

**Toward a safety goal : dicussion of preliminary policy considerations
/ Office of Policy Evaluation, U.S. Nuclear Regulatory Commission.**

Washington, D.C. : The Commission, 1981.

<http://hdl.handle.net/2027/mdp.39015041241111>

HathiTrust



www.hathitrust.org

Public Domain, Google-digitized

http://www.hathitrust.org/access_use#pd-google

This work is in the Public Domain, meaning that it is not subject to copyright. Users are free to copy, use, and redistribute the work in part or in whole. It is possible that heirs or the estate of the authors of individual portions of the work, such as illustrations, assert copyrights over these portions. Depending on the nature of subsequent use that is made, additional rights may need to be obtained independently of anything we can address. The digital images and OCR of this work were produced by Google, Inc. (indicated by a watermark on each page in the PageTurner). Google requests that the images and OCR not be re-hosted, redistributed or used commercially. The images are provided for educational, scholarly, non-commercial purposes.

UMMU
TK
9152
.T69
1981

C 581,268

Digitized by Google

Original from
UNIVERSITY OF MICHIGAN

Y3.N88:10/0764

1057-H-2

NUREG-0764
For Comment

Toward a Safety Goal: Discussion of Preliminary Policy Considerations

**U.S. Nuclear Regulatory
Commission**

Office of Policy Evaluation



NORTH ENGINEERING LIBRARY
1001 S.T. BLDG.
THE UNIVERSITY OF MICHIGAN
ANN ARBOR, MICHIGAN 48109

Digitized by Google

Original from
UNIVERSITY OF MICHIGAN



Available from

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

and

National Technical Information Service
Springfield, Virginia 22161

Toward A Safety Goal: Discussion of Preliminary Policy Considerations

Manuscript Completed: March 1981
Date Published: March 1981

Office of Policy Evaluation
U.S. Nuclear Regulatory Commission
Washington, D.C. 20555



umma

TK

152

169

171

UMMU
DEPOS.

10/ '96

ABSTRACT

This report includes a Commission statement on preliminary policy considerations in development of an NRC safety goal and a supporting report providing a more detailed discussion. The discussion begins with an exposition of criteria by which the quality of different approaches to a safety goal may be judged. It next addresses alternative approaches to safety-goal formulation and presents some safety-goal proposals. Finally, there is a brief discussion of degree of safety.

CONTENTS

	<u>Page</u>
COMMISSION STATEMENT:	
DEVELOPMENT OF A SAFETY GOAL: PRELIMINARY POLICY CONSIDERATIONS.....	vii
SUPPORTING REPORT:	
TOWARD A SAFETY GOAL: DISCUSSION OF PRELIMINARY POLICY CONSIDERATIONS	
I. INTRODUCTION.....	1
A. Background.....	1
B. Purpose and Scope of this Report.....	1
C. Organization of the Report.....	2
II. CRITERIA FOR ASSESSING AN APPROACH TO A SAFETY GOAL.....	2
A. Considerations in Formulating Criteria.....	2
B. Proposed Criteria: Formulation 1.....	6
C. Proposed Criteria: Formulation 2.....	12
D. General Observations.....	16
III. APPROACHES TO SAFETY GOAL FORMULATION.....	18
A. Methods of Approach to Acceptable Risk.....	18
B. Possible Characteristics of Safety Goals.....	21
C. Types of Quantitative Approach.....	23
D. Approaches to Dealing with Uncertainty.....	26
E. Considerations in Balancing of Values.....	28
IV. SOME SAFETY GOAL PROPOSALS.....	29
A. The ACRS Proposal.....	31
B. Other Proposals.....	42
V. DEGREE OF SAFETY.....	45

Commission Statement, as published in the Federal Register:

DEVELOPMENT OF A SAFETY GOAL -
PRELIMINARY POLICY CONSIDERATIONS

In accordance with its previously announced Plan for Developing a Safety Goal (45 FR 71023, October 27, 1980), the Nuclear Regulatory Commission hereby publishes for public comment a preliminary statement of policy considerations concerning that subject. A detailed discussion supporting this statement is being published separately as NUREG-0764, "Toward a Safety Goal: Discussion of Preliminary Policy Considerations." A copy of NUREG-0764 and of NUREG-0739, "An Approach to Quantitative Safety Goals for Nuclear Power Plants," submitted to the Commission by the Advisory Committee on Reactor Safeguards in October 1980 are available for inspection at the Commission's Public Document Room, 1717 H Street, N.W., Washington, D.C. Single copies of NUREG-0764 and NUREG-0739 are available upon written request and at no cost to persons who wish to comment; requests should be made to the NRC-GPO Sales Program, Attention: Sales Manager, Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, D.C. 20555 (Phone 301-492-9530). Copies may also be purchased from the NRC-GPO Program and the National Technical Information Service, Springfield, Virginia 22161. Comments on the statement and on NUREG-0764 and on NUREG-0739 will be considered in work leading to the preparation of a policy paper for submittal for Commission consideration. That work includes further studies, public meetings, and workshops. Salient activities will be publicly announced in order to secure broad participation. (For a description of that further work, see NUREG-0735, included in 45 FR 71023, October 27, 1980.)

Comments and suggestions are solicited concerning the entire subject matter of this statement and NUREG-0764 or any of their aspects. In addition, the Commission solicits comments on certain additional, specific questions appearing at the end of this notice.

Written comments should be addressed to the Secretary of the Commission, U.S. Nuclear Regulatory Commission, Washington, D.C. 20555, Attention: Docketing and Service Branch, and should be received by May 26, 1981. Late comments will be considered, but only to the extent practicable to do so.

For further information: Contact Mr. George Sege, Office of Policy Evaluation, U.S. Nuclear Regulatory Commission, Washington, D.C. 20555, (202) 634-3295.

Dated at Washington, District of Columbia, this 23rd day of March 1981.

For the Nuclear Regulatory Commission.

Samuel J. Chilk
Secretary of the Commission

vii

PRELIMINARY POLICY CONSIDERATIONS

I. INTRODUCTION

A. Background

In accordance with the Plan for Developing a Safety Goal (NUREG-0735; also published in 45 FR 71023, October 27, 1980), the Commission is issuing for public comment this preliminary statement of considerations involved in developing a safety goal. This preliminary statement has been prepared in furtherance of the Commission's intent to seek to define more clearly the level of protection of the public health and safety that it believes is adequate.

In accordance with further provisions of the Plan, the comments received will constitute a part of the considerations leading to preparation of the eventual safety goal. In addition to consideration of the written comments, workshops and public meetings will be held. A subsequent paper will be prepared by the Office of Policy Evaluation and is scheduled to be submitted to the Commission for its consideration in August 1981. After consideration of that paper, the Commission intends to seek additional public comment on the proposed policy statement. Those comments will also be considered in the preparation of the final safety goal itself.

B. Purpose and Scope of the Preliminary Statement

The purpose of this preliminary statement is to elicit comments by indicating the kinds of considerations which may enter into an articulation of the Agency's safety goal.

The Advisory Committee on Reactor Safeguards has prepared "An Approach to Quantitative Safety Goals for Nuclear Power Plants" (NUREG-0739, available on request). Commenters are urged to obtain and use that document as one example of a concrete application of the concepts discussed here. The Commission has at this time not formed any views as to the merits of the approach described by the ACRS or any other single policy. The ACRS has characterized its proposal as intended to "serve as one focus for discussion." It is the only specific proposal that has so far been formally submitted to the Commission, and is detailed and accompanied by background information.

The scope of early policy development efforts has focused primarily on radiological accidents associated with power reactors. It is, however, intended to include considerations of sabotage and other external events and to seek adaptation of the safety goal to other facilities and operations where the need is evident.

C. Organization of the Statement

The articulation of a safety goal must not only stipulate a required degree of protection that is sound, but must also be formulated in a way that is useful in a practical, regulatory sense.

This statement begins by discussing alternative approaches to safety-goal formulation. The discussion includes types of safety goals, types of quantitative approach, approaches for dealing with uncertainty, and comments about some aspects of balancing of risks and of safety and other values.

Presented next are references to five specific safety-goal proposals in addition to the aforementioned volume by the Advisory Committee on Reactor Safeguards.

Finally, there is a short discussion of degree of safety, since all or most of the potential approaches can be adapted to establishing a degree of protection at any level over a wide range of possible levels. Selecting that level on a sound basis is, of course, the paramount objective of the effort to develop a safety goal.

II. APPROACHES TO SAFETY GOAL FORMULATION

In this section, discussions progress from a characterization of possible safety-goal forms, to an identification of types of quantitative approaches, to a brief account of means of dealing with uncertainty, and, finally, to comments about balancing disparate values.

A. Safety Goal Forms

Safety-goal forms have been classified according to seven pairs of contrasting characteristics which are given below. (Based on Mattson et al., 1980. See NUREG-0764 for further discussion). The next steps in the goal development program should address these characteristics. The seven pairs of characteristics, with brief comments, are as follows:

1. Single vs. Multiple Goals - Advantageous combinations may be possible, consisting of an interrelated hierarchy of overall and lower-order goals.
2. Quantitative vs. Qualitative Goals - Overall quantitative goals introduce severe verifiability problems with respect to low-probability high-consequence events. On the other hand, qualitative goals may defy assignment of clear, ascertainable meaning. Different specific aspects of safety-goal formulation may be better suited to one or the other of these alternatives, or may be expressible in forms combining qualitative and quantitative elements.
3. Ends-Oriented vs. Process-Oriented Goals - Ordinarily, goals are thought of principally as ends to which strategies and resources are directed. A process-oriented alternative focuses on processes that command confidence. For example, one of the basic reactor safety principles of the Naval Nuclear Propulsion program and our own is that if the means of providing for nuclear safety are pursued with the goal of excellence of performance, then acceptable safety will be attained.

4. Absolute vs. Relative Goals - Some comparative analysis must underlie relative goals, though comparisons may be fraught with considerable difficulties.
5. Individual vs. Society-Oriented Goals - In the NRC regulations (10 CFR Parts 20, 50/Appendix I, and 100), both individual and societal goal forms are used in complementary fashion for limiting radiation exposure.
6. Site (or Region) Dependent vs. Site (or Region) Independent Goals - Location may affect the probability of a serious accident, the probability of a failure of a reactor containment, and the radiation exposure of people in event of an accident.
7. Time-Related vs. Atemporal Goals - A time-related goal may call for achievement of some specified safety standard by a specified year and a more stringent safety standard by some later year. Also, a goal may be interim or settled (though even a "settled" goal may need to be changed, because of technological progress, new information, or new perceptions). An interim goal may offer a trial period.

B. Types of Quantitative Approach

Recent efforts towards development of proposed safety goals have had a predominantly quantitative emphasis, notwithstanding evident problems of verification of compliance with any quantitative goals. The Advisory Committee on Reactor Safeguards (ACRS), as background for its own development of a quantitative approach, accepted a survey of some proposals previously presented in the United States and abroad, according to which such approaches can be roughly categorized into three groups: those that set limits on individual risk of death only; those that consider frequency of accidents and magnitude of the consequences; and those that imbed the criteria in risk management frameworks that, at least in part, consider risks from alternatives or other societal endeavors (NUREG-0739, Part 1).

C. Approaches to Dealing with Uncertainty

Verifiability of whether a nuclear plant meets an overall quantitative safety goal is subject to considerable uncertainties surrounding probabilities and consequences of severe but low-probability accidents. Various approaches can be considered for mitigating the impact of those uncertainties. We describe five approaches: arbitration, modeling, restrained use of quantitative criteria, conservatism, and non-quantitative approaches. Combinations of approaches are, of course, possible and may well have merit.

1. Arbitration - The decisions about whether a goal is met can be entrusted to a group of specially empowered experts.
2. Modeling - Verification could be tied to a prescribed mode of calculation.

3. Restrained Use of Quantitative Criteria - Use of the top-level probability target can be limited to decisions concerning the clearly prohibitive and the clearly trivial outliers. Another mode of restraint is to use an overall plant target only loosely and to provide some general guidance for setting quantitative targets at the level of systems and equipment for which failures can be tolerated with a frequency sufficient for possible verification.
4. Conservatism - Conservatism can be embedded in the calculations or explicitly applied to risk estimate results. In the "embedded" approach, estimates of actual risk are based on conservative assumptions, in order to increase assurance that decisions concerning the issue at hand will be on the safe side.

As an alternative, one can require calculations to be realistic -- the best one can do. Margins of due prudence would be introduced into administrative decisions. Sensitivity analyses may help discern the consequences of alternative judgments and thus help to set an appropriate margin.

5. Nonquantitative Approaches - Nonquantitative approaches available to reduce vulnerability due to uncertainties include independent lines of defense; use of requirements known to be beneficial, though not quantitatively; extra circumspection when abandoning known, established practices for promising innovations; and professional judgment, with emphasis on the qualifications and experience of decision makers and advisors.

D. Approaches to Balancing of Values

As the work of safety-goal development progresses, we plan to seek a sound balancing of the values relevant to goal setting. At this time we limit our discussion to noting in a preliminary way two specific topics: (1) individual and social risks and (2) safety and economic values.

1. Individual and Social Risks - A safety goal may focus on the risk of death or illness to which any one individual may be exposed, including permissible differences among individuals in different circumstances, such as workers, nearby residents who may receive some benefit as compensation for incremental risk, and the general population. The goal may also include consideration of the numbers of people exposed to a risk.

