

7
RECEIVED BY TIC OCT 16 1978

NUREG/CR-0400

RISK ASSESSMENT REVIEW GROUP REPORT TO THE U. S. NUCLEAR REGULATORY COMMISSION

MASTER

H. W. Lewis, Chairman

R. J. Budnitz

W. D. Rowe

H. J. C. Kouts

F. von Hippel

W. B. Loewenstein

F. Zachariasen

Ad Hoc Review Group

Prepared for
U. S. Nuclear Regulatory Commission

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency Thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

NOTICE

This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Nuclear Regulatory Commission, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, nor assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, nor represents that its use would not infringe privately owned rights.

Available from
National Technical Information Service
Springfield, Virginia 22161
Price: Printed Copy \$5.25 ; Microfiche \$3.00

The price of this document for requesters outside of the North American Continent can be obtained from the National Technical Information Service.

**RISK ASSESSMENT REVIEW GROUP REPORT
TO THE
U. S. NUCLEAR REGULATORY COMMISSION**

H. W. Lewis, Chairman
R. J. Budnitz W. D. Rowe
H. J. C. Kouts F. von Hippel
W. B. Loewenstein F. Zachariasen

Manuscript Completed: September 1978
Date Published: September 1978

Ad Hoc Risk Assessment Review Group
H. W. Lewis, Chairman

Prepared for the
U. S. Nuclear Regulatory Commission
Washington, D. C. 20555

September 8, 1978

The Honorable Joseph M. Hendrie
Chairman
U. S. Nuclear Regulatory Commission
Washington, D.C.

Dear Chairman Hendrie:

I am pleased to transmit herewith the final report of the Risk Assessment Review Group. You will note that it contains a number of recommendations directed toward better use of risk assessment methodology within the Commission, an assessment of the achievements and limitations of WASH-1400, and a number of disjoint items.

I hope that the report will be useful to you, and would like to take this occasion, on behalf of my colleagues and myself, to thank you and your staff for your cooperation and help. In particular, I wish to highlight the outstanding support provided by John Austin, your Designated Federal Officer for this task.

Finally, I would like to note, for the record, that this was truly an independent study. You and all the Commissioners encouraged us at all times to call the shots as we saw them, and never tried to influence us in any way. Our report is as objective as we are able to make it.

Best Regards,



H. W. Lewis

TABLE OF CONTENTS

	<u>PAGE</u>
Summary.....	vi
Findings.....	viii
Recommendations.....	xi
Introduction.....	1
I. Risk Assessment Methodologies.....	4
II. Statistical Issues.....	6
III. Completeness.....	14
IV. The Data Base.....	16
V. The RSS Assessment of the Damage to Human Health From Radiation After a Postulated Accident.....	17
Disjoint Items.....	24
1. Browns Ferry.....	24
2. Common Cause Failures.....	27
3. Human Factors.....	29
4. Format and Scrutability.....	32
5. The Peer Review Process and the Adequacy of Treatment of Peer Comments.....	34
6. Earthquakes.....	37
7. Risk Perception.....	39
8. Allegations by UCS Concerning RSS Treatment of Quality Assurance/Quality Control.....	41
9. Current Role of Probabilistic Methods in the Regulatory Process.....	44
10. Acts of Violence: Sabotage and War.....	45
11. ATWS (Anticipated Transients Without Scram).....	46
12. The Influence of Design Defects and Quality Assurance Failures.....	48
13. Calculation of Population Doses From Given Releases of Radionuclides.....	50
Annex I - The Risk Assessment Review Group.....	51
Annex II - The Charter.....	53
Annex III - Correspondence Regarding Formation of the Review Group.....	54
Annex IV - Meetings of the Review Group.....	66

SUMMARY

The Risk Assessment Review Group was organized by the U.S. Nuclear Regulatory Commission on July 1, 1977, with four elements to its charter:

- (1) Clarify the achievements and limitations of WASH-1400, the "Rasmussen Report."*
- (2) Assess the peer comments thereon, and responses to those comments.
- (3) Study the present state of such risk assessment methodology.
- (4) Recommend to the Commission how (and whether) such methodology can be used in the regulatory and licensing process.

The group was formed to represent a wide spectrum of views about nuclear safety, though each member was chosen for his technical expertise. We have profited from a year of study and testimony, and wish to acknowledge the outstanding cooperation we have received from the staff of the Nuclear Regulatory Commission, the nuclear industry, and concerned scientists and citizens.

We find that WASH-1400 was a conscientious and honest effort to apply the methods of fault-tree/event-tree analysis to an extremely complex system, a nuclear reactor, in order to determine the overall probability and consequences of an accident. We have reviewed the methodology, the data base, the statistical procedures, and the results.

We have found a number of sources of both conservatism and nonconservatism in the probability calculations in WASH-1400, which are very difficult to balance. Among the former are inability to quantify human adaptability during the course of an accident, and a pervasive regulatory influence in the choice of uncertain parameters, while among the latter are nagging issues about completeness, and an inadequate treatment of common cause failure. We are unable to define whether the overall probability of a core melt given in WASH-1400 is high or low, but we are certain that the error bands are understated. We cannot say by how much. Reasons for this include an inadequate data base, a poor statistical treatment, an inconsistent propagation of uncertainties throughout the calculation, etc.

Also, both the dispersion model for radioactive material and the biological effects model should be improved and updated before they are applied in the regulatory and licensing process.

*U.S. Nuclear Regulatory Commission, Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants, WASH-1400 (NUREG-75/014), October 1975. Available from National Technical Information Service, Springfield, VA 22161.

We do find that the methodology, which was an important advance over earlier methodologies applied to reactor risks, is sound, and should be developed and used more widely under circumstances in which there is an adequate data base or sufficient technical expertise to insert credible subjective probabilities into the calculations. Even when only bounds for certain parameters can be obtained, the method is still useful if the results are properly stated. Proper application of the methodology can therefore provide a tool for the NRC to make the licensing and regulatory process more rational, in more properly matching its resources (research, quality assurance, inspection, licensing regulations) to the risks provided by the proper application of the methodology. NRC has moved somewhat in this direction, and we recommend a faster pace.

Among our other findings are the well-known one that WASH-1400 is inscrutable, and that it is very difficult to follow the detailed thread of any calculation through the report. This has made peer review very difficult, yet peer review is the best method of assuring the technical credibility of such a complex undertaking. In particular, we find that the Executive Summary is a poor description of the contents of the report, should not be portrayed as such, and has lent itself to misuse in the discussion of reactor risks.

