Sturges, of the Nuclear Regulatory Commission, for their technical guidance throughout the research. Their comments, along with Mr. Niyogi's, also of the NRC, on the two drafts of this document, have helped significantly in its final formulation.

The authors also express their appreciation to Mr. E. Lofgren, of SAI, for his help in Chapters 5 and 12, Mr. J. Fragola for his technical input to Chapter 11, and Dr. A. Tingle, of our staff, for his advice in the area of meteorology and the CRAC code. They also are indebted to Ms. S. Dyroff and Ms. S. Monteleone for the final preparation and typing of the manuscript.

NRC FORM 335 (7-77)

U.S. NUCLEAR REGULATORY COMMISSION

BIBLIOGRAPHIC DATA SHEET

1. REPORT NUMBER *(Assigned by DDC)*
NUREG/CR-2040
BNL-NUREG-51367

4. TITLE AND SUBTITLE *(Add Volume No., if appropriate)*
A Study of the Implications of Applying Quantitative Risk Criteria in the Licensing of Nuclear Power Plants in the United States

2. *(Leave blank)*

3. RECIPIENT'S ACCESSION NO.

7. AUTHOR(S)
S. Mitra, R. Hall, A. Coppola

5. DATE REPORT COMPLETED
MONTH: March | YEAR: 1981

9. PERFORMING ORGANIZATION NAME AND MAILING ADDRESS *(Include Zip Code)*
Department of Nuclear Energy
Brookhaven National Laboratory
Upton, New York 11973

DATE REPORT ISSUED
MONTH: May | YEAR: 1981

6. *(Leave blank)*

8. *(Leave blank)*

12. SPONSORING ORGANIZATION NAME AND MAILING ADDRESS *(Include Zip Code)*
Division of Systems and Reliability Research
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555

10. PROJECT/TASK/WORK UNIT NO.

11. CONTRACT NO.
FIN A3222

13. TYPE OF REPORT
Technical

PERIOD COVERED *(Inclusive dates)*
10/1/80 - 4/30/81

15. SUPPLEMENTARY NOTES

14. *(Leave blank)*

16. ABSTRACT *(200 words or less)*

Describes the results of an investigation into the feasibility of developing and using a set of probabilistic risk criteria to help judge the safety of nuclear power plants. The principal aim of this report was to critically review and examine the implications and ramifications of the various proposals for a numerical risk criterion from a unified viewpoint.

17. KEY WORDS AND DOCUMENT ANALYSIS

17a. DESCRIPTORS

Quantitative Risk
Risk Criteria
Component Availability
System Availability
Release Criteria
Numerical Risk

17b. IDENTIFIERS/OPEN-ENDED TERMS

18. AVAILABILITY STATEMENT
Unlimited

19. SECURITY CLASS *(This report)*
Unclassified

21. NO. OF PAGES

20. SECURITY CLASS *(This page)*
Unclassified

22. PRICE
S

NRC FORM 335 (7-77)

UNITED STATES
NUCLEAR REGULATORY COMMISSION
WASHINGTON, D. C. 20555