Social risk considerations also include such issues as equities of distribution of risks and benefits, possible transfer of risk to future generations (through long-lived radioactive materials), distinctions between prompt and delayed fatalities, and whether and how property damage risks (which may indirectly entail other health-and-safety risks) should be taken into account.

A special aspect of the social risk issue is involved in the concept of risk management, that is, whether and how to consider the risks of alternatives in setting a nuclear-plant safety goal.

2. Safety and Economic Values - Safety and economic values are not necessarily or invariably in conflict. Safety may protect plant investment and property of others as well as save medical and other costs of illness. In this positive respect, the issue is the extent to which economic incentives for safety should be reflected in safety-goal formulation. However, in many questions about safety, safety and economic values conflict or appear to conflict; the issue then involves, first, whether under the circumstances at hand economic impacts can be taken into account in a safety decision and, second, what amount of resources may be devoted (or what benefits foregone) to achieve a particular safety goal.

Decisions concerning possible retroactive applications of new requirements, including the safety goal itself, may involve special safety-cost trade-offs since the costs of specific safety improvements for plants already being built or operating may be substantially higher than for plants yet to be built.

III. SOME SAFETY-GOAL PROPOSALS

NUREG-0764 describes six safety-goal proposals, all urging a quantitative formulation of safety goals and suggesting specific numerical values for the parameters that they propose. The diverse frameworks can, however, be employed in conjunction with numerical values that differ from those suggested in the specific proposals. The intellectual task of devising a sound framework and the policy task of setting appropriate levels are distinct, though they must in due course merge. Both tasks are essential, as is their proper merging.

All six of the proposals are based on the use of probabilistic risk assessment. They depend for their effectiveness on the potentialities of that technique and are subject to its limitations -- gaps and uncertainties in data bases and difficulties of constructing adequate sets of underlying assumptions and of accounting for human factors in safety. Some of the approaches proposed include provision for accomodating or mitigating the effects of uncertainties.

With all the approaches, consideration would have to be given to (a) the appropriate role for other methods where probabilistic risk assessment cannot dependably resolve an issue and (b) problems of transition from current practices, including decisions with respect to possible retroactive application of new requirements, as well as (c) the development of sufficiently wide availability of an adequate level of expertise in the new probabilistic methods.

The proposals described in NUREG-0764 include one submitted to the Commission on October 31, 1980, by the Advisory Committee on Reactor Safeguards and

described by the ACRS as "a preliminary proposal for a possible approach to quantitative safety goals ... intended to serve as one focus for discussion on the subject ... and ... expected to be ... a first step in an iterative process." (NUREG-0739, "An Approach to Quantitative Safety Goals for Nuclear Power Plants," ACRS, 1980.)

The ACRS proposal is the only safety-goal proposal that has been formally presented to the Commission for consideration or discussion. However, we are aware of other proposals relating to development of a safety goal for use in the United States that have been elaborated in varying degree and have been published or widely communicated. Five such proposals are discussed briefly in NUREG-0764. Their sources are the Atomic Industrial Forum; Dr. Chauncey Starr, of the Electric Power Research Institute; Dr. Vojin Joksimovic, of General Atomic; Dr. Edwin Zebroski, of the Nuclear Safety Analysis Center; and Mr. Robert Bernero, of the NRC Office of Nuclear Regulatory Research.

IV. DEGREE OF SAFETY

The essence of a safety goal is the degree of safety that it establishes. The generic discussion of approaches to safety-goal formulation and of the criteria by which approaches may be judged has as its purpose the creation of a vehicle for specifying the degree of safety sought in a way that would make it meaningful and useful.

The assignment of specific proposed degrees of safety in the specific proposals summarized above illustrates the use of the proponents' frameworks and the degrees of safety suggested by the proponents as appropriate.

The Commission reserves judgment about any specific degree of safety, pending further progress in its safety-goal development efforts.

ADDITIONAL QUESTIONS

The Commission solicits comments and suggestions concerning the entire subject matter of the preliminary statement of policy considerations and the supporting discussion paper or any parts of that subject matter. In addition, comments are solicited concerning the following specific questions:

1. (a) Among the criteria for selection of an approach to safety-goal formulation that are presented, which are particularly important? Unimportant? Illustrative criteria are that the goal be comprehensive, logical, verifiable, practical, and publicly acceptable.

(b) Should additional or different criteria be considered? What criteria and with how much emphasis?
2. Which of the following are particularly important to include in a safety goal:
 - (a) Some general approach to risk acceptability?
 - (b) Quantitative safety goals?
 - (c) Qualitative --even subjective -- standards?
 - (d) Approach to safety-cost trade-offs?
 - (e) Goals for future safety improvements?
 - (f) Standards for determining when new requirements should be applied retroactively?
3. (a) Among the approaches to safety goal formulation that are discussed, what approach or combination of approaches is particularly appropriate? Inappropriate? Why?

(b) Should any other approach be considered? What approach?
4. (a) Among the approaches to dealing with uncertainty that are discussed, what approach or combination of approaches is particularly appropriate? Inappropriate?

(b) Should any other approach be considered? What approach?
5. What should be some of the characteristics of safety requirements:
 - (a) What should be the role of safety-cost trade-offs?
 - (b) To what extent should benefits of nuclear power, absolute and relative to alternatives, enter safety-requirements decisions?
 - (c) to what extent is it appropriate for requirements for new and previously approved plants to differ?
 - (d) Should a safety goal be applied directly to cases in order to attain a similar degree of safety from case to case (even though that may result

in specific design and operation requirements differing according to circumstances)? Or should the goal be applied generically and have requirements, rather than estimated degree-of-safety results, be uniform?

(e) To what extent should the goal reflect protection of individuals regardless of numbers of persons affected, and to what extent should it reflect total, integrated population or societal effects?

(f) To what extent should equities of distribution of benefits and adverse impacts influence requirements?

(g) Should the safety goal reflect increased aversion to risk of high consequences even at low probability?

(h) What is the proper balance between stability of requirements and flexibility for modification as knowledge develops and insights change.

6. (a) How should the stringency of nuclear-power-plant safety requirements compare with current practice?

(b) How should stringency of the safety goal compare with risks accepted from other (non-nuclear) electrical energy sources and with risks arising in various other contexts?
7. The Commission invites comments on the ACRS proposal and on the other specific proposals described in the paper, and would welcome any alternative proposals or suggestions.

Supporting Report

TOWARD A SAFETY GOAL: DISCUSSION OF PRELIMINARY POLICY CONSIDERATIONS

I. INTRODUCTION

A. Background

In accordance with the Plan for Developing a Safety Goal (NUREG-0735; also published in 45 FR 71023, October 27, 1980), the NRC Office of Policy Evaluation submitted for Commission consideration, on December 29, 1980, a preliminary policy paper, including a draft statement of preliminary policy considerations involved in developing a safety goal. After consideration of the preliminary policy paper, the Commission issued a statement for public comment. The Commission statement, as published in the Federal Register, is included at the front of this volume. This supporting report provides a more detailed discussion of the considerations underlying the statement. The draft statement and supporting report were submitted to the Commission by Edward J. Hanrahan, Director, Office of Policy Evaluation (OPE). Their principal author was George Sege, Senior Policy Analyst, OPE. An NRC inter-office steering group provided counsel.

In accordance with further provisions of the Plan, the comments received on the preliminary statement and this more detailed report will constitute a part of the considerations leading to preparation of a policy paper, including a draft policy statement. This paper will receive Commission consideration later in 1981. Subsequently, the Commission intends to seek additional public comment on the policy statement based upon the August 1981 paper.

B. Purpose and Scope of this Report

The purpose of this Discussion of Preliminary Policy Considerations is to indicate and examine the kinds of considerations which may enter into an articulation of the Agency's safety goal.

The scope of early policy development efforts has focused primarily on radiological accidents associated with power reactors. It is, however, intended to include considerations of sabotage and other external events and to seek adaptation of the safety goal to other nuclear facilities and operations where the need is evident.

Furthermore, a safety goal to be adopted by the Commission must be consistent with the Atomic Energy Act and the Energy Reorganization Act of 1974, which created the NRC. The essence of any useful safety goal is the balance struck between safety and other competing interests. Thus, the evident intent of Congress places at least some polar limits on the safety goals we might consider. A goal of "zero risk," for example, would ensure safety at the cost of sacrificing all other goals which motivated Congress's approach to nuclear power regulation. Courts have made it clear that "zero risk" is not the standard for the NRC to apply. At the other pole, a safety goal which gave decisive weight to the interest of promoting nuclear power development, even at the cost of high risk to the public, would clearly be inconsistent with the non-promotional posture established for the NRC by the Energy Reorganization Act of 1974. Apart from these obvious extremes,

however, the legal constraints on what safety goal might be appropriate for NRC adoption are not clearly defined.

C. Organization of the Discussion

A safety goal must be formulated in a way that is useful in a practical, regulatory sense. The Discussion begins with an exposition of the criteria by which the quality of a candidate approach to safety-goal formulation may be judged. No one definitive set of criteria has been selected; rather, criteria are at this time discussed as options regarding which public comment is sought.

The Discussion next addresses alternative approaches to safety-goal formulation. This section includes alternative general methods of approach, types of safety goals, types of quantitative approach, approaches for dealing with uncertainty, and comments about some aspects of balancing of risks and of safety and other values.

Presented next are some safety-goal proposals. These include a proposal recommended to the Commission for preliminary consideration as an option by the Advisory Committee on Reactor Safeguards and five other proposals known to the Commission.

Finally, there is a short discussion of degree of safety, since all or most of the potential approaches can be adapted to establishing a degree of protection at any level over a wide range of possible levels. Selecting that level on a sound basis is, of course, the paramount objective of the effort to develop a safety goal.

II. CRITERIA FOR ASSESSING AN APPROACH TO A SAFETY GOAL

A. Considerations in Formulating Criteria

For a sound, reasonable safety goal to result, the approach to its formulation must be reasonable as well. To arrive at this determination requires an explicit discussion of what constitutes a suitable approach to the problem. The formal result of that discussion is guidance in the form of a set of criteria for assessing an approach to safety goals.

Such a set of criteria can be stated, ordered, and emphasized in various ways. While several sets of criteria have been proposed, there is not a wide divergence of opinion among the proposers on the key criteria. Authors have named and stated their proposed criteria somewhat differently, but the proposals are similar in essentials. Since the range of possible criteria and sets of criteria is great, some reason for choosing among them is also necessary. At the very least, sets of criteria should embody five basic characteristics: they should require the goal to be comprehensive, logical, verifiable, practical, and publicly acceptable.

Drawing on the various proposals, a set of criteria which incorporates these five characteristics for assessing an approach to a safety goal is

given below. As far as we can now determine, these represent the key criteria on which there is apparent agreement. The following description is brief since we do not foreclose additions, deletions, and changes to the formulation; but we do consider this formulation as appropriate for the next phases of our program to develop a safety goal. These criteria are not in conflict with the various sets of criteria previously proposed; they represent a synthesis of the ideas in the previous proposals and they do not differ greatly from them.

1. Comprehensive - The approach should accommodate a broad problem definition addressing the significant risks of social concern, including early and delayed fatalities, injury, and economic effects. The comprehensiveness of the safety-goal formulation involves two related basic questions:

- To what extent does the formulation address significant elements of the issue?
- How comprehensive is it in the facilities and activities that it covers?

A part of this characteristic relates to the generality of the policy; another, to what special issues are included, notably:

- Safety-cost trade-offs.
- Relation between goals for plants in operation, being built, or planned.
- Goals for future improvements.

A major current issue in safety-goal formulation relates to power reactor accident considerations, more particularly, major accidents. Specific criteria within the accident-safety issue could relate to whether and to what extent the stated goal policy includes within its sphere of determination or influences: rules and case decisions on plant design, operation, siting, emergency planning, degraded-core provisions, anticipated transients without scram, offsite contingencies, etc. They could relate to whether it includes a policy on decisions under uncertainty and the extent to which it excludes accident considerations that pose difficult policy problems (e.g., any implications for dilemmas of operational discipline vs. individual rights).

2. Logical - There should be a clear and justified decision rule which produces coherent and rational results.

It is an underlying assumption in efforts to develop a safety goal that a cohesive structure connecting the important interrelated issues in safety regulation can lead to better safety and better management of resources devoted to safety.

A result-oriented test of a proposed policy is whether it can be applied to different aspects of a set of issues and whether, when so applied, it leads to reasonable results in each (e.g., siting, containment features, degraded core, ATWS, operator training).

These are the characteristics that watch for pitfalls in terms of unintended effects. Specific characteristics include these:

- Does the policy direct attention and resources to issues that are significant?
- Does it encourage or stifle initiatives to improve safety where warranted?

A stark articulation of the reasonableness-of-results tests is:

- Though the goal statement appears acceptable, does it lead to results that are not? Notably, could it be used to justify licensing a reactor where that result is not intended? Would it lead to shutting down a reactor or denials of applications and do so not for good cause, but as an unintended result of a rigid goal formulation?

3. Verifiable - The approach should make it possible to establish the relationship between the performance of a facility or operation and the safety goal, in order to measure or otherwise determine compliance. An approach must be clear in its underlying assumptions, philosophical roots, matters included and excluded, inputs required, analytic methods, and uncertainties as well as lead to actions consistent with the goal.