In summary we find that the fault-tree/event-tree methodology is sound,* and both can and should be more widely used by NRC. The implementation of this methodology in WASH-1400 was a pioneering step, but leaves much to be desired.

*One of us (F.v.H) is doubtful that the methodology can be implemented so as to give a high level of confidence that the probability of core melt is well below the limit set by experience.

FINDINGS

- WASH-1400 was a substantial advance over previous attempts to estimate the risks of the nuclear option. The methodology has set a framework that can be used more broadly to assess choices involving both technical consequences and impacts on humans.
- WASH-1400 was largely successful in at least three ways: in making the study of reactor safety more rational, in establishing the topology of many accident sequences, and in delineating procedures through which quantitative estimates of the risk can be derived for those sequences for which a data base exists.
- We are unable to determine whether the absolute probabilities of accident sequences in WASH-1400 are high or low, but we believe that the error bounds on those estimates are, in general, greatly understated. This is true in part because there is in many cases an inadequate data base, in part because of an inability to quantify common cause failures, and in part because of some questionable methodological and statistical procedures.
- There are many sources of both conservatism and nonconservatism in WASH-1400 and we have no way of judging their relative magnitudes in the WASH-1400 probability estimates. For example, the treatment of common cause failures may be nonconservative, while the pervasive intrusion of regulatory conservatism throughout the report is a conservative factor. In addition, inability to quantify human adaptability during the course of an accident (as illustrated at Browns Ferry) and failure to take credit for this is a major source of conservatism.
- The achievements of WASH-1400 in identifying the relative importance of various accident classes have been inadequately reflected in NRC's policies. For example, WASH-1400 concluded that transients, small LOCA, and human errors are important contributors to overall risk, yet their study is not adequately reflected in the priorities of either the research or regulatory groups.
- Despite its shortcomings, WASH-1400 provides at this time the most complete single picture of accident probabilities associated with nuclear reactors. The fault-tree/event-tree approach coupled with an adequate data base is the best available tool with which to quantify these probabilities.
- It should be noted that the dispersion model for radioactive material developed in WASH-1400 for reactor sites as a class cannot be applied to individual sites without significant refinement and sensitivity tests.

- The biological effects models should be updated and improved in the light of new information.
- After having studied the peer comments about some important classes of initiating events, we are unconvinced of the correctness of the WASH-1400 conclusion that they contribute negligibly to the overall risk. Examples include fires, earthquakes, and human accident initiation.
- It is conceptually impossible to be complete in a mathematical sense in the construction of event-trees and fault-trees; what matters is the approach to completeness and the ability to demonstrate with reasonable assurance that only small contributions are omitted. This inherent limitation means that any calculation using this methodology is always subject to revision and to doubt as to its completeness.
- WASH-1400 made clear the importance to reactor safety discussions of accident consequences other than early fatalities. Although potential consequences other than early fatalities of large releases of radioactivity into the atmosphere (thyroid damage, land contamination, delayed cancers, genetic defects, etc.) may not have been prominently displayed among the findings of WASH-1400, they are systematically calculated there for the first time. Our resulting increased understanding of the full spectrum of reactor accident consequences has implications for nuclear power plant design, siting, and planning for mitigation of consequences.
- The statistical analysis in WASH-1400 leaves much to be desired. It suffers from a spectrum of problems, ranging from lack of data on which to base input distributions to the invention and use of wrong statistical methods. Even when the analysis is done correctly, it is often presented in so murky a way as to be very hard to decipher.
- For a report of this magnitude, confidence in the correctness of the results can only come from a systematic and deep peer review process. The peer review process of WASH-1400 was defective in many ways and the review was inadequate.
- Lack of scrutability is a major failing of the report, impairing both its usefulness and the quality of possible peer review.
- The Executive Summary to WASH-1400, which is by far the most widely read part of the report among the public and policy makers, does not adequately indicate the full extent of the consequences of reactor accidents; and does not sufficiently emphasize the uncertainties involved in the calculation of their probability. It has therefore lent itself to misuse in the discussion of reactor risk.
- WASH-1400 was directed to make a "realistic" estimate of risk. In the regulatory process, the usual conservatisms must be incorporated.

There have been instances in which WASH-1400 has been misused as a vehicle to judge the acceptability of reactor risks. In other cases it may have been used prematurely as an estimate of the absolute risk of reactor accidents without full realization of the wide band of uncertainties involved. Such use should be discouraged.

RECOMMENDATIONS

- Re-evaluate NRC's inspection and quality assurance system and licensing criteria to determine the extent to which they incorporate those things that have been learned from WASH-1400 and other relevant literature.
- Use WASH-1400 probabilistic methodology more effectively to guide the reactor safety research program so as to reduce the uncertainties in analysis, and to gain greater understanding of those points of risk uncovered.
- Where there is an inadequate data base, the methodology of WASH-1400 can still be used to uncover the topology of accident sequences. In such cases the limits of knowledge should be stated, without pressure to quantify (other than bounding) that which is unquantifiable.
- Communicate to the relevant branches of Government (e.g, Department of Energy) the desirability of performing risk assessments on electric generating technologies alternative to light-water reactors.
- In general, avoid use of the probabilistic risk analysis methodology for the determination of absolute risk probabilities for subsystems unless an adequate data base exists and it is possible to quantify the uncertainties. However, the methodology can also be used for cases in which the data base will only support a bounding analysis, and for other cases in the absence of any better information if the results are properly qualified.
- Fault-tree/event-tree analyses should be among the principal means used to deal with generic safety issues, to formulate new regulatory requirements, to assess and revalidate existing regulatory requirements, and to evaluate new designs.
- NRC should encourage closer coordination among the research and probabilistic analysis staff and the licensing and regulatory staff, in order to promote the effective use of these techniques.
- The consequence model used in WASH-1400 should be substantially improved, and its sensitivities explored, before it is used in the regulatory process.

INTRODUCTION

The Risk Assessment Review Group had its origin in the fall of 1977, in the wake of a long controversy about the technical credibility and uses of WASH-1400, the Rasmussen report, also referred to as the Reactor Safety Study or RSS. This was a pioneering effort to apply formal risk assessment methodology (through fault-tree/event-tree techniques) to the issues of reactor safety, and is the centerpiece of the action on that issue. Although the social question of the extent to which nuclear power ought to play a role, in the provision of energy in the United States devolves at least as much on such other issues as nuclear proliferation, the disposal of radioactive waste, the problem of reprocessing, other energy sources, etc., the issue of safety plays a key role and is also important in its own right. In addition, the study of reactor safety occurs against a background of inadequate experience to quantify the issues on that basis alone (there have now been accumulated some 300 reactor-years of commercial experience with the largest reactors, and some 2,000 reactor years of Naval experience of questionable applicability), and is accompanied by widespread public anxiety.

In addition to the role of quantitative safety analysis for the purpose of providing a portion of the information necessary for society to match risks and benefits more rationally, improved understanding of accident sequences can also serve to improve safety itself by enabling the NRC and the industry to more properly match resources to the weak points uncovered. (For example, the principal contributor to the PWR core melt probability uncovered during the Rasmussen study, a check valve problem, has already been addressed in those reactors to which it was applicable.) Thus the regulatory process itself can be made more rational.

The predecessor to WASH-1400 in assessment of risk from nuclear power plants was WASH-740, which in 1957 provided an estimated upper bound to the consequences of hypothetical large accidents in nuclear reactors. At that time it was not regarded as possible to estimate the probability of these large accidents. About 15 years later, WASH-1400 undertook to calculate both probabilities and consequences, for a wide spectrum of accidents. It will be clear in the sequel that we have many points of difference with the content of WASH-1400, but we do solidly support the effort.

Our charter has three elements, and it is clear from the above that they are interwoven and do not lend themselves to neat separation. We were charged to clarify the achievements and limitations of WASH-1400 - not only as an end in itself, but also as a means of assessing what can now be said quantitatively about reactor safety. In addition, we were charged to assess the "peer comments" on the Rasmussen report. Finally, we were charged to assess the current state of such risk assessment methodology, and to make a recommendation to the Commission on whether and how such methodology can be introduced into the regulatory and licensing process.

To accomplish this, the Commission has tried to assemble a balanced Review Group, including some who have been in the past critics of RSS, some previous defenders, and some who have not previously been involved. Not all "points of view" are represented since our purpose is to write a technically accurate report, rather than to debate the benefit of, or evils of, nuclear power. Further, we have confined ourselves to the study of the final version of WASH-1400, and not the draft, because to study the latter is to study history, not reactor safety.

Finally, we would be remiss if we did not acknowledge the outstanding cooperation we have had from the nuclear community, the NRC, the critics of both, and others. Many people have given up inordinate amounts of their time to be with us, or to prepare testimony for us, and we are grateful. If there are flaws in our understanding, it is not because the community has shirked in its effort to educate us.

On the other hand, our interaction with the community has revealed two sour notes which need to be mentioned. The term "peer" means equal, and (as will be discussed in detail below) effective peer review is the only method known to the technical community for quality assurance in its product. However, in the arena of reactor safety, a peer comment has come to mean anything written by anybody to anybody asserting anything about anything. The comment need not contain a better analysis, demonstrated technical expertise in the subject, evidence of error, or conformance to any of the normal standards by which the technical community assesses a peer comment. Thus, responsible peer comments, which do exist, are immersed in a sea of others, and this degrades the peer review process. In addition, part of the NRC staff, though completely cooperative with our panel, appears to have drifted into a stance which can at best be described as defensive, and at worst as a "siege mentality". Perhaps it is asking too much of people who are subjected to both rational and irrational attack to distinguish the two, but it would still be in the best interests of NRC to do so. (This is well illustrated in the BWR control rod controversy, to be discussed in grisly detail below. It would have been in the best interests of NRC had the staff long ago given up its stubborn defense of this indefensible calculation.)

Our report is organized as shown in the Table of Contents. It was clear at the outset that there were a large number of discrete, or disjoint, subjects which we needed to address, but which were not part of a coherent discussion of risk assessment and its possible applications. We have arbitrarily separated those and labeled them "disjoint items," though there is great arbitrariness in the selection thereby implied. Others may well have chosen differently.

We have many conclusions which appear in the main body of the report. For those who wish very concise summaries of our principal conclusions on two of the elements of our charter, they are:

1. The RSS has been most successful in the provision of a logical framework for the discussion of reactor safety, information about the relative probabilities of various accident sequences, and the beginning of an effort to provide absolute probabilities.

2. The information provided by probabilistic risk assessment about the relative importance of different accident sequences should be, to a much greater extent than is currently the case, incorporated into the determination of the priorities of the reactor safety research program, the determination of regulatory and enforcement priorities, etc. On the other hand, the methodology should only be incorporated into the regulatory process for those subsystems for which an adequate data base exists, and this can only be determined on a case-by-case basis. The methodology is the best guide we have to what is important to reactor safety.

3. WASH-1400 is defective in many important ways. Many of the calculations are deficient when subjected to careful and probing analysis, with the result that the accuracy of many of the absolute probabilities calculated therein is not as good as claimed. One key deficiency is the use by the study team of some methodological and statistical assumptions that lack credibility. Therefore, the absolute values of the risks presented by the Report should not be used uncritically either in the regulatory process or for public policy purposes.

I. RISK ASSESSMENT METHODOLOGIES

The methodology used in WASH-1400 for the determination of the probability of an accident begins with the construction of a logic tree leading from initiating events to final development. The logic structure consists entirely of "and" and "or" gates, causally connected. For example, if two open valves were in parallel, it would require a failure of both to close a fluid line, and the failures would be fed into an "and" gate before the consequence (closure of the line) could be assumed to occur. Clearly, there is an enormous complex of such structures, but (apart from questions of continuous variables, which will be discussed later) the method in principle is able to cope with all possible sequences of events. Unfortunately, the number of interconnects soon becomes impossibly large in a real system, leading to perhaps billions or more of possible logic tracks through the system from beginning to end. Because this is so, all the so-called different methods of analysis, or different methodologies, are simply alternate means of extracting from this intricate topology of sequences small and meaningful subsections which are tractable, and the application of differing levels of quantitative analysis to those subsets excised.

For example, event-tree analysis consists in limiting the universe of discourse by beginning at one particular event, and tracing forward in causal sequence through those paths which derive from it, so that one considers all possible consequences (within reason) of a particular initiating event. Fault-tree analysis is exactly the same, except that the event provides a starting point from which one traces connections backward in time, thereby discovering all those earlier circumstances which can lead to the event under consideration. (In a sense, the entire probabilistic part of WASH-1400 is equivalent to a fault-tree analysis, in that the many sequences that do not lead to core melt in the reactor were discarded. Only those which terminated in a particular event, core melt, were then included.)