The question of how clearly a safety goal can be interpreted in terms of regulatory actions (on which hinges the predictability of such actions) involves not only the clarity of the goal statement but also the possibility of verification, of measurement. The basic question is:

- How clearly can it be known whether a particular action serves or is consistent with that goal?

How can one know, or reasonably judge, whether some low probability target (say, 10^{-6} per reactor-year) is met? It should be noted that absence of strict verifiability, in a statistically meaningful sense, does not doom all quantitative goals for high-consequence, low-probability events to failure on this criterion. However, available options to deal with the verification problem may introduce weaknesses with respect to other criteria by leaving troublesome decisions outside the purview of the goal or by introducing some fallible arbitrariness into its interpretation.

4. Practical - It must address and lead to the resolution of real problems by real people within real resource constraints. Moreover, the approach

should be compatible with social interests and institutions. It should not require drastic changes in the usual operations of the institutions involved in its implementation, but it should also be able to influence those operations, where necessary, toward the implementation of the safety goal.

The practicality of a safety-goal approach depends on how adequately the pertinent technical, industrial, institutional, and human capacities and limitations are taken into account. Significant aspects include:

- Implementability. To what extent does it build on the most useful available technical knowledge? To what extent does it build on industrial capacity to comply? To what extent can the policy be in fact expected to be effectively translated into useful regulatory and licensing results by NRC staff, Licensing Boards, the industry, intervenors, the courts?
 - Realism in the face of uncertainty. Does it properly reflect recognition of uncertainties concerning probabilities and consequences of unwanted events, including uncertainties that can be narrowed by experience, research, and analysis, and those that cannot?
 - Flexibility. How well can it accommodate new knowledge, new insights, new potentials for worthwhile improvements, correction of gradually discovered error?
 - Efficiency in use of resources. How well does it channel resources for hardware, operation, analysis, and regulatory and licensing processes to areas of high safety benefit?
 - Transition problems. What problems are involved in transition from current practices to the new practices demanded by the goal? Can it accommodate reasonable safeguards against counter-productive dislocations in going from known, established approaches to approaches that are more promising but less tried, less known?
5. Publicly Acceptable - The application of the approach should be a process that is sensitive to and accounts for people's views and ensures their early and continued participation.

The question of the extent to which a safety goal is acceptable to affected groups is complicated by the heterogeneity of the public. Industry, nearby residents, the general public and its various interest groups are affected differently. Values and perceptions differ, but the ultimate approach should be acceptable to a majority and not be totally unacceptable to any significant fraction of the public.

There is a particular burden on the agency to seek to impart to the public a sound perception of risks addressed and of the relation

between policy and result. A significant criterion in the public-acceptability category is whether a policy is distorted by dichotomies in technical and public perceptions or by damaging errors in public understanding of the facts or in agency predisposition. The agency must not be an advocate of nuclear technology; rather its job is to set out what is known and, under Congressional oversight, establish and implement working safety goals.

B. Proposed Criteria: Formulation 1

A set of criteria for selection of an approach to acceptable risk for NRC consideration was prepared by a team led by Dr. Paul Slovic, of Decision Research. The Decision Research team identified seven qualities desired of an approach to making acceptable-risk decisions: an approach should be comprehensive, logically sound, practical, open to evaluation, politically acceptable, compatible with institutions, and conducive to learning. These criteria are intended by the authors for use in an analysis of possible approaches, to examine how each approach could, in principle, satisfy the criteria and how well each approach currently does so in practice. They note, however, that although such an analysis evaluates the decision options from various perspectives, it does not tell which to choose. Such judgments of importance might reflect personal values, legislative mandates, or the exigencies of particular situations. The approach preferred for one problem might be rejected in another situation for which its particular strengths (e.g., political acceptability) were not essential.

The Decision Research team's discussion of the seven criteria follows. This discussion is an edited version of a report adapted by the authors from Chapters 2 and 3 of their Approaches to Acceptable Risk: A Critical Guide, NUREG/CR-1614; ORNL Sub-7657 (U.S. Nuclear Regulatory Commission; Washington, D. C., 1980)

Like other decisions, deciding how to decide involves a choice among alternatives. The alternatives are the various possible decision-making methods, including the "do nothing" option of letting matters run their own course. The choice among them is difficult, in part, because the alternatives are not readily compared. Each embodies a somewhat different concept of what concerns need to be addressed and what constitutes a rational procedure for amalgamating those concerns into a decision specifying a course of action.

One way to compare diverse methods is to evaluate each in terms of a common set of criteria. This discussion offers one such set, representing what a society might want out of an approach. It describes the importance of each criterion and how it might be applied to various decision-making approaches. Having evaluated how well each of a set of approaches satisfies each criterion, in theory and in practice, does not by itself indicate which approach to choose. Unless one approach surpassed all others in all respects, it would still be necessary to decide which criteria are most important. These importance weights might be different for different individuals and institutions. In that case, some procedure would be needed

to establish "official" priorities. Weights may also vary across decisions, with the approach preferred for one problem being rejected in another for which its particular strengths were not as essential. In this case, some procedure would be needed for identifying classes of decision problems with their respective criterion priorities.

1. Comprehensive

Acceptable-risk decisions are among the most difficult faced by modern society. Addressing all the issues they raise requires expertise from the natural, mathematical, and social sciences, as well as from engineering and the humanities. These issues may be described in terms of five generic sources of uncertainty or complexity facing any decision-making approach. Failure to address all of them explicitly and persuasively means that an approach is, at best, solving only part of the problem.

Uncertainty about problem definition. The problem definition establishes the universe of discourse for the decision-making process. It determines (a) that there is a decision to be made; (b) what action options and consequences (e.g., risks, cost, benefits) are valid considerations; and (c) what kinds of information and uncertainty are worthy of note. Often the decision has effectively been made once its definition is set. A method should encourage a creative and comprehensive definition. Failing that, it should at least indicate what has been left out and the implications of those omissions.

Uncertainty about the facts. Incomplete and conflicting information is the lot of most acceptable-risk decisions; an approach must be able not only to live with that uncertainty, but to treat it in an explicit and credible fashion. It needs to elicit and consider the range of opinion within the relevant technical fields as well as the limits to the methodologies used by those fields, so as to provide non-technical decision makers with a realistic appraisal of what is known.

Uncertainty about values. Acceptable-risk decision makers are typically entrusted with acting in accordance with some vaguely defined "society's best interest." Their task is complicated by the conflicting arguments regarding how that definition should be sharpened and the natural tendency for arguers to mix their own best interests with those of society. An approach should acknowledge these uncertainties about how to evaluate different consequences and explicate the assumptions made by various evaluation procedures (e.g., different economic analyses). Furthermore, it should identify cases in which the issues are so novel or complex that society or its citizens do not have articulated positions.

Uncertainties about the human element. Acceptable-risk problems involve a variety of human actors. As consumers, voters, legislators, regulators, operators, and promoters, people shape technologies, the

world within which they operate, and the effective degree of risk and benefit that society derives from them. To be useful, a method must make explicit, realistic assumptions about how these various people perceive and respond to risks and to information about them.

Uncertainties about decision quality. An appraisal of the overall quality of the decision reached by an approach tells users how much confidence they should place in its conclusions. It tells the purveyors whether they should try again before reaching any conclusions, perhaps by recruiting more information, assessing value issues more thoroughly, consulting additional individuals, changing the problem definition, or using an alternative method. In principle, an approach should be capable of reporting that it is not up to the task, either because the uncertainties are so great as to render its conclusions indeterminate or because crucial uncertainties lie in areas that the method does not address.

2. Logically Sound

Delineating the problem is not synonymous with providing guidance. Indeed, comprehensiveness alone can lead to confusion and frustration. To be useful, an approach must provide a timely and logically defensible summary of all that it encompasses. Without such summaries, "analyses" can unfairly discourage projects by inducing a feeling that "we should not mess with anything that is so poorly understood," breed mistrust by making observers think that "they must be hiding something in that morass," or encourage capricious action by suggesting that "we might as well go ahead with this project since there is not convincing evidence against it." Thus, a viable approach must produce some conclusion, if only "collect more data; we do not know enough to decide at the moment."

In addition, that conclusion must be derived via a defensible decision rule. Such a rule would be:

- (a) Sensitive to the various aspects of a decision problem; changes in the alternatives, facts, values, or uncertainties considered should be capable of leading to different recommendations;
- (b) Reliable (or reproducible), in the sense that repeated application to the same problem should produce the same result;
- (c) Justifiable, in terms of either theoretical arguments, demonstrating why it should lead to good decisions, or empirical evidence showing that it has worked in the past;
- (d) Suitable to societal risk problems and not uncritically imported from the other realms (e.g., from corporate decision making or problems without potential loss of life); and

- (e) Unbiased in its recommendations, not giving undue weight to any interest or type of consideration.

3. Practical

Like the technologies they are meant to manage, decision-making methods must work in reality, as well as look good on the drawing board. It must be possible to implement the approach with real problems, real people, and real resource constraints.

Real problems. To apply an approach, one must establish a reasonable correspondence between its terms and equivalents in reality. Cost-benefit analysis, for example, has limited usefulness when generally accepted operational definition of "cost" are lacking. Any approach could fail if it used one statistical summary of risk (e.g., expected annual fatalities) when policy makers were interested in others (e.g., catastrophic potential), or if it were able to consider only a fixed set of alternatives in a reality that persisted in creating new ones.

Real people. Weighing strategies for the management of a technological hazard is a labor-intensive enterprise drawing upon a select pool of skilled individuals, including substantive experts (those who know most about a particular hazard) and normative experts, specializing in decision making per se. Can enough of these special people be recruited for a reasonable facsimile of the approach to be implemented? If experts can be found, does their task use them to best advantage? Is it too novel and complicated to be comprehended? Do its questions fit the cognitive structure of the expert's knowledge? Finally, an approach must consider any possible bias on the part of the experts, due to pecuniary or scientific vested interests. These questions become more acute as the scientific data base shrinks and experts are asked to create instant knowledge, in the form of educated intuitions, rather than draw upon the fund of knowledge that has undergone peer review.

Resource constraints. An approach must work adequately in the time and money constraints of real-time decision making. Often, decision makers turn to a new method when they need help quick, in order to respond to a crisis in which traditional decision-making procedures have failed. Even when time is available, decision makers may be reluctant to spend hard cash for the probabilistic benefits of good advice, which at best increases one's chances of making the right choice. If an approach needs to be done right or not at all, it should at least identify situations in which some analysis is not better than none at all.

4. Open to Evaluation

As discussed earlier, assessing the quality of a decision or method is hard under the best of circumstances. An approach should not make matters worse by obscuring its internal functioning. All those

whose fate it is deliberating have a right to ask: What are its underlying assumptions? What are its political and philosophical roots? What options does it foreclose or prejudice? Where are fact and value issues mixed? What inputs were used? What computational procedures were followed? How much uncertainty surrounds the entire enterprise?

Providing answers to these questions is essential to the validity of an approach. Many acceptable-risk questions are so complex and multidisciplinary that no one can expect to get the right answer on the first try. In such cases, an approach should indicate that it is only an approximate answer. Moreover, it should invite constructive criticism designed to spot omissions, errors, and hidden assumptions that can be treated in a subsequent iteration. Even destructive criticism may be better than none at all, if it catches some problems and adds some new perspectives.

Evaluation is particularly frustrated by poorly defined procedures and lack of conceptual clarity. The unexamined approach is hardly worth using. An approach that fails to test its effectiveness and clarify its prejudices is not to be trusted.

5. Politically Acceptable

An approach can fail in the harsh, politicized world within which hazards are managed because it works too poorly, works too well, or works in a vacuum.

If an approach has serious conceptual flaws (e.g., because it misdefines the problem or has no defensible integration rule), critics will readily impeach any displeasing recommendations it produces. For example, the fuzzy logic of some environmental impact statements exposed them to interminable litigation by parties dissatisfied with their conclusions (or intent on procrastinating). At the other extreme, an approach may encounter little resistance because it is so poorly operationalized that all interested parties see how they can manipulate it to their own purposes. In time, combatants may learn to conduct their debates in, say, the nomenclature of cost-benefit analysis, transforming the techniques into a rhetorical device and voiding its impact.

Conceptual strength can also encourage political complications. If an approach produces a clear, persistent, and unwanted signal, the offended parties may choose to discredit it, rather than just fight one particular conclusion. For example, an approach that redressed an existing imbalance of power between producers and consumers, employers and employees, or laborers and the general public might be attacked by the side whose advantage it jeopardized.

Finally, an approach can fail by disregarding means in its quest for the optimal end. In any participatory system, recommendations must

be sold as well as generated. One aspect of that selling job is to insure that people's views have been accommodated. Usually that means asking them early and sincerely enough to affect even the problem definition. Attention to the process of decision making may also facilitate the creation of solutions, by negotiating settlements between opposing parties. Moreover, a good process may itself have positive consequences, such as helping participants live and work together, reducing social alienation, and enabling participants to monitor a decision's implementation by educating them in its rationale and technical details. With a successful approach, process may be its most important product.

6. Compatible with Institutions

For better or worse, hazards are being managed today. To accommodate this management, a complex of social institutions has evolved. An approach's chances of survival drop as it departs from the standard operating procedures of the institutions. Even a method that satisfies the other six criteria might not get very far if no one is empowered or ordered to heed its recommendations, if legal precedents bind the hands of crucial actors, if it fails to produce the paperwork required for documentation, or if the personnel it requires are neither found nor wanted in the relevant agencies.

On the other hand, an approach may fit too well. Institutions have their own agendas, which need not coincide with those of the people they represent. Decision makers in some institutions may prefer an approach that cloaks their decisions in ambiguity, reduces their accountability, establishes a position for them in hazard management, defers difficult value questions to external "experts," or studies hard issues forever. Hence, the institution as well as the approach may need adapting.

7. Conducive to Learning

Attempting to satisfy these criteria encounters a fundamental conflict: the need to respect political and institutional realities without being overwhelmed by them. A final objective is to change those realities. An approach should educate its participants, eliminate opportunities for obstructionism, and build up its own record of precedents. Somehow society should become better or wiser for its adoption. Achieving this objective might even lead one to sacrifice short-term benefits, such as efficiently solving a particular problem, for the sake of long-term goals, such as developing generic standards.

Features that make an approach conducive to learning include: (a) leaving a clear record of deliberations and assumptions, to facilitate evaluation and the cumulation of knowledge; (b) encouraging two-way communication between scientists and decision makers, to improve understanding of one another's problems and uncertainties; (c) educating lay observers, to enhance their ability to follow the process

and develop expertise in the subtleties of acceptable-risk questions; (d) having enough generality to be used on many problems, allowing users to acquire an in-depth understanding of one technique, rather than a superficial grasp of many problem-specific methods.

One, more active, role an approach might fulfill is recruiting talented scientists and lay people to a problem. Another is alerting users to recurrent oversights. A third is indicating generic categories of hazards that can be managed in a consistent fashion, drawing on the same decision-making effort. A fourth is increasing the credibility of society's decision-making bodies by offering them more trustworthy tools. Perhaps the most general criterion for judging the contribution of an approach to long-term effective management may be whether it raises the level of debate.

C. Proposed Criteria: Formulation 2

An alternative taxonomy is presented in a paper prepared by Roger Mattson, Malcolm Ernst, Warren Minners, and Miller Spangler, all of the NRC Office of Nuclear Reactor Regulation (NRR). Their paper, "Concepts, Problems, and Issues in Developing Safety Goals and Objectives for Commercial Nuclear Power" (draft dated August 13, 1980), has been submitted for publication in Nuclear Safety. The discussion of its taxonomy of six criteria appears below.

Before NRC chooses any particular set of safety goals, it needs to consider alternatives. The safety goals that are selected for more in-depth consideration as alternatives will then need to be evaluated comparatively against some established criteria. Generally, it might be expected that the same set of criteria will more or less serve both the selection and the evaluation of alternative goals. First of all, criteria for candidate goal selection and evaluation would comprise estimates of their effectiveness in serving their societal purposes. A re-examination of the five purposes suggested above [in their paper] reveals that they are highly subjective. Moreover, as previously noted [in their paper], the effectiveness of the goals will in part depend upon program strategies and resource allocations. Also, perceptions of how well various goal candidates serve these purposes may differ widely, if past experience is a guide. With these thoughts in mind, we propose that it would be helpful if the primary purposes for a safety goal formulation were buttressed by additional selected criteria of the type discussed below.

1. Comprehensiveness in Covering the Various Risks of Significant Social Concern

It is difficult, on its face, to imagine how a single, overall safety goal could speak to all of the various risks of significant concern as these are presently perceived by different segments of the public, the nuclear industry, the NRC and other agencies of government. The complexity introduced by the need to speak to multiple audiences is

compounded by the question of how the safety goal might be perceived by future generations whose interests are, to a significant degree, entrusted to our decisions. Both the National Environmental Policy Act of 1969 (Sec. 101(a) and (b)(1)) and the Energy Reorganization Act of 1974 (Sec. 2(a)) provide policy directives regarding impacts on future generations.

The variety of audiences is not the only reason that a single overall safety goal might not suffice. The overall risk of nuclear reactors would not cover the overall risks of the nuclear fuel cycle which, in turn, are important to a comparative analysis of risks of alternative sources of energy for generating electricity. On the other hand, an overall goal which addresses the risk for all elements of the fuel cycle would conceal the spatial and temporal distribution of risk. It therefore would deny an exploration of equity considerations in safety policy decisions.

Moreover, an overall safety goal for nuclear reactors at each given site of not exceeding some particular probability of excess deaths per year (or shortening of life expectancy) would not, of itself, address concerns for the kinds of serious economic damages to consumers suffered in the TMI accident nor concerns for land use or property damages external to the plant boundaries in the event of an accident involving substantial offsite releases of radioactivity. Nor would such a goal address the costs in human anxiety often associated with the catastrophic and involuntariness aspects of nuclear risks. Thus, the comprehensiveness of the proposed safety goal, or the candidate set of safety goals, in covering the various risks of significant social concern is an important criterion in their selection and making comparative evaluations for deciding on safety goal formulation.

2. Verifiability of Goal Attainment

Safety goals which have a low order of verifiability would not be expected to produce a high level of acceptance. Certain kinds of quantitative as well as qualitative goals may suffer from a low order of verifiability. In some instances, qualitative goals may be stated in such vague or ambiguous terms that various parties will interpret them differently or perhaps be confused as to how to interpret their practical significance. Similar confusion and disagreement may also accompany the ex post facto interpretation of historical experience in the verification of previously established goals. Likewise, certain quantitative goals may pose difficulties of verification. For example, 30 tosses of a coin may provide a suitable number of observations to provide reasonable verification that the actual probability of tails is 50 percent for a single toss of a coin. However, a safety goal involving accident rates of low order probabilities (such as once in a hundred thousand or a million reactor years of experience) is scarcely verifiable with 500 or even 1,000 years of reactor operating experience.

Nor, if such an event did happen within such a period of experience, can it be verified that the correct assessment of probabilities should have been once in 500 or once in 1000 reactor years. There still remains a broad range of uncertainty that the true value is greater or less than such an empirically derived frequency. The verification problem is even worse when one considers the heterogeneity of reactor designs and site or regional characteristics of importance to safety events, as well as variations in human factors affecting the design, construction and operation of each particular nuclear plant. The problem, of course, may be made less severe if an overall quantitative goal involving a low-frequency event does not stand by itself but is part of a more comprehensive goal set in which supporting goals and standards have a higher order of verifiability. Accordingly, close attention to the criterion of verifiability in the design of a candidate goal, or goal set, can appreciably improve its comparative rating. For example, an intermediate level of detail in a particular goal set might treat the rate and the consequences of anticipated operational occurrences as they are presently defined in 10 CFR. Such a goal set would get a higher rating on this attribute than a goal which treated only the rate and consequences of core-melt accidents.

3. Clarity and Understandability of the Goal

To an important degree, this criterion interfaces with the previous criterion of verifiability. A goal statement that suffers from ambiguous or imprecise wording (such as "adequate protection" or "no unreasonable risk") is likely to have a low order of verifiability and at the same time deserve a low rating for clarity and understandability. Nevertheless, there can be an important difference between the two attributes when one considers the wide variations of education, experience and personal biases affecting preference and understandings of word choices and numerical forms. For example, stating a safety goal in negative powers of 10 may have clear and precise meaning to technologists or others trained to understand such terminology, but produce a bewildering effect on people who lack such training. The same may be true for supporting safety standards whose attributes are expressed only in terms of millirads, picocuries or person rems. In some cases, only simple adjustments to the goal statement would need to be made to improve the rating of this criterion. The evaluation of this criterion requires that one not simply ask whether the goal formulation is clear and understandable, but to whom it is clear and understandable. This may suggest that an acceptable goal set must have several equivalent wordings at various levels of scientific sophistication.

As discussed below, goal achievement is often subject to a number of factors that are partially, if at all, controllable by the parties held responsible for goal attainment. For example, the incidence and magnitude of certain natural or man-made events such as earthquakes, floods, fires, or explosions leading to common mode failures may be beyond the full control of nuclear regulators, equipment

manufacturers, or plant operators. Population densities and land use patterns in the vicinity of nuclear plants may also change. Even when the subject of control is not at issue, there may be significant deficiencies of information for which certain assumptions or engineering or management judgements are then made.

Thus, for improved clarity and understandability, the goal, or goal set, must be conditioned by explicitly stating the underlying assumptions so that the practicality and degree of realism of goal attainment can be better appraised. In the long run, a conditioned goal will probably instill more trust than if seemingly unconditioned goals are presented and the essential underlying assumptions are unarticulated, hidden, or neglected in the goal formulation process.

4. Harmonization of the Safety Goal, or Goal Set, with Other Societal Interests

The degree to which the expressed safety goals are compatible with other societal interests is an important attribute. A large number of writers on risk assessment and risk management view the fundamental problem as one of balancing risks with other costs and benefits and especially in comparison of risk-acceptance of risk-rejection decisions vis-a-vis alternatives. The statement of goals and objectives in NRC's first Annual Report recognizes the vital aspect of examining such trade-offs as does the President's Commission recommendation (for NRC) to establish safety-cost trade-offs. It is clear that the term "cost" should not be viewed in the narrow sense of direct dollar cost associated with safety measures. Rather, cost should also include the full gamut of adverse societal effects envisioned in NEPA requirements. Moreover, there is the additional perspective of "opportunity costs," which are a measure of the benefits foregone in the alternative use of resources.

5. Defensibility of the Goal, or Goal Set

The defensibility rating of the candidate goals would be determined by a variety of attributes, some of which are overlapping the above criteria. Additional aspects of defensibility are conceptual soundness, practicality of implementation, adaptability to available resources and institutional constraints, and flexibility in accommodating diverse situations.

The defensibility of the goals will also be determined by the acceptability and explainability of the procedures by which they were formulated. A goal-setting procedure involving public input at all stages and providing an explicit account of the considerations that enter into the subjective judgements will stand a superior chance of being defensible.

6. Safety Goal Advocacy and the Division of Responsibility

The establishment of effective, defensible safety goals will demand a certain measure of statemanship in taking into account the initial pattern of safety goal preferences of the principal advocates as well as known predispositions of other influential segments of society. Several advocacies of quantitative safety goals have been cited in the introduction, and an even larger number is found in the literature dealing with technological risks. There are also a number of advocates of qualitative safety goals, particularly for those kinds of risks for which numerical verification is questionable or where there is a failure of a single, or several, quantitative safety goals to be adequately comprehensive in reflecting the key risks and interests of societal concern.

Obviously, there is more room for accommodation of various stake holders if, in lieu of a single goal, a goal set is chosen which may have a combination of goal forms organized in a consistent and logical fashion. For example, the supplementing of an overall safety goal with a well chosen set of supporting objectives of safety standards may provide a strength and acceptance in their unity that would not be achievable when viewed separately. The design of an effective goal set will require close attention to the profile of strengths and weaknesses of each member of the set in the light of the above evaluation criteria, including the known predispositions or positions of advocacy among influential segments of the public. No such goal harmonization efforts, however, can be expected to receive acclaim or even acceptance from all quarters.

The NRC should not establish all of the goals or supporting standards by which safety is promoted. Some part of this effort should be conducted by nuclear equipment vendors and nuclear utilities because of their primary responsibility for designing, constructing, and operating nuclear plants safely. Other reasons for a division of effort in safety goal formulation between NRC and the industry are those of resource allocation and utilization as well as division of authority (e.g., for "hands on" operations). However, the responsibilities among the various participants in nuclear safety goal formulation need to be clearly defined and a schedule of complementary goal formulation efforts needs to be established so that a reasonable degree of agreement on the goal formulation process can be reached in a timely way.

D. General Observations

Slovic et al. (op. cit.) say the following concerning the criteria that they propose -- and this can be said of alternative formulations as well:

Like the approaches they are designed to evaluate, these criteria are not entirely independent. Weakness in some respects may preclude strength in others. An approach is unlikely to be comprehensive if

it does not elicit competent criticism from a variety of perspectives. Without openness to evaluation, there is little opportunity to learn from experience and increase understanding over the long term. An approach with obvious logical flaws is unlikely to fare well politically. As a result, an approach that stumbles in one respect is likely to encounter other difficulties as well.

On the other hand, some of these goals may be in conflict. It may be easier to find a logically sound integration rule if one leaves out certain awkward issues, thereby failing to address parts of the problem. Political acceptability may require involving so many parties in the decision-making process that the constraints of the responsible institutions are overwhelmed. Openness to evaluation may mean vulnerability to cheap shots and unfair criticism, thus impairing political acceptability.

If no approach does, or even can, satisfy all of these criteria and if their respective strengths and weaknesses lie in different realms, then we must decide what we really want. As a result, the choice of an approach is a value-laden and political act, reflecting our preferences of how society should look and function.

Whatever the basic formulation of criteria, a wide range of evaluation options exists in terms of the degree of emphasis accorded to the various categories of criteria or to specific criteria within each category. Thus, one might emphasize comprehensiveness of the policy, at the cost of toleration of interpretive difficulties. Alternatively, clarity and ease of interpretation could be emphasized, at the cost of limiting the policy's reach to the more readily tractable parts of the issue.

On the basis of the foregoing discussion, specific sets of criteria, with assigned degrees of emphasis, could be structured into a selected number of discrete evaluation-method options. We do not do so at this time, and we leave open for now the question of whether such discrete options should be constructed later in the Commission's safety-goal program.