With this picture, it is clear that there are many things one can do to carve out tractable methodologies. We have mentioned that one can proceed forward in time (event-trees) or backward in time (fault-trees) or one can concentrate on those sections of the logic tree which affect a particular component (failure modes and effects analysis) etc. All these are efforts to extract from an impossibly large collection of sequences a subset, chosen to be both important and finite, and treat it. The best choice will depend upon such matters as the completeness of the data base, the time and effort it is desired to expend, the objective (to estimate risk or to estimate a limit to it), etc., and it is feckless to discuss whether one method is better than another. In particular, it is incorrect to say that the event-tree/fault-tree analysis is fundamentally flawed, since it is just an implementation of logic.

On the other hand, the successful implementation of such a methodology rests heavily upon the availability of an adequate data base, effective computing tools, a sufficiently deep understanding of the detailed engineering of the system to permit construction of the logic

trees, and some logically sound procedure for limiting the universe of discourse. These are issues in the implementation of the methodology, rather than in the methodology itself, and are clearly less tractable as the systems become more complex.

It was mentioned earlier that there does exist a class of considerations which go beyond the dichotomous logic structures so far discussed, and that has to do with the existence of continuous variables. For example, in the real world, a valve is not limited to the two options of open or closed but sometimes simply leaks. Issues that involve partial degradation of the performance of components, questions of mitigation of the effects of failure, and such matters cannot be schematized in a dichotomous structure of this sort, but it can at least provide an imperfect model of the real situation. For example, no one would suggest the treatment of a fire through a logic diagram which branches with the probability that a fire starts versus the probability that it does not. There are all kinds of fires, some serious and some not serious, some which are put out early and some which are put out late, etc., and this class of situation is not easy to fit into the straitjacket of fault-tree/event-tree analysis. It requires an extension of such analysis, and efforts to develop such extensions are under way.

Apart, however, from considerations such as those contained in the last paragraph, the methodology used in RSS can be an effective tool for the analysis of accident sequences. It is in the data base, the aggregation of sequences, the statistical treatment and the success with which the system has been modeled that there are difficulties, and we deal with these separately.

II. STATISTICAL ISSUES

A. Introduction

RSS was faced with the problem of estimating the probability of occurrence of an extremely rare event - core melt - in a system of great complexity, a nuclear power reactor.

Since the event has never occurred in a commercial reactor, there are no direct experimental data on which to base an estimate. The only datum that exists is the observation that there have been no core melts in several hundred reactor-years of light water power reactor operation, and this fact provides at best an upper bound on the probability to be estimated. Therefore, it is necessary to resort to a theoretical calculation of the probability. But since the system is so complex, a complete and precise theoretical calculation is impossibly difficult. It is consequently necessary to invoke simplified models, estimates, engineering opinion, and in the last resort, subjective judgments. No one should be surprised, therefore, that RSS has generated considerable controversy over how these models, estimates, approximations, judgments and the like were chosen or made.

Theoretical calculations of the probability of an event are not, of course, automatically suspect. For sufficiently simple systems theoretically obtained probabilities are entirely reliable - for example, a symmetrical coin has a probability of 1/2 of coming up heads, and no experimental data are needed to make one believe that this theoretical probability is reasonable.

By the same token, subjective judgments as to the probability of an event can also be reliable. When the judgment is made by a person or a group highly experienced with the system in question, the predicted probability is sometimes quite accurate. A case in point is a horse race. The post-time odds are precisely the subjective probabilities of winning estimated for each horse by an experienced group - the handicappers and to some extent the betting public - and one finds empirically that the post-time odds and the race outcomes agree with one another to a remarkable degree over the long haul.

The overall methodology adopted by RSS to carry out its theoretical calculation of the probability of core melt has been outlined in Section I of this report. The complex overall system is broken down into its component subsystems, and probabilities of failure are generated for these. In some cases it is possible to find experimental data on subsystem failure probabilities, or at least data on other similar systems. In these cases the issue is: are there enough data, and are the data from sufficiently similar systems? In other cases models of the system are employed, which are simple enough to permit a theoretical calculation of failure probabilities. The issue in this case is whether or not the model is sufficiently realistic.

Once failure probabilities for the components are arrived at, they are then combined through the fault-tree and event-tree flow charts to generate an overall probability of core melt. Assumptions must be made regarding the independence of various events and various subsystems. This generates another area of concern: the possibility of an event influencing more than one component or fault-tree flow line at once. These common cause failures are treated separately in RSS, through a model constructed for this purpose. The model chosen has generated more controversy.

In what follows we shall address a number of the major criticisms of the statistical and probabilistic calculations in RSS to which we have been exposed, give our views on the issues so raised, and our assessment of their impact on the RSS conclusions.

B. The "Square-Root Bounding Model"

Perhaps the probabilistic issue which has raised most furor about RSS is the model used to estimate common cause failure probabilities. This is not because this particular calculation is more important to the overall conclusions in RSS than anything else; rather it is because, in contrast to so much else in RSS which has generated controversy, it is so well-defined mathematically. Other issues are much more subjective, and therefore argument over them tends to end up with each side simply saying the other is wrong, but unable to prove it.

The model adopted by RSS (called the "square-root bounding model") runs as follows. The true system is too complex to calculate a failure probability. Therefore, we shall make simple models, simple enough to permit calculation of the failure probability. Let M label the model, and let P(M) be the failure probability calculated in the model. Next we must assign a probability that the model is right. This probability cannot be calculated; instead one must make a subjective judgment about it. The judgment should of course be made by people experienced with the system being modeled; in this case there is hope that their assignment of a model probability is a good one. Call Q(M) the probability (in this case the subjective degree of belief) that model M is correct. The calculated probability of failure, P, is now given by

$$P = \sum_M P(M)Q(M)$$

provided that the set of models is exhaustive and exclusive. That is, P is the average of P(M) over the distribution Q(M).

So far so good; the only argument up to this point will be over the choice of Q(M). In RSS, the choice is to take Q(M) as log-normal in whatever parameters characterize the model.