Indeed, we recognize as an open question to what extent it is desirable to articulate a full set of criteria, perhaps with weights assigned, in advance of the approach-selection decision process. An alternative to full a priori articulation is a process in which definitive articulation of specific criteria and decisions about the relative weights to be assigned them are arrived at as an integral part of the decision process.

Finally, we note that, in the program's early stages, the criteria should be applied tolerantly so that candidate approaches would not be too easily rejected from further consideration. As the program progresses towards the narrowed set of options of the later policy paper, a more demanding interpretation of criteria will, it is hoped, lead to the desired narrowing -- and perhaps eventually to a single recommended approach.

III. APPROACHES TO SAFETY GOAL FORMULATION

An approach to developing safety goals presupposes the existence of and compliance with a set of criteria which validate it. It also presupposes its competence to address the complexities of the issue of nuclear safety and of assessing proposed safety goals or amending existing ones. These presumptions allow a wide range of possibilities, however, since approaches may differ in method, form, and emphasis, and be formal or informal, simple or complex, in their application to the issue before them.

In this section, discussions progress from a characterization and assessment of three broad kinds of approaches, to a consideration of certain features inhering in possible safety goals, to an examination of types of quantitative approaches, to a brief account of means of dealing with uncertainty, and finally, to comments about balancing disparate values.

A. Methods of Approach to Acceptable Risk

Slovic et al. proposed classification of approaches to acceptable risk into three broad methodological categories: formal analysis, bootstrapping, and professional judgment. These are described in NUREG/CR-1614 (Chapter 3) as follows:

1. Formal Analysis

Formal analysis assumes that intellectual technologies can help us manage the problems created by physical technologies. Cost-benefit analysis and decision analysis are the most prominent techniques for thinking our way out of whatever troublesome situations we have created for ourselves. Evolving from economic and management theory, these approaches share a number of common features:

- (a) Conceptualization of acceptable-risk problems as decision problems, requiring a choice between alternative courses of action. For example, cost-benefit analysis attempts to identify the option with the greatest preponderance of benefits over costs.
- (b) A divide-and-conquer methodology. Complex problems are decomposed into more manageable components which can be assessed individually and then combined to provide an overall assessment.
- (c) A strongly prescriptive decision rule. The components are combined according to a formalized procedure; if one accepts the assumptions underlying the analysis and its implementation, then one should follow its recommendations.
- (d) Explicit use of a common metric. Decisions are hard when one must make value tradeoffs between conflicting objectives. In order to compare different consequences, formal methods reduce them to a common unit (e.g., dollar value).

- (e) Official neutrality regarding problem definition. These techniques are intended to be applicable to all problems with clearly delineated consequences, measurable options, and identifiable decision makers.

Purveyors of formal analysis tout its potential rigor, comprehensiveness, and scrutability. Skeptics wonder how often this potential is realized. Are analyses accessible to interested observers? Can all consequences and options of interest be accommodated? Don't actual applications have a more ad hoc flavor than the theory would suggest? Critics also worry about power being concentrated in an intellectual elite, analysts failing to appreciate the organizational impediments to implementing recommendations, and ideological biases lurking in the ostensibly neutral assumptions underlying the methods.

2. Bootstrapping Approaches

Whatever theoretical appeal formal analysis may have, the technical difficulties encountered in trying to conduct an analysis have led some observers to despair of ever devising a comprehensive formula for acceptable-risk decisions. An alternative approach, which produces a quantitative answer without recourse to a complicated formula, relies on first identifying and then continuing policies that have evolved over time. Proponents of this family of approaches argue that society achieves a reasonable balance between risks and benefits only through a protracted periods of hands-on experience. The safety levels achieved with old risks provide the best guide to how to manage new risks. Assuming that one has identified such an equilibrium state, the balance between costs and benefits achieved there should be enshrined in future decisions, short-circuiting the learning and adjustment process and, in effect, lifting ourselves up by our own bootstraps.

One member of this family, the revealed-preferences approach, uses the cost-benefit tradeoffs effected by our market, social, and political institutions in the recent past as prescriptions for future balances. Another member, the natural-standards approach, looks to the geologic past; it argues that the ambient levels of pollution during the development of a species is the level to which that species is best suited and the level to be sought when setting future tolerances. In either case, a description of past policies is taken as a prescription for the future. The resultant policy should be consistent with existing decisions and be sensitive to complex tradeoffs that are hard to accommodate in formal computations. One conceptual limitation of bootstrapping is that for new hazards, which are often the most troublesome, there may be no relevant experience to which to refer. Another is that these methods pass judgment on the acceptability of individual options, without explicitly considering the alternatives. One possible political limitation is bootstrapping's strong bias toward the status quo; it assumes, in effect, that whatever is (or was), is right for the future.

3. Professional Judgment

Another response to the possibility that there is no one formula for determining "how safe is safe enough?" is to rely on the judgment of the technical experts most knowledgeable in a field. Professional judgment is exercised whenever a physician decides that a by-pass operation or immunization program is worth the risk, a civil engineer decides that soil porosity has been adequately handled in the design of a dam, or a boilermaker decides not to reinforce further a potentially leaky joint. In making their decisions, professionals might avail themselves of formal analyses, if such existed, but they are not bound by the conclusions of those analyses nor need they articulate the reasons for their decision. Their own "best judgment" is the final arbiter of whether to accept the risks associated with an option.

Although one might balk at even the suggestion of letting technical experts make decisions about value issues, technicians are trained to be servants responsive to their clients' needs. If society as a whole is defined as the client, professional judgment may be the best way to devise creative and balanced solutions, considering what is desirable, feasible, and practical. When professionals deliberate, they may not only summarize existing knowledge, but also create new knowledge in the form of new and better options. A physician may finesse the question of whether a drug is safe enough for a patient who is sloppy about taking pills by devising a therapeutic regime that circumvents the problem; similarly, a safety engineer may alter traffic patterns so as to increase the effective safety of an aging bridge with fixed load-bearing capacity.

Professionals may stumble in some areas where formal methods are strongest. An inarticulable rule frustrates critics and colleagues attempting to assess the professional's performance and spot errors. Under the cloak of professional wisdom may lie only a vague notion of what options are available or even a failure to consider more than one traditional solution. Finally, there is no necessary link between expertise in a substantive area and expertise in decision making.

Similiarities and Contrasts

These three approaches are not as conceptually distinct as they might initially appear. Formal analyses require a large element of professional judgment, whereas professionals can (and at times do) base their judgments on formal analyses. Bootstrapping requires risk and benefit measurements resembling those in formal analysis; for their part, formal analyses often turn to the historical record for critical measures, making assumptions like those underlying bootstrapping. Professionals are often tradition oriented, attempting to do what has been done in terms of policy making; the past studied by bootstrapping has largely been created by the actions of professionals.

The difficulties the approaches face also have similarities. Characterizing a proposed technology for comparison with a historically derived standard encounters many of the same technical problems as characterizing it for comparison with alternative courses of action in a formal analysis. Both bootstrapping and professional judgment may falter by failing to consider alternatives. Furthermore, the prescriptive validity of each is contingent upon their descriptive validity. Professionals should be allowed to make acceptable-risk decisions only if they do know more; the cumulative record of evolutionary processes should be consulted for guidance only if such processes properly accommodate social pressures and realities. These correlated weaknesses may decrease the possibilities for hybridizing approaches to compensate for one another's vulnerabilities.

B. Possible Characteristics of Safety Goals

Mattson et al. (*op. cit.*) classify varieties of form of safety-goal formulation according to seven pairs of contrasting characteristics. The pairs, with brief discussion based on the work of Mattson and coworkers, are as follows:

1. Single vs. Multiple Goals

A single overall goal has the merit of simplicity. A set of more specific, narrower goals may, on the other hand, lend themselves better to verification of compliance. Advantageous combinations may be possible, consisting of an interrelated hierarchy of overall and lower-order goals.

2. Quantitative vs. Qualitative Goals

Overall quantitative goals introduce severe verifiability problems with respect to low-probability high-consequence events. On the other hand, general qualitatively stated goals may defy assignment of clear, ascertainable meaning. (Goals such "as safe as possible" or "as safe as reasonably possible" are included among the qualitative forms.) Different specific aspects of safety-goal formulation may be better suited to one or the other of these alternatives, or may be expressible in forms combining qualitative and quantitative elements. The authors see merit in hybrid, quantitative/qualitative goal forms.

3. Ends-Oriented vs. Process-Oriented Goals

Ordinarily, goals are thought of principally as ends to which strategies and resources are directed. However, uncertainties and difficulties in verifying goal attainment may be so substantial that it may be preferable to substitute processes that command confidence. The latter alternative is reflected in the basic reactor-safety principles of the Naval Nuclear Propulsion program, as recounted by Admiral Rickover in testimony before the President's Commission on the Accident

at Three Mile Island. Those principles included strong central technical control; management measures to achieve strong technical competence; conservatism of design; compliance with detailed operating procedures; prompt correction of safety deficiencies; and a host of do's and don'ts in the selection and training of plant operators and in the maintenance of formality and discipline in their performance. The implication of Admiral Rickover's comments is that, if the means of providing for nuclear safety are pursued with the goal of excellence of performance, then acceptable safety will be attained.

4. Absolute vs. Relative Goals

A goal may be expressed as some absolute level that risk should not exceed. Alternatively, it may be related to some other risks to which people are exposed (risks of other energy sources, other activities, natural phenomena). The relation may be expressed quantitatively as a ratio or percentage, or in non-quantitative terms (as equal to, less than, substantially less than, etc.). In any event, some comparative analysis must underlie relative goals, though comparisons may be fraught with considerable difficulties.

5. Individual vs. Society-Oriented Goals

In the NRC regulations (10 CFR Parts 20, 50/Appendix I, and 100), both individual and societal goal forms are used in complementary fashion for limiting radiation exposure. For example, Section C limits emission of radioiodine and radioactive particulates in terms of a maximum permissible individual exposure (15 millirem to any organ), while Section D of the same Appendix establishes an interim criterion of aggregated societal impact for cost-benefit trade-off purposes (\$1,000 per man-rem).

6. Site (or Region) Dependent vs. Site (or Region) Independent Goals

There are many variable locational factors which may affect: (i) the probability of a serious accident, (ii) the probability of a failure of a reactor containment, and (iii) the radiation exposure of people in event of an accident. Included are such factors as: population density and distribution; character of emergency evacuation routes and related locational features; natural phenomena such as meteorology, tornadoes, earthquakes, and floods; and proximity to man-made hazards such as airports, military installations, LNG terminals, and transport systems involving movements of hazardous chemicals or other explosive materials.

The issue of whether goals should vary with specific location includes the extent if any to which various possible site weaknesses may be compensated by design and operational features. The issue of possible regional variation of goals includes whether less exacting goals should be allowed in regions presenting region-wide difficulties (e.g., a relatively densely populated area of the country).

7. Time-Related vs. Atemporal Goals

A time-related goal may involve a schedule of time horizons for successive improvements in safety standards or goals. It may call for achievement of some specified safety standard by a specified year and a more stringent safety standard by some later year.

Within this temporal aspect is the question of interim vs. settled goals. A "settled" goal is not necessarily permanently established, but it has a low expectation of near-term change compared to an "interim" goal. However, even a "settled" goal may need to be changed, because of technological progress, changing social values, institutional changes, or the emergence of new information by which the uncertainty or some other factor in a goal is perceived.

Establishing a goal as "interim" may contribute to its acceptability; the interim goal would offer a trial period while the process of learning, technological progress, and the results of additional research, studies, and continued debate will serve to clarify desirable goal changes after the interim period has elapsed.

C. Types of Quantitative Approach

Recent efforts towards development of proposed safety goals have had a predominantly quantitative emphasis, notwithstanding evident problems of verification of compliance with any quantitative goals. The Advisory Committee on Reactor Safeguards (ACRS), as background for its own development of a quantitative approach, accepted a survey of some proposals previously presented in the United States and abroad, according to which such approaches can be roughly categorized into three groups: those that set limits on individual risk of death only; those that consider frequency of accidents and magnitude of the consequences; and those that imbed the criteria in risk management frameworks that, at least in part, consider risks from alternatives or other societal endeavors (NUREG-0736, Part 1). Some (though not all) of the criteria considered in the survey apply specifically to nuclear reactors.

1. Individual Risk Criteria

One of the early proposals for quantitative risk criteria for nuclear reactors was made by Adams and Stone (1967) of the Central Electricity Generating Board in Great Britain at an IAEA Symposium on Siting and Containment. They proposed that the parameter determining acceptable siting be taken as individual risk. They suggested that an incremental increase in an individual's chance of death per year that is smaller than the demographic variation in the United Kingdom of that chance would be inappreciable and acceptable on those grounds. Differences significantly greater than 10^{-5} per year occur among England, Wales, Scotland, and Ireland, and they proposed that an incremental individual

risk of 10^{-5} chance of death per year would be acceptable. For immediate deaths and a plant lifetime of 30 years, this risk would correspond to a statistical loss of life expectancy of about 6 days, while for death delayed until 10 years after exposure the statistical loss is about 3 days. Of course, the loss is much larger for the actual victims and zero for all the others.