But it is in practice too hard to make an exhaustive and exclusive set of models, so RSS makes two models, an upper bound model M_U and a lower bound model M_L . These are

again selected subjectively, presumably on the basis of engineering judgment. Then it is assumed that $P(M_U)$ and $P(M_L)$ are symmetrically placed; hence, since $Q(M)$ was chosen to be log-normal, the average P is

$$P = \sqrt{P(M_U)P(M_L)}$$

A noteworthy example of this unacceptable procedure in RSS is the calculation of the probability of failure of three given adjacent control rods in a BWR. M_L is chosen as a model in which all rods fail independently; thus $P(M_L) = P_1^3$ where (P_1) is the failure probability for a single rod, experimentally determined to be 10^{-4} . M_U is chosen as a model in which on 1% of the occasions one control rod fails, all fail. Thus $P(M_U) = 10^{-2} P_1$. (The 10^{-2} is claimed to have some experimental basis.) Hence RSS says

$$P = \sqrt{(10^{-2} P_1) P_1^3} = 10^{-9}$$

The degree of arbitrariness in this procedure boggles the mind. The lower bound model gives a bound which is so low as to be absurd, and there is no reason to believe that M_U is in any sense a symmetrically placed upper bound model. Nor is there any reason to believe that $Q(M)$ is log-normal. And the results are very sensitive to these arbitrary choices. For example, had $Q(M)$ been normal, RSS would have found

$$P = \frac{1}{2} (P_1^3 + 10^{-2} P_1) \sim 10^{-6}$$

In our view, criticism levelled at the "square-root bounding model" is entirely justified.

Could RSS have done anything better? Probably yes. The analysis in the ATWS report (admittedly performed later than RSS) is in our opinion a far better review of BWR scram failures. The NRC Division of Systems Safety staff has pointed out that a scram system was found inoperable during a test in an operating U.S.-designed BWR in Germany.

The staff calls this a scram system failure and obtains a (maximum likelihood) scram system failure probability of 1.5×10^{-3} /reactor-year based on 659 reactor-years of U.S. and foreign LWR experience. The probability for system failure during an actual demand would be somewhat lower, however, since it would depend on a demand occurring while the scram system was inoperable. Assuming monthly testing, a period of inoperability of two weeks per failure (for monthly testing), and six "significant" transients in LWRs per year, the probability would be approximately 6×10^{-5} /demand - 4 to 5 times higher than the WASH-1400 estimate.

In the absence of sufficient data, and in the face of such large theoretical uncertainties as are displayed in the BWR control rod problem, we feel that it is

preferable not to try to come up with a point estimate - a single number - for a failure probability, but rather to content oneself with bounds. To be useful, of course, the bounds cannot be too wide (thus a model like the "lower bound model" described above is useless). The sensitivities in the bounds can be estimated by varying the methods of calculation. And finally the errors and uncertainties must be made explicit, and carried through succeeding stages of the calculation to see how they affect the final conclusion.

C. Use of the Log-Normal Distribution for Fitting Data

RSS has a nearly universal practice of fitting every piece of data it has to a log-normal distribution, and, whenever an entirely unknown distribution needs to be subjectively chosen, to choose a log-normal distribution. (An example was given in Section B.) This choice is based on "engineering judgment" - most failure probabilities in other systems with which people are familiar are alleged to follow, roughly at least, a log-normal distribution.

Some of the data displayed in RSS do not look in the least log-normal, and yet are fitted to log-normal. Also, some of the log-normal fits to data that look log-normal seem to be bad fits - in that all the experimental points lie on one side of the fitted curve, for example. Further, in other cases there are so few data that we cannot say whether a log-normal fit looks reasonable or not.

Yet it is our view that these problems don't matter very much (except as they leave questions as to the validity of the statistical treatment in RSS). Errors in fitting may indeed occur, due to the pressure of time in producing so monumental a document as RSS, or simply due to the draftsman's sloppiness. And in some cases other distributions may indeed be better fits. But none of the fits we have seen is so far from the available data that use of the log-normal form is likely to misestimate truth by more than a factor of two or at most three. So within the errors listed in RSS, we accept the log-normal as an acceptable summary of most data.

D. The Use of Subjective Probabilities

RSS had to use subjective probabilities in many places. Without these, RSS could draw no quantitative conclusions regarding failure probabilities at all. The question is raised whether, since subjective probabilities are just someone's opinion, this has a substantial impact on validity of the RSS conclusions.

It is our view that use of subjective probabilities is necessary and appropriate, and provides a reasonable input to the RSS probability calculations. But their use must be clearly identified, and their limits of validity must be defined.

It is true that a subjective probability is someone's opinion. But, as the horse race example mentioned in the introduction illustrates, some people's opinions can be very

accurate, even in a quantitative sense. So the question really being raised is: how valid are the opinions of those whose subjective judgments were used in RSS? It must be kept in mind that the subjective probabilities used are for some of the individual steps in a sequence of events. No one provided a subjective probability for the overall probability of core melt. For many of the steps in which a subjective probability was used it was the output of experienced engineering judgment on the part of people familiar with events of that type. This, of course, does not guarantee the accuracy of the probabilities so generated, but if properly chosen makes them the best available. The only situation in which one might be concerned about a subjective probability leading to really major errors is where there is no experience at all. An example of this in RSS is nuclear pressure vessel failure. Such cases should be clearly identified, and results sensitively dependent on them treated with extra caution.

E. The Use Of Models

As we explained in the introduction, RSS was required to model many situations in order to be able to carry out what would otherwise be prohibitively difficult theoretical calculations. One extreme example of such a model was given in Section B in the "square-root bounding model." But there are many other places where models are used. For example, in the analysis of the Browns Ferry fire, to calculate the probability that the fire could have led to core melt, it is necessary to have the probability $P(t)$ that a given valve still operates at time t after onset of the fire. Some experimental information was available - namely, that several such valves did fail after a time t_0 . RSS then took $P(t) = (1/t_0) \exp(-t/t_0)$. This is a model; there is no a priori reason to believe it to be right. Though such a model may perhaps be arguable for failure rates, where little is known about the initiating mechanism, it is extremely dubious for repair times, where there is, after all, a minimum required. Further, even with the assumption of an exponential model (as, for example, for the relief valve failures), the assignment of the exponential coefficient in terms of the single observed failure time also involves a model assumption. This is a relatively subtle statistical point which we do not wish to belabor, but feel deserves mention.

Nevertheless, most models will not give wildly different answers. As with subjective probabilities, the use of models is necessary to get an answer. The choice of one model over another generates an uncertainty, but within that uncertainty the use of the model is justified, provided the uncertainty is estimated and indicated.

F. Variations Between Reactors

Only two reactors were studied in detail in WASH-1400: one BWR and one PWR. Yet the results were extended to 98 other reactors without estimating the additional uncertainties associated with this extension. Reactors are different. For example, many reactors didn't, in 1975, have the automatic recirculation pump trip which reduced, by a factor of 10, the calculated probability of core melt following a transient without scram in the BWR considered in WASH-1400. Some still don't.

It is true that RSS made no pretense to a careful analysis of more than two reactors, and that some allowances for individual site variability using demographic and meteorological data were included in extending the analysis to 100 reactors. But such an extrapolation is difficult, and has to introduce further uncertainties.

G. Completeness and Relevance of Data

Many of the constituent failure probabilities and distributions employed in RSS are based on experimental data, sometimes on actual reactor data and sometimes on data describing systems "similar" to those in reactors.

We have already (in Section C) commented on the use of the log-normal distribution to describe the data; at issue here is whether there are enough data and whether they are relevant data. The first of these questions is a statistical one; the second is an engineering one requiring a subjective judgment.

There are many cases in which the input data quoted in RSS are very skimpy. There are cases in which an entire distribution is "derived" from only a few data points. If this is all the information that exists, of course, then it is all that one can use. But the uncertainties associated with using it are then correspondingly large, and need to be taken into account and propagated through the entire calculation.

The relevance issue is more complicated. Can data on component failures, for example, from experience in a nonnuclear environment, be extrapolated to the extreme conditions which would be associated with a reactor accident? Here engineering judgment comes into play, with the difficulties already referred to in Section D. At least the identification of important kinds of uncertain data can generate enthusiasm for the experimental testing of some kinds of equipment under extreme conditions.

H. Propagation of Errors

In principle, of course, the methodology adopted by RSS in calculating the probability of core melt - or, more generally, the consequences of reactor accidents - requires assigning errors to each component probability and carrying them through to the end to arrive at an error estimate for the output numbers. To an extent this was done, but there are some omissions. Three examples are the BWR control rod failure uncertainties which were masked by use of the arbitrary model described in Section B, errors associated with some containment failure modes, and the errors in biological effects.

I. Probability "Smoothing"

These words describe a procedure employed in RSS in which rapid variations in release probabilities as a function of release category were suppressed. This was done by arbitrarily adding to the probability in each category 10% of the probability in the adjacent next larger and next smaller categories. For categories with small intrinsic

probabilities; this 10% addition is overwhelmingly the dominant contribution; thus the procedure artificially increases the overall risk estimate and lowers the precision as well.

While the numerical effect of "smoothing" is not large, the procedure is nevertheless arbitrary, unnecessary, and plain wrong, and had it produced the effect of reducing rather than increasing the risk estimate, would have generated an outcry.

J. Use of Median Instead of Mean

This should not have been a major issue, but since it has become one it is worthy of mention. RSS produces many probability distributions, and for convenience chooses to characterize them in terms of parameters rather than giving the entire distribution.

A given type of distribution can be uniquely characterized in terms of a finite number of parameters - for example, a mean and a variance for a normal or a log-normal distribution. Any independent set of parameters can be used for this purpose; thus one can also uniquely specify a log-normal distribution by giving its median instead of its mean. The mean and median are of course, different - the median being smaller - but given either, one can find the other if the distribution is known (e.g., log-normal) and a second characteristic parameter (e.g., variance) is also known. Therefore, as long as the parameter used to describe the distribution is clearly identified, there is no confusion or lack of precision about the entire distribution.

Nevertheless RSS has been criticized for giving the median instead of the mean of its log-normal distributions. Without adequate warning, some readers of RSS confused the two, and assumed that RSS was actually presenting a mean. They therefore inferred that RSS was predicting a lower failure rate than it actually did. In our opinion this is more an issue of scrutability than of statistics, so we shall not belabor it further here.

K. Conclusion

Most of our conclusions have already been listed in the preceding sections. They may be summarized as follows:

First, in view of Sections B-J, we feel that the uncertainty in the final probability produced by RSS is understated. RSS assigns a multiplicative factor of 5 uncertainty to its result; we cannot here, in the short time available, produce a different number, but we can, and do, feel that 5 is too small.

Second, we feel that RSS could have done a far better job of clarifying just what is assumed, what is calculated, and how. It is, of course, always easy to make criticisms of this sort after the fact; nevertheless the lack of care in labelling just what is a subjective probability, what is an experimental probability, what is a model and so on

in RSS is quite annoying to the reader. RSS could perhaps have saved itself a non-trivial amount of criticism had it made more effort to be precise about what it was doing.

III. COMPLETENESS

WASH-1400 isolated, from an impossibly large number of accident sequences, a relatively small number which it concluded were of significance to risk through their contribution to core melt. It is important to ask whether the WASH-1400 identification process was complete, and indeed whether the probabilistic methodology is capable, in practice or even in principle, of performing a complete analysis. It is, however, very important to be careful about what is meant by the words "complete" and "significant."

A discussion of completeness can very easily lead to a superficial answer, namely "by definition all such analyses are incomplete." But no one pretends that any such analysis can treat all possible accident sequences, so the question of completeness must be defined more carefully. By significance we mean a factor of two or more in overall risk.

In the context of WASH-1400, the completeness question can mean: Did the study team have enough engineering insight to write down all potential significant sequences? It can also mean: In the WASH-1400 analysis, did they stop calculating too soon? In the context of possible weakness or inapplicability of the methodology in general, it can also mean: Do multi-point failures contribute significantly to the risk, in a manner that cannot be quantified through the methodology? Or, are reactor systems somehow so complex that, as a practical matter, the methodology cannot provide risk estimates quantified with enough confidence to be of any use?

Let us begin with the generic methodological issues. On the question of multi-point failures, the issue is whether the methodology is capable of coping with accident sequences involving multi-point failures, which bridge numerous event-trees but are not "common-caused" in the sense of being caused by some common initiating event such as an earthquake or loss of off-site power. There is no way to answer this question without reference to a specific attempt to use the methodology, which of course means reference to WASH-1400. The WASH-1400 experience has been reasonably favorable, although by no means conclusively so: WASH-1400 did identify some important instances involving multi-point failures, demonstrating that the methodology has no difficulty in principle with issues of this kind. In fact, an event-tree formed by bridging other event-trees in this way looks like any other event-tree in the analysis; at this stage one realizes that the event-trees generally consist of multiple failures, and the question reverts to the previous one as to whether the analysts have been wise enough to find all significant sequences. The issue of whether WASH-1400 itself might be flawed in this regard is not a generic methodological issue, and will be discussed just below.