The apparently positive correlation between standard of living and health has been used by Bowen (1975) to develop a general risk acceptance criterion for technological activities in the United Kingdom. He suggests that the risks imposed upon society should be negligible or balanced by benefits. However, risk levels that can be scientifically supported, say, a 10^{-5} chance of death per year, cannot be considered negligible in all situations, and balancing by direct individual benefits is not possible in cases where the victim cannot be readily identified in advance, for example, the one excess cancer fatality that might be expected from the TMI accident. Bowen argues that the balance should be done macroscopically.

He assumes that the observed annual increase in life expectancy in the U.K. is due to overall societal efforts, i.e., its investment in "the industrial machine" of which any technological facility forms a part. An additional yearly risk of death of 10^{-5} from a new facility roughly balances the expected increase of an individual's life expectancy during one year. With regard to accidents having a potential for a major disaster, Bowen argued that the 10^{-5} limit should be demonstrated to a high confidence level.

2. Frequency-Consequence Approaches

The previous criteria dealt specifically with individual fatality risks without directly including limits on other types of risk or addressing the effects of a large scale accident. In the frequency-consequence type of approach, special attention is given to the magnitude of an accident. A basic common assumption in variants of this approach is that the limiting frequency of a particular accident should depend in some way upon its magnitude. One early British proposal (Farmer, 1967) suggests a limit on the frequency of accidental release of radioactive material; a more recent Canadian one (Atomic Energy Control Board of Canada, 1978) on frequency of individual exposure; and a recent British proposal (Kinchin, 1978 and 1979) is concerned with limits on the fatalities due to accidental exposure.

In an early paper ("Siting Criteria - A New Approach," IAEA Symposium on Siting and Containment, 1967), F. R. Farmer proposed that probabilistic analysis be employed in reactor safety assessment and suggested that the safety criterion of less than 0.01 premature deaths per reactor year be adopted. In addition, he proposed that a risk acceptance limit line relating release magnitude to frequency be used to judge the acceptability of the estimated occurrence frequency for any particular accident. The severity of the accident was measured

by the release in curies of iodine-131, one of the volatile fission products thought to be of greatest importance in thermal reactor accidents.

The Farmer limit line does not deal specifically with effects dependent upon population density and other conditions around the site. Therefore, the actual limits on effects, such as risk to individuals, property damage, or number of expected fatalities, must be estimated from site-specific analyses.

3. Risk Management Approaches

Two common premises of risk management approaches are: that society has a limited amount of resources to allocate for the reduction of the risks that accompany the benefits of its endeavors and that these resources should be allocated wisely. Such approaches reflect concern that improper actions to reduce risks may not minimize risk and may even give rise to an increase in overall risks.

Okrent and Whipple (1977) described a quantitative approach to risk management which incorporated the following principal features:

Risk assessment. Each risk-producing facility, technology, and the like would have to undergo assessment of risk both to the individual and to society.

Graduated limits on individual risk. Societal activities would be divided into major facilities or technologies, all or part of which are categorized as essential, beneficial, or peripheral to society. There would be a decreasing level of acceptable risk to the most exposed individual (for example, 2×10^{-4} additional risk of death per year for the essential category, 2×10^{-5} for the beneficial category, and 2×10^{-6} for the peripheral category).

Internalization of residual risk costs. To provide incentive to reduce risk and balance some inequities between those who receive the benefits and those who are burdened by risk, the cost of the residual risk would have to be internalized, perhaps by means of a tax paid to the Federal government.

Modest risk aversion. Risk aversion to large events would be built into the internalization of the cost of risk, but with a relatively modest penalty.

Cost effective reduction of residual risk. A limit on the marginal cost of risk reduction could be imposed.

The late C. L. Comar's editorial for Science (1979) entitled "Risk: A Pragmatic De Minimis Approach" suggested the following as potentially useful, though oversimplified guidelines to perspective in risk policy rationalization:

- (a) Eliminate any risk that carries no benefit or is easily avoided.
- (b) Eliminate any large risk (about 1 in 10,000 per year or greater) that does not carry clearly overriding benefits.
- (c) Ignore for the time being any small risk (about 1 in 100,000 per year or less) that does not fall into category (a).
- (d) Actively study risks falling between these limits, with the view that the risk of taking any proposed action should be weighed against the risk of not taking action.

Risk management considerations are involved in the ACRS proposal and other recent U.S. proposals, addressed below (in Section IV).

D. Approaches to Dealing with Uncertainty

Verifiability of whether a nuclear plant meets an overall quantitative safety goal is subject to both reducible and irreducible uncertainties surrounding probabilities and consequences of severe but low-probability accidents. Various approaches can be considered for mitigating the impact of those uncertainties. We describe five approaches: arbitration, modeling, restrained use of quantitative criteria, conservatism, and non-quantitative approaches. Combinations of approaches are, of course, possible and may well have merit.

1. Arbitration

The decisions about whether a goal is met can be entrusted to a group of specially empowered experts. (We avoid the term "science court.") This group could provide a mechanism for clear, authoritative resolution of disputes involving data gaps and possible conflicting interpretations. Selection of the arbiters would be guided by the aim of high quality in their decisions. However, there is here some inherent danger to the rationality of the regulatory results, since decisions would be made by methods that would necessarily in part transcend the scientific method. Predictability could also suffer, as experts and their opinions change.

2. Modeling

Verification could be tied to a prescribed mode of calculation. (For example, "Consider the aggregated probability to be the sum of the probabilities of the ten dominant sequences described below.") The approach would aim at consistency and predictability. With such approaches the quality of the results depends heavily on the fidelity of the calculational model; one needs to consider whether there is, in the model or its use context, enough respect for uncertainties (including possibilities of crucial oversight) to make its dependability commensurate with the consequences of misjudgment.

3. Restrained Use of Quantitative Criteria

Use of the top-level probability target can be limited to decisions concerning outliers, the clearly prohibitive and the clearly trivial. By leaving many decisions outside its purview, this approach detracts from the policy's value on the criterion of reach. But its use would be focused on those parts of the safety-goal issue for which it is most dependable.

Another mode of restraint is to use an overall plant target only loosely and to provide some general guidance for setting quantitative targets at the level of systems and equipment for which failures can be tolerated with a frequency sufficient for possible verification (say, once per 10 or 100 reactor-years). Such an approach, while permitting verification, is subject to the effects of unavoidable gaps in allocation (where measurable failure frequencies cannot be accepted) and of common-cause and interactive failures.

4. Conservatism

Conservatism can be embedded in the calculations or explicitly applied to risk estimate results. In the "embedded" approach, estimates of actual risk are based on conservative assumptions, to increase assurance that decisions concerning the issue at hand will be on the safe side. One could, for example, say that estimated risk should be ensured at some specified high confidence level for each critical element (say, 90 percent of 99 percent confidence, rather than a 50 percent confidence-level "best estimate").

The conservatism-in-calculations approach is subject to difficulties of two sorts. The first is that the most significant and troubling uncertainties often do not lend themselves to the making of estimates to which particular confidence levels can reasonably be assigned, and, in any event, some residual uncertainty remains.

The second difficulty is that, with conservatism introduced into the elements of an overall estimate, the degree of residual uncertainty in the result and the overall extent to which it includes conservativeness margins are not readily apparent to the decision maker. This handicaps trade-offs against side effects, not only economic ones, but also those of adverse safety impacts elsewhere of a safety feature introduced to solve a particular problem.

As an alternative to the "embedded" conservativeness approach, one can require calculations to be realistic -- the best one can do. Margins of due prudence would be introduced into administrative decisions. This explicit approach to conservatism offers visibility of margins, for consideration and debate. The decision maker still needs, however, a "feel" for the uncertainties entailed by the assumptions, facts, and methods of the original risk estimation. Sensitivity analyses may help discern the consequences of alternative judgments and thus help to set an appropriate margin.

5. Nonquantitative Approaches

Nonquantitative approaches available to reduce vulnerability due to uncertainties include the following:

- Independent lines of defense. For example, remote siting no matter what degree of plant safety the calculations show.
- Use of requirements known to be beneficial, though not quantifiable. For example, emphasis on qualifications and training of operators, other plant personnel, and managers. Economic incentives for safety (or removal of disincentives) could be in this category.
- Extra circumspection when abandoning known, established practices for promising innovations.
- Professional judgment, with emphasis on the qualifications and experience of decision makers and advisors.

E. Considerations in Balancing of Values

Much of what has preceded in this paper has reflected consideration of disparate and sometimes conflicting values involved in development of a safety goal -- ethical, political, and economic values, along with technical considerations.

As the work of safety-goal development progresses, we plan to address the philosophically complex and demanding task of seeking a sound balancing of the values relevant to goal setting. The scope of this section, however, is limited to noting in a preliminary way two specific topics: (a) individual and social risks and (b) safety and economic values.

1. Individual and Social Risks

A safety goal may focus on the risk of death or illness to which any one individual may be exposed, including permissible differences among individuals in different circumstances such as workers, nearby residents who may receive some benefit as compensation for incremental risk, and the general population.

The goal may also include consideration of the numbers of people exposed to a risk. An individual living near a plant in a densely populated area may not be exposed to greater individual risk than a person living near a plant remote from population centers, but the total number of probabilistically "expected" casualties (i.e., the aggregated social risk) would be greater.

Social risk considerations also include such issues as equities of distribution of risks and benefits, possible transfer of risk to future generations (through long-lived radioactive materials),

distinctions between prompt and delayed fatalities, and whether and how property damage risks (which may indirectly entail other health-and-safety risks) should be taken into account.

A special aspect of the social risk issue is involved in the concept of risk management, that is, whether and how to consider the risks of alternatives in setting a nuclear-plant safety goal. A strict safety goal for nuclear power may result in some potential nuclear plants not being built and thereby lead to the safety risks imposed by alternative power sources (perhaps coal) or resulting from deprivation.

2. Safety and Economic Values

Safety and economic values are not necessarily or invariably in conflict. Safety may protect plant investment and property of others, as well as save medical and other costs of illness. In this positive respect, the issue is the extent to which economic incentives for safety should be reflected in safety-goal formulation.

However, in most questions about safety, safety and economic values conflict or appear to conflict; the issue involves, first, whether, under the circumstances at hand, economic impacts can be taken into account in a safety decision, and second, what amount of resources may be devoted (or what benefits foregone) to achieve a particular safety goal. A possible approach involves setting some minimum acceptable safety level that cannot be compromised for economic reasons, but taking economic factors into account in seeking a higher safety goal. The safety-cost trade-offs involved may be aided by guidelines about what degrees of safety should be viewed as equivalent to what costs (e.g., \$1,000 per man-rem probabilistically expected), notwithstanding problems in applying a common metric to values that may not be philosophically commensurate. Special difficulties in such trade-offs include methods of reflecting wide uncertainties in safety consequences in evaluations relative to generally better known economic impacts.

Decisions concerning possible retroactive applications of new requirements may involve special safety-cost trade-offs since the costs of specific safety improvements for plants already being built or operating may be substantially higher than for future plants. By contrast, safety benefits are likely to be similar for old plants, and perhaps less than for future plants, in view of lesser remaining operating life and possible side effects of modifications of existing systems.

IV. SOME SAFETY-GOAL PROPOSALS

Six frameworks are currently used by society to regulate risks. The first is market regulation. Until recently, caveat emptor (let the buyer beware) was the mechanism for making safety decisions. The second is no-risk. In the Delaney Clause to the Food, Drug and Cosmetic Act, Congress instructed

the FDA to allow no carcinogens to be added to the food supply. In attempting to implement these instructions, FDA has attempted to define a negligible level of risk to make the absolute nature of the Delaney Clause more flexible. Most of the proposals that follow attempt to define a negligible risk by using various comparisons. The third is technology-based standards, such as best available control technology. This framework is an attempt to shift attention from risks and costs over to available technology. The fourth is risk-benefit analysis. This framework attempts to trade off increased benefits against possible increases in risk to achieve a better social solution. The fifth framework is cost-effectiveness analysis. This framework is used a great deal in regulating the nuclear fuel cycle to ensure that emissions requirements are equally stringent throughout the cycle. ALARA and ALAP are variants of this framework. The sixth framework is benefit-cost analysis. This framework seeks an explicit quantification of all major benefits and costs of some decision and a formal balancing.

Each of these frameworks is a possible candidate for future regulation of the nuclear fuel cycle, along with new proposals. Any proposal must be evaluated in terms of the criteria set out in a previous section. The following proposals for quantitative safety goals are set out as current proposals known to the NRC. We hope and expect to see new proposals formulated as a result of this request for comment. We also hope to receive evaluations of the extent to which these proposals fulfill the criteria set out in this document, as well as other criteria which are relevant.

This section contains six safety-goal proposals, all urging a quantitative formulation of safety goals and suggesting specific numerical values for the parameters that they propose. The diverse frameworks (the ACRS "goal levels" and "upper limits," the AIF criteria, the Joksimovic "limit line," etc.) can, however, be employed in conjunction with numerical values that differ from those suggested in the specific proposals. The intellectual task of devising a sound framework and the policy task of setting appropriate levels are distinct, though they must in due course merge. Both tasks are essential, as is their proper merging. All six of the proposals are based on use of probabilistic risk assessment. They depend for their effectiveness on the potentialities of that technique and are subject to its limitations -- gaps and uncertainties in data bases, difficulties of constructing adequate sets of underlying assumptions, and of accounting for human factors in safety. Some of the approaches proposed include provision for accommodating or mitigating the effects of uncertainties.