Are reactor systems somehow so complex that as a matter of either principle or practice the methodology cannot provide overall risk estimates quantified completely enough, and with enough confidence, to be of any use? This is not an issue of principle but rather of

implementation. As our Group has deliberated together on this issue, we find we are of mixed minds. Although several sequences have been recalculated since the issuance of WASH-1400 with some change in estimated contribution to risk, in this interval no new sequences have been shown to be important (in the sense that they double or triple the overall risk estimated in WASH-1400). Yet a new and important sequence could be established as important at any time. Some of us have a nagging uncertainty that something will turn up. Yet here the reactor accident record to date comes into play: with several hundred reactor-years of relevant experience, the WASH-1400 result on core melt probability is less than two orders of magnitude below the current experience record, and precursors of major accidents unidentified in WASH-1400 might perhaps be expected to begin to show up. If they have, they have so far escaped recognition (with the possible exception of Browns Ferry). Thus, some of us conclude that unidentified event sequences significant to risk might contribute at most something between a small increment and an increment expressed as a small multiplicative factor above the WASH-1400 result, as opposed to an order of magnitude or more.* It is important, in our view, that potentially significant sequences and precursors, as they appear, be subjected to the kind of analysis contained in WASH-1400, in such a way that the analyses are subjected to peer review.

On the subject of whether the WASH-1400 analysis stopped too soon (in the process of separating the many inconsequential sequences from the few important ones), a partial answer to the question can be obtained by rough estimates of the probability and numbers of identified sequences that were omitted. We have no reason to doubt the results, but we have not independently verified the assertion by the WASH-1400 study team that the contributions are small, in the aggregate, when one does an estimation of the sum of the known improbable sequences.

*One of us (F. v. H.) questions whether, for a system as complex as a nuclear power plant, the methodology can be implemented to give such a high level of confidence that the summed probability of many known and unknown accident sequences leading to an end point such as a core melt is well below the limit set by experience.

IV. THE DATA BASE

The frequency of occurrence of low probability events is by definition difficult to obtain empirically. Although the RSS team used existing reactor history, this was inadequate by itself for many low probability events and faults, and the team had to use data on failures of other similar systems and independent components. In many cases, transferability to reactor situations is straightforward and justifiable. In other cases, difference in design, use, or circumstances could alter the failure rate in nuclear service. But in general, reliance on the expanded data base drawn from non-nuclear experience is an improvement on use of the much more limited data base from purely nuclear use. This was the best that could be done.

However, two procedural considerations significant for follow-up must be addressed. First, areas in which there is a paucity of data should be particularly examined to uncover how better data can be obtained. Second, as new data, including additional reactor-years, recorded events and failures, and better component reliability estimates are made, these must be entered into the process in a formal and continuing manner. New data of significance are being developed in many important areas. Reanalysis based on the improved data should be done as the occasion arises. In particular, it is important that any regulatory application of RSS methodology rest at all times on the best and most current data and methodology.

For instance, the estimate of failure probability distribution of pipes, which spans several orders of magnitude, might be reevaluated. A study of the degradation of component performance in LOCA environments is now being conducted.

V. THE RSS ASSESSMENT OF THE DAMAGE TO HUMAN HEALTH
FROM RADIATION AFTER A POSTULATED ACCIDENT

A. Introduction

The authors of the RSS were guided in estimating the deleterious effects of radiation on humans by an advisory group on health effects, consisting of eminent physiologists, physicians, geneticists, and biological specialists from a broad range of disciplines. The advisory group was charged with assessing the disease effects, including placing uncertainty bands or ranges on the estimates. The goal was to make as realistic an estimate as possible, without resorting either to conservative assumptions that would tend to overestimate the consequences or to the converse.

The problem, as posed, was a considerable challenge to the several experts who worked on it. Over the years, many of these scientists and physicians had been involved in various national and international committees and panels concerned with "standard-setting" in the context of radiation protection. In this latter activity, it is appropriate to mask ignorance or uncertainties by leaning toward the more "conservative" side, thereby assuring that the public (or occupational) health is protected by a greater margin rather than a lesser margin. The challenge of making a realistic estimate was non-trivial.

B. Difficulties in Making the Assessments

Of course, the key difficulty with any realistic assessment of the effects of radiological exposure is the paucity of human data: the most useful data are from those few unfortunate circumstances that arose in years past from accidents, ignorance, or war, situations that everyone hopes will not recur. This difficulty is important even when assessing the consequences of exposures quite similar to those represented by the existing data. Thus, even when assessing the likely consequences, from a reactor accident, of exposures similar to those received at Hiroshima or Nagasaki in 1945, the human data are not wholly adequate to the task. For exposures quite different from those represented by the data, problems are all the more difficult. For example, variations across demographic sectors (age, race, sex, prior disability, place of residence) are key issues. Also, the assessments made in RSS require understanding of the dose-response relationship at relatively low doses (below about 50 rads). Because the data at lower doses are few and inadequate, it is necessary to estimate this relationship using data at much higher doses, which is a very difficult task.

The paucity forces consideration of other classes of data: both animal data in vivo, and data in vitro on human and animal systems. Here the problem is to derive sufficient understanding of these directly useful data to make possible the leap to human response in a quantitative and reliable way.

Another key issue is that all radiation-induced effects are also produced by insults other than radiation: indeed, some of them may be multi-factorial in origin, requiring more than one kind of insult to produce the effect. (Occasionally, multiple factors may afford protection too!) Since almost none of the effects has a well-understood etiology, damage estimates are almost always based upon combining empirical observations with limited physiological understanding.

Equally problematical is the fact that the doses which might induce cancers after a postulated reactor accident include some which are delivered over a protracted period of many years; this protraction may have an important effect on the response, especially for cancer induction, but the exact nature of that response is not known.

C. The Key Question

The main question to be addressed is: How well did the RSS team succeed? How well did the RSS estimates represent the consensus of the community of experts? Did the uncertainties encompass the spectrum of genuinely respectable differences of opinion?

D. Methodology Used in RSS

The methodology used was different for the three main types of end-point: acute high exposures leading to death or severe disability; disease, mainly cancer and typically involving a long latency period, from lower or chronic exposures; and genetic effects that produce disease or malfunction in subsequent generations.

1. The Acute Lethal Dose Issue (LD-50/60)

The RSS analysis of early mortality from high acute radiation doses is based mainly on bone marrow injury. The analysis takes into account not only the direct mortality data themselves but also the availability and application of medical therapeutic intervention. This intervention would, if applied, allow large numbers of acutely exposed persons to survive and recover from these high doses following a calculated accident.

Basically, the RSS analysis relies upon three different response functions for mortality within 60 days: one response function for cases with little or only minimal treatment; one for cases with careful supportive treatment; and one for cases with "heroic" treatment, involving bone marrow transplants and so on. It was the judgment of the RSS team that, in the event of a serious accident, medical resources throughout the nation would be mobilized to treat the exposed individuals, and that the middle response function (careful supportive treatment) would be the appropriate one to use. This seems to us likely, although the ability to carry out such intervention not only has not been demonstrated, but isn't even well planned at this time.