With all the approaches, consideration would have to be given to (a) the appropriate role for other methods where probabilistic risk assessment cannot dependably resolve an issue and (b) problems of transition from current practices, including decisions with respect to possible retroactive application of new requirements. A transitional problem of a different sort is the development of sufficiently wide availability of an adequate level of expertise in the new probabilistic methods.

A. The ACRS Proposal

On October 31, 1980, the Advisory Committee on Reactor Safeguards submitted to the Commission what it described as "a preliminary proposal for a possible approach to quantitative safety goals ... intended to serve as one focus for discussion on the subject ... and ... expected to be ... a first step in an iterative process." (NUREG-0739, "An Approach to Quantitative Safety Goals for Nuclear Power Plants," ACRS, 1980.) The ACRS letter containing the Committee's proposal was accompanied by a three-part report prepared by its Subcommittee on Reliability and Probabilistic Assessment, discussing aspects of the subject in detail. The ACRS letter (without the background report) is presented below.



- 32 -
UNITED STATES
NUCLEAR REGULATORY COMMISSION
ADVISORY COMMITTEE ON REACTOR SAFEGUARDS
WASHINGTON, D. C. 20555

October 31, 1980

Honorable John F. Ahearne
Chairman
U.S. Nuclear Regulatory Commission
Washington, D.C. 20555

SUBJECT: AN APPROACH TO QUANTITATIVE SAFETY GOALS FOR NUCLEAR POWER PLANTS

Dear Dr. Ahearne:

In a letter dated May 16, 1979, the ACRS recommended that consideration be given by the Nuclear Regulatory Commission to the establishment of quantitative safety goals for nuclear power reactors. The ACRS acknowledged the difficulties and uncertainties in the quantification of risk but stated its belief that quantitative safety goals and criteria can provide an important yardstick for the engineering judgment that would still be required. The ACRS further recommended that the Congress be asked to express its views on the suitability of proposed NRC quantitative safety goals and criteria in relation to other relevant aspects of our technological society.

In a letter dated June 11, 1979 to the ACRS you noted that you would appreciate any further development of the concept of quantitative safety goals that the ACRS could provide. In a memorandum dated August 14, 1979 the ACRS advised you that it was assigning the project of developing a possible approach to quantitative safety goals to its Subcommittee on Reliability and Probabilistic Assessment and that it was anticipated that about a year would be needed to develop recommendations.

The Subcommittee has now developed a preliminary proposal for a possible approach to quantitative safety goals. The proposed approach is intended to serve as one focus for discussion on the subject of quantitative safety goals and as such is expected to be only a first step in an iterative process.

The Subcommittee has prepared its discussion of the subject in the form of a report which consists of three parts, as follows:

- Part 1 "On the Development of Quantitative Risk Acceptance Criteria,"
J. M. Griesmeyer and D. Okrent.
- Part 2 "Risk Management and Decision Rules for Light Water Reactors,"
J. M. Griesmeyer and D. Okrent.
- Part 3 "Applications and Implications of Trial Risk Acceptance Criteria,"
D. H. Johnson and W. E. Kastenber.

Honorable John F. Ahearne

-2-

October 31, 1980

Part 1 is primarily a review of several prior or current proposals for quantitative risk criteria which have been developed by others. Part 2 provides the preliminary proposal for a possible approach to quantitative safety goals. Part 3 provides a brief evaluation of several technologies, including nuclear power plants, in terms of criteria like those proposed in the report.

The ACRS recognizes that there are several other ongoing efforts to examine the development of such criteria. The Committee hopes that this report will contribute material useful in the process of developing an approach.

The trial approach to quantitative safety criteria, which is described in Part 2 of the report, is divided into two major tasks: the predominantly social and political task of setting the safety criteria (termed decision rules herein) and the technical task of estimating the risks and deciding whether the safety criteria have been met.

The safety criteria or decision rules are as follows:

- Limits are placed on the frequency of occurrence of certain hazardous conditions (hazard states) within the reactor.
- Limits are placed on the risk to the individual of early death, or delayed death due to cancer arising from an accident.
- Limits are placed on the overall societal risk of early or delayed death.
- An "as low as reasonably achievable" approach is applied with a cost-effectiveness criterion that includes both economic costs and a monetary value of preventing premature death.
- A small element of risk aversion is applied to infrequent accidents involving large numbers of early deaths compared to a similar number of deaths caused by many accidents each involving one or two deaths.

Each decision rule on hazard states and on individual and societal risk consists of a pair of numbers: an upper, non-acceptance limit on risk and a lower, safety goal level of risk. Compliance with the upper limit would be required for extended operation of the plant; otherwise, it must be improved within a certain period of time (to be determined) that depends upon the severity of the risk involved. On the other hand, any risk value lower than the safety goal level would be considered in compliance for the particular category of risk. However, risks must be further reduced below these safety-goal levels whenever improvements are possible that meet certain cost effectiveness criteria for risk reduction. Between the upper, non-acceptance

Honorable John F. Ahearne

-3-

October 31, 1980

limit and the lower safety-goal level of risk, there is a digressionary range in which case by case consideration of uncertainties, regional need for power, and alternative risks is required in the decision as to whether the plant should be allowed to operate for an extended time without modification.

The preliminary numerical values which have been suggested for use in the decision rules are primarily a matter of judgment and are intended to help stimulate discussion and evaluation in concrete terms.

Ultimately the NRC and the Congress must consider a wide range of socio-political and economic factors, of which direct risk to the public health and safety is but one, in arriving at a judgment on suitable risk acceptance levels.

The quantitative values suggested for use in the proposed decision rules are intended to be applicable for new light water power reactors and may be more stringent than is deemed appropriate for existing plants.

DECISION RULES

HAZARD STATES

Accidents that damage the facility represent possible forerunners of more severe accidents. A tentative set of hazard states of progressive severity has been defined and a preliminary set of limits on their rate of occurrence has been proposed, as is shown in Table 1. The limit on the frequency of a large offsite release, assuming that a fuel melt has occurred, places emphasis on mitigation as well as prevention of serious accidents. Such a division between accident prevention and accident mitigation is believed to be necessary because of the difficulty in demonstrating with a very high degree of confidence that a frequency of large scale fuel melt much less than the proposed goal of 10^{-4} per reactor-year can be achieved in view of the complexities introduced by consideration of matters such as sabotage, earthquakes, and other potential multiple failure scenarios.

INDIVIDUAL RISKS

The limits on risk to individuals living closest to the reactor site have been set well below the sum of all other risks for any age group and below those from the principal competing source of generating electricity. Lower levels were chosen (by a factor of five) for the risk of early death than for delayed death from cancer many years after an accident.

Table 2 summarizes the proposed decision rules for risks of delayed death from cancer and of early death. Note that relatively few people will have risks as high as the most exposed individuals who presumably reside close to the plant site boundaries. Most people will be exposed to risks lower than the goal levels.

Honorable John F. Ahearne

-4-

October 31, 1980

SOCIETAL RISK

It has been suggested in the literature that society is risk averse when comparing a single, infrequent large accident with a number of small accidents leading to the same total number of fatalities in the same time period. A simple approach which assesses an equivalent social cost that increases faster than the actual consequences for events involving multiple deaths uses an equation of the form

$$\text{Equivalent social cost} = \sum_{\text{accidents}} (\text{Frequency})(\text{Consequence})^{\alpha}$$

in which α is greater than unity. If α is equal to one, the equivalent social cost would be the same as the expected costs (frequency times consequence). Although values of α as high as 2 or 3 have been proposed in the literature for fatalities from accidents, such values would prohibit many existing technological endeavors because of extremely high equivalent social cost, e.g., dams or large quantities of hazardous chemicals stored close to population centers. Studies performed by the Subcommittee and summarized in Part 3 of the report, indicate that society does not consciously place such high risk aversion penalties on needed activities.

In this proposal it is suggested that the social cost for delayed cancer deaths should be assessed as equal to the calculated number of fatalities (i.e., $\alpha = 1$). The range on the estimated number of people who die from the pollution arising from a coal-fired plant which generates 10^{10} kWh is about 10 to 200 (see Part 3 of the report); 10 is proposed here as the upper, non-acceptance limit on the delayed cancer deaths due to a nuclear power plant; the goal level is that there be less than two cancer fatalities per 10^{10} kWh.

To provide incentives to reduce the catastrophic potential of accidents, it is proposed to assess the equivalent social cost of early deaths with a value of α slightly larger than unity, namely $\alpha = 1.2$; hence, the equivalent early death cost of the plant, E_{ed} , would take the form

$$E_{ed} = \sum_{\text{accidents}} (\text{Frequency})(\text{Early Deaths})^{1.2}$$

The limits on equivalent early deaths are reduced by the same factor of five from the delayed death limits as was done for the limits on individual risk. Table 3 summarizes the decision rules for societal health risks.

SOCIETAL IMPACT REDUCTION - ALARA

It is proposed to use an "as low as reasonably achievable" (ALARA) cost-effectiveness criterion to judge whether additional risk reduction is required beyond

Honorable John F. Ahearne

-5-

October 31, 1980

that level of safety required to meet the other decision rules. The cost of an improvement would be balanced against the combined change in economic losses and in the risk of delayed cancer deaths, and equivalent early deaths.

While there is some limit on how much the United States can afford to spend to reduce risk from all of its technological activities, lest economic instability lead to greater risk directly or indirectly, the current perspective on nuclear reactors may be such that society is willing to spend more for LWR safety than for many other things.

It is tentatively proposed that the marginal cost limit on expenditures be set at \$1 million per delayed cancer death averted and \$5 million per equivalent early death averted, when "equivalent" early deaths are calculated using the coefficient $\alpha = 1.2$ for risk aversion.

It is anticipated that careful study will be required to quantify the economic losses due to property and resource damage. Because of uncertainties and the fact that some impacts cannot be quantified, it is proposed that the marginal cost limit on expenditures to reduce adverse economic impacts be twice the expected reduction in impact when applying the ALARA criterion. This also stresses prevention rather than repair of possible damage.

Table 4 summarizes the quantified ALARA criterion.

RISK QUANTIFICATION

The rest of the proposed framework deals with the technical tasks of risk quantification, which will by no means be simple. It has to be acknowledged from the beginning that there will be both large uncertainties in such risk estimates and significant differences between independent estimates of the same risk. The form of the decision rules is intended to compensate in part for some of this uncertainty. Limits are placed on the expected values of the various risks. These expected values are the weighted average of the probabilities and therefore reflect some of the uncertainties. Also, limits are placed on both the risk of a damaging accident to the fuel and on the risk of a large release of radioactive material assuming the occurrence of fuel damage, thereby requiring both prevention and mitigation.

A major tool for this effort will be a plant and site specific quantitative risk analysis which is essentially a probabilistic estimate of the distribution of risks. The details of the analysis will form a safety profile of the particular plant and site that can be used to make risk-based decisions on design and/or procedural changes. The estimated risk distribution will explicitly express the range of uncertainties and will be used in the application

Honorable John F. Ahearne

-6-

October 31, 1980

of the decision rules. Special attention must be given to quality assurance of the risk assessment. There must be full and explicit identification of the assumptions and limitations of the analysis, and peer review will be required. In addition, it is proposed that a procedure be established to provide a legally binding determination of those risk distribution values to be used with the decision rules. A possible approach to this aspect is the establishment of a Risk Certification Panel. After peer review of the analyses had been completed, the panel would be given the statutory authority to make a finding on the risk values to be used in the application of the decision rules.

The ACRS hopes this report proves to be useful in the ongoing effort on the development of quantitative safety goals. The Committee plans to continue to pursue the matter actively.

Sincerely,



Milton S. Plesset
Chairman

Table 1. Limits on Occurrence of Hazard States

Hazard State	Probability Goal	Decision Rules on Mean Frequency	
		Goal Level	Upper Limit
Significant Core Damage (> 10% of noble gas inventory leaking into primary coolant)	Less than 1/100 per reactor lifetime	$f_{cd} < 3 \times 10^{-4}$ per reactor year	$f_{cd} < 1 \times 10^{-3}$ per reactor year
Large Scale Fuel Melt - LSPM (> 30% of oxide fuel becoming molten)	Less than 1/300 per reactor lifetime	$f_m < 1 \times 10^{-4}$ per reactor year	$f_m < 5 \times 10^{-4}$ per reactor year
Large Scale Uncontrolled Release from Containment given LSPM (> 10% of Iodine inventory and 90% of noble gas)	Small, given a Large Scale Fuel Melt	$f_{R/m} < 0.01$ per LSPM	$f_{R/m} < 0.1$ per LSPM

f_{cd} is the frequency of Significant Core Damage per reactor year.

f_m is the frequency of Large Scale Fuel Melt per reactor year.

$f_{R/m}$ is the frequency of Large Scale Uncontrolled Release per Large Scale Fuel Melt.

The upper non-acceptance limits must be satisfied for extended operation of a new plant or for issuance of a construction permit. Between the upper limits and the goal levels is a discretionary range for case-by-case consideration of uncertainties and competing risk. Once the risk level decision rules have been applied, risk must still be reduced if such reduction is reasonably achievable within the cost-effectiveness criterion of Table 4.

Table 2. Limits on Risks to Most Exposed Individual

Probability Goal	Mean Frequency per Site-year		Decision Rules on Mean Frequency per Large Scale Fuel Melt-LSFM	
	Goal Level	Upper Limit	Goal Level	Upper Limit
Probability of delayed death from cancer due to all reactors at a site over lifetime of individual <0.0005	$f_d < 5 \times 10^{-6}$ per site-year	$f_d < 2.5 \times 10^{-5}$ per site-year	$f_{d/m} < 0.01$ per LSFM	$f_{d/m} < 0.05$ per LSFM
Probability of early death due to a reactor accident over lifetime of individual < 0.0001	$f_{ed} < 1 \times 10^{-6}$ per site-year	$f_{ed} < 5 \times 10^{-6}$ per site-year	$f_{ed/m} < 0.002$ per LSFM	$f_{ed/m} < 0.01$ per LSFM

f_d is the individual risk of delayed cancer death per site year.

$f_{d/m}$ is the individual risk of delayed cancer death per large scale fuel melt.

f_{ed} is the individual risk of early death per site year.

$f_{ed/m}$ is the individual risk of early death per large scale fuel melt.

The upper non-acceptance limits must be satisfied for extended operation of a new plant or for issuance of a construction permit. Between the upper limits and the goal levels is a discretionary range for case by case consideration of uncertainties and competing risks. Once the risk level decision rules have been applied, risk must still be reduced if such reduction is reasonably achievable within the cost-effectiveness criteria of Table 4.

Table 3. Societal Health Risk Limits

Measure of Risk	Decision Rules on Societal Health Risks	
	Goal Level	Upper Non-Acceptance Limit
E_d = the expected value of: \sum (Frequency) (Delayed Cancer Deaths) accidents and normal operation	$E_d < 2$ per 10^{10} kWh	$E_d < 10$ per 10^{10} kWh
E_{ed} = the expected value of: \sum (Frequency) (Early Deaths) ^{1.2} accidents	$E_{ed} < 0.4$ per 10^{10} kWh	$E_{ed} < 2$ per 10^{10} kWh

E_d is the average number of delayed cancer deaths per 10^{10} kWh of electricity generated.

E_{ed} is the average number of equivalent early deaths per 10^{10} kWh of electricity generated.

10^{10} kWh is the amount of electricity generated by a large (1200 MWe) power plant operating at full capacity for one year.

The upper non-acceptance limits must be met for extended operation of a new plant or for issuance of a construction permit. Between the upper limits and the goal levels is a discretionary range for case by case consideration of uncertainties and competing risk. Once the risk level decision rules have been applied, risk must still be reduced if such reduction is reasonably achievable within the cost-effectiveness criteria of Table 4.

Table 4 Quantified ALARA Cost-Effectiveness Criteria.

Expenditure Limits for Impact Reduction	
\$1 million per delayed cancer death averted	$\$1 \times 10^6 / (\Delta E_d L)$
\$5 million per early equivalent death averted	$\$5 \times 10^6 / (\Delta E_{ed} L)$
2 times the economic loss (due to resource damage) averted	$2 / (\Delta E_r L)$
<p>A particular improvement is "cost-effective" and required if</p> $\text{Cost} \leq [2\Delta E_r + (\$5 \times 10^6) (\Delta E_{ed}) + (\$1 \times 10^6) (\Delta E_d)] L$	

ΔE_d is the change (due to the proposed improvements) in the expected value of:

$$\sum_{\substack{\text{accidents} \\ \text{and} \\ \text{normal} \\ \text{operation}}} (\text{Frequency}) (\text{Delayed Cancer Deaths})$$

ΔE_{ed} is the change (due to the proposed improvements) in the expected value of:

$$\sum_{\text{accidents}} (\text{Frequency}) (\text{Early Deaths})^{1.2}$$

ΔE_r is the change (due to the proposed improvements) in the expected value of:

$$\sum_{\text{accidents}} (\text{Frequency}) (\text{Economic Losses})$$

L is the remaining lifetime of the plant in units of 10^{10} kWh to be generated and the frequencies are calculated per 10^{10} kWh. This is the amount of electricity generated by a large (1200 MWe) plant operating at full capacity for one year.

B. Other Proposals

The ACRS proposal is the only safety-goal proposal that has been formally presented to the Commission for consideration or discussion. However, we are aware of other proposals that have been elaborated in varying degree and have been published or widely communicated. Five such proposals are discussed briefly below. All five, like the ACRS proposal, involve quantitative safety goals and use of probabilistic risk assessment. (The proposals are on file in the Public Document Room, Nuclear Regulatory Commission, 1717 H Street, N.W., Washington, D.C. 20555. Copies may be purchased from the Public Document Room.)

1. Atomic Industrial Forum

The Atomic Industrial Forum (AIF) Subcommittee on Probabilistic Risk Assessment presented an approach to the ACRS Subcommittee on Reliability and Risk Assessment on July 1, 1980. The AIF proposal involves quantitative safety goals addressed in four elements: individual health effects, population health effects, cost-benefit ratio, and core degradation probability.

The proposed individual health effects criterion is that the incremental risk of adverse health effects to the maximally exposed individual in the vicinity of a nuclear plant site should not result in a significant increase in annual mortality risk or in significant shortening of expected statistical life span. This is interpreted in terms of a suggested goal of 10^{-5} /year individual mortality risk (mean value).

The proposed population health effects criterion is that the incremental cumulative risk of adverse health effects to the exposed population per 1,000 MW(e) of nuclear power capacity, considering the probability and consequences of events integrated over the spectrum of potential accidents, should be no more than a small fraction of the average background incidence of health effects. This is reflected in a suggested goal of 0.1 fatality/year per 1,000 MW(e) (mean value).

The AIF subcommittee suggests that for improvements beyond the safety goal and for possible exemptions from it decisions should be guided by cost-benefit analyses. A numerical criterion of \$100/man-rem is suggested. The criterion is ascribed to the principle that the benefit, in terms of population risk reduction, afforded by a change in plant design or operating procedure should be comparable to that which is generally achievable through alternative investments of the cost of the change in other areas of public risk reduction. The numerical criterion (which is lower than the \$1,000/man-rem chosen by the NRC to guide interpretation of the as-low-as-reasonably-achievable criterion in radiation protection) is said to be equivalent to \$1 million per life saved and to be comparable to median cost-benefit ratios for other health and safety protective measures.

As the core degradation probability criterion, the AIF group proposes that a limit should be established for the probability of accidents involving serious core degradation so that, for an expected population of reactors, the recurrence interval for accidents as serious as Three Mile Island is on the order of one per several decades. They argue that a core degradation criterion is needed, as a supplement to the direct public radiation-impact criteria, to establish minimum requirements for accident prevention, prevent undue emphasis on mitigation, reduce frequency of stress-provoking events for populations near plants, and limit the economic risks of accidents.

The AIF group calls for use of probabilistic risk assessment (PRA) to support deterministic requirements. PRA should be used in interpretation of the goals in terms of generic requirements, but not as a licensing condition for individual plants. A common PRA methodology should be developed.

The AIF proposes three cautions about quantitative safety goals:

- The suggested numerical values should be used with mean value (50 percent confidence level) estimates of risk. Higher values would be appropriate for more conservative (high confidence level) estimates of risk.
- The initial set of values should be interim, for trial use for a period of three years.
- Qualitative judgment must supplement quantitative goals. This is particularly important in borderline cases.

2. Starr

Dr. Chauncey Starr, of Electric Power Research Institute, in a paper presented before the American Nuclear Society (November 1980), advocated a safety-goal approach that would build heavily on electric utility companies' economic self-interest in safety. He argues that the economic requirement for protection of the investment in the plant may often be a more demanding safety constraint than social acceptability based on hazard to the public. He presents the view that economic self-interest would motivate utilities to work towards a probabilistic design target that is 10,000 times stricter than a regionally acceptable upper bound for risk to the public. (The nearby public's risk acceptance would be influenced by benefits; Dr. Starr would compensate nearby residents for the small risk increment by reduced electric rates.) The 10,000-fold "protective range" would provide a margin for uncertainties in probabilistic estimates.

Dr. Starr's proposal includes recognition that the licensees' self-motivated safety goals would need to be supplemented by NRC requirements, in two respects: NRC should (a) establish design criteria

for systems that protect the public but not the plant (remote siting and reactor containment) and (b) check that the industry acts in accordance with its own financial interest (by requiring good engineering and management practices in siting, design, construction, and operation).

3. Joksimovic and O'Donnell

Dr. Vojin Joksimovic and L. F. O'Donnell, of General Atomic, in a paper entitled "Quantitative Safety Goals for the Regulatory Process (October 1980)," present a "limit-line" approach to safety goals.

The limit lines are plots of number of fatalities or public property damage as a function of the annual probability with which such effects may acceptably occur.

Two alternative lines are presented for latent cancer fatalities: one, reflecting a "balanced risk policy," would establish a 10^{-4} /reactor-year probability limit on accidents causing 100 casualties; the other, based on an "emphasized risk policy," would limit casualties at that probability level to 0.1. Both limit lines would impose lower consequence-probability products for high-consequence accidents (the increased emphasis being reflected by an exponent of 1.5 on the number of fatalities in an accident).

4. Zebroski

Dr. Edwin Zebroski, Director of the Nuclear Safety Analysis Center at the Electric Power Research Institute, in a paper entitled "A Proposed National Nuclear Safety Goal" (March 1980), described his proposal as "a tentative and probably incomplete formulation" and as a personal (rather than an EPRI) proposal. Dr. Zebroski stated the proposed goal as follows:

- (a) Reactor design, operation, and regulation should insure that accidents which reach the stage of core melting have a probable frequency of no more than one such occurrence in 30 years, taking account of the actual population of civilian reactors in operation in the U.S.
- (b) Reactor safety systems and containments shall be maintained and operated so that even if core meltdown were to occur, there would be less than one chance in 1,000 that a radiation release leading to a dose of 1 R or more to any member of the public will result.

5. Bernero

The NRC and others have underway a number of probabilistic plant analyses which are producing estimates of the probability of significant accident sequences (severe core damage or core melt sequences).

Robert Bernero, Director of NRC's Division of Systems and Reliability Research, has suggested the use of interim criteria and priorities for taking corrective action pending formal development of numerical criteria. (Presentation before ACRS Subcommittee on Reliability and Probabilistic Assessment, July 1, 1980.) The interim criteria suggested by Mr. Bernero are as follows:

<u>Estimated Probability of Severe Core Damage, Per Year</u>	<u>Action</u>
Greater than 10^{-2}	Correct in days
10^{-2} to 10^{-3}	Correct in months
10^{-4} to 10^{-5}	Correct in years
Less than 10^{-5}	No action

Users of these interim criteria are urged to weigh possible bias in the analysis, the quality of the analysis, the potential scale of consequences, and other significant factors in applying them.

V. DEGREE OF SAFETY

The essence of a safety goal is the degree of safety that it establishes. The generic discussion of approaches to safety-goal formulation and of the criteria by which approaches may be judged has as its purpose the creation of a vehicle for specifying the degree of safety sought in a way that would make it meaningful and useful.

The assignment of specific proposed degrees of safety in the specific proposals summarized above illustrates the use of the proponents' frameworks and the degrees of safety suggested by the proponents as appropriate.

The Commission reserves judgment as to any specific degree of safety, pending further progress in its safety-goal development efforts.

NRC FORM 335 (7 77)		U.S. NUCLEAR REGULATORY COMMISSION BIBLIOGRAPHIC DATA SHEET		1. REPORT NUMBER (Assigned by DDC) NUREG-0764 For Comment	
4. TITLE AND SUBTITLE (Add Volume No., if appropriate) Toward a Safety Goal: Discussion of Preliminary Policy Considerations				2. (Leave blank)	
				3. RECIPIENT'S ACCESSION NO.	
7. AUTHOR(S) Office of Policy Evaluation				5. DATE REPORT COMPLETED MONTH YEAR March 1981	
9. PERFORMING ORGANIZATION NAME AND MAILING ADDRESS (Include Zip Code) Office of Policy Evaluation USNRC Washington, D.C. 20555				DATE REPORT ISSUED MONTH YEAR March 1981	
				6. (Leave blank)	
				8. (Leave blank)	
12. SPONSORING ORGANIZATION NAME AND MAILING ADDRESS (Include Zip Code) Same as above				10. PROJECT/TASK/WORK UNIT NO.	
				11. CONTRACT NO.	
13. TYPE OF REPORT Regulatory Report			PERIOD COVERED (Inclusive dates)		
15. SUPPLEMENTARY NOTES				14. (Leave blank)	
16. ABSTRACT (200 words or less) <p>This report includes a Commission statement on preliminary policy considerations in development of an NRC safety goal and a supporting report providing a more detailed discussion. The discussion begins with an exposition of criteria by which the quality of different approaches to a safety goal may be judged. It next addresses alternative approaches to safety-goal formulation and presents some safety-goal proposals. Finally, there is a brief discussion of degree of safety.</p>					
17. KEY WORDS AND DOCUMENT ANALYSIS					
17a. DESCRIPTORS					
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 \$

NRC FORM 335 (7-77)



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

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE, \$300

POSTAGE AND FEES PAID
U.S. NUCLEAR REGULATORY
COMMISSION



Digitized by Google

Original from
UNIVERSITY OF MICHIGAN