The issue is usually cast in terms of the "LD-50/60" concept. The LD-50/60 is defined as that lethal dose which results in the death of 50% of the exposed population within 60 days. For the three classes of treatment, the LD-50/60 used in RSS is 340 rads (minimal), 510 rads (supportive), and 1050 rads (heroic). The LD-50/60 is, of course, only one data point on a complete response function. This function is shown in RSS in the figure in Appendix VI on p. F-37, and covers the full range of response.

One key problem is the possible inapplicability of the data that were used. For example, if a human is exposed not only to bone marrow dose from inhalation and ingestion, but also to external radiation, then the LD-50/60 might be significantly lower. There are animal data to support some LD-50/60 reduction, and thus a corresponding increase in the fatalities from any exposure profile. This type of effect would be a counterpoint to the benefit of therapeutic intervention.

The data available to support the RSS analysis are weak, and this is widely recognized in the medical community. Among the experts, there is a wide range of opinion about what LD-50/60 to use. The RSS team appears to have used a combination of available data and the judgment of experienced physicians.

It appears to us, based upon conversations with members of the expert community, that the LD-50/60 value of 510 rads used in RSS for the case of supportive medical treatment is as reasonable an estimate as any other. However, there seems to be a very wide range of uncertainty in opinion, with values cited in the range from perhaps 400 to 600 rads.

The use of the response function for only minimal intervention (LD-50/60 = 340 rads) would result in approximately three times as many fatalities. Thus the range of 400 to 600 rads, which we feel encompasses the reasonably likely spectrum of scholarly opinion for supportive intervention, implies differences in fatalities of about a factor of 2 increase (for 400 rads) or 2 decrease (for 600 rads).

2. Disease (Mainly Cancer) From Doses Below Acute Lethal Dose

The starting point for the RSS assessment seems to have been the conclusions of the BEIR Report, published by the National Academy of Sciences in 1972. At the time of the RSS study, this well-known effort was the latest comprehensive examination of the effects of ionizing radiation, containing thorough examination of essentially all of the then-extant data.

The RSS team did not stop with the BEIR Report and its conclusions, however. It was a different group from the BEIR Committee, and it made some different judgments.

a. Relative or Absolute Risk

There is a normally existing cancer rate, resulting from both natural and man-made environmental causes. Reactor-accident-induced cancers would be an increment to this existing rate, and a question generating some controversy is whether or not the amount of increment is a function of the existing rate.

Two simple models exist: in one, the radiation-induced cancer rate is said to be an increment independent of the existing rate; in the other, it is taken as a multiple of the existing rate. These are known as the absolute-risk and relative-risk models, respectively. Since the existing rate varies with, among other things, age, the total number of accident-induced cancers over a given period will be greater in the relative than in the absolute risk model. The RSS analyses used the absolute-risk model, and therefore, the estimate of latent cancer fatalities from a reactor accident has been criticized by some as an underestimate.

The issue of the "plateau period" (that period after exposure during which cancer incidence continues to show an excess over nonradiogenic rates) is tied in with the relative-risk/absolute-risk issue. In particular, the incidence of almost all nonradiogenic cancers is widely known to increase dramatically with increasing age. Thus, a relative-risk model with a plateau period extending throughout life would predict that the radiation-induced cancer incidence would also increase with age, while an absolute-risk model with a finite plateau period (say, 30 years) would predict that the radiogenic cancer incidence would fall off after the end of the plateau period.

It is not yet possible, either on the basis of our existing understanding of the physiological systems affected by radiation or on the basis of empirical data, to make a clear choice between the two models, or more generally (since probably neither model is in fact universally correct) to decide what the functional relationship between existing and induced cancer rates really is. One difficulty is that the relationship apparently depends upon the kind of cancer in question.

This is one issue for which the existing nonhuman data are of little use. Also, the human data are still inadequate. The human data upon which the estimates rest are mainly for exposures (such as of the atomic bomb victims and the uranium miners) several decades ago, but for which the study populations are still mostly alive. Follow-up of these cases through their full life spans will ultimately allow better understanding of this issue. It appears at this time that at least a few end points (for example, leukemia and possibly bone cancer) show some fall-off after perhaps a 30-year plateau period; for these, the absolute-risk model would be appropriate. On the other hand, for many others the recent human epidemiological data show no

such fall-off; indeed, the incidence seems to continue with age, and perhaps to increase with older age similarly to the incidence of nonradiogenic cancers.

The RSS analysis used an absolute-risk model. If all cancers obeyed a relative-risk model, the predicted increase over the RSS estimates would be about a factor of 1.5 for a 30-year plateau period and a factor of 4 for a lifetime at risk. Several criticisms of the RSS analysis have focussed on this point, saying that both the RSS central estimate and the uncertainty range are too low and do not encompass the likelihood that the lifetime-plateau, relative-risk model might ultimately turn out to be closer to the actual data as they come in. In our view, the use of the absolute-risk model in RSS without taking into account the possibility that some end points might be better described by a relative-risk model, is an incomplete description. Our study group as nonbiologists has not been able to decide that one approach is better than another. However, if the relative-risk model turns out to be a better descriptor for some of the end points, then the RSS estimates for carcinogenesis would have to be increased.

b. Dose Rate Reduction Factors

A distinction was made between the effects of high doses of radiation delivered at a high rate, on the one hand, and radiation at low dose or dose rate on the other. It was recommended in WASH-1400 that in most respects the carcinogenic effects be assumed reduced in the latter case compared to the former.

The BEIR report, although it discussed this issue, did not make this recommendation. It used a single value of cancer incidence per million man-rem, regardless of dose level or dose rate. At the same time, the BEIR report states that this assumption may well overestimate the risk when the radiation is delivered at a very low rate. Other radiation advisory bodies (NCRP, UNSCEAR) have also stated the possibility that radiation may be less carcinogenic at low doses and low dose rates.

The advisory group to RSS recommended that for realistic analysis it should be assumed that low-LET radiation at low dose and low dose rate be assigned a dose reduction factor. This factor was recommended to be as low as 0.2 at the lower dose rate, rising to 1.0 for doses above 10 rem/day when total exposure exceeds 25 rem. These factors were to be applied in calculating a central estimate of incidence of most types of cancer, with the notable exception of the female breast, where specific data on incidence of breast cancers in mice indicate no dose reduction in this tissue.

The RSS in fact did use these dose rate factors in calculating the central estimates of cancer incidence. There were also calculations of upper limit

estimates which used a dose effectiveness factor of unity throughout, and lower limit estimates which used purely quadratic dose effectiveness factors of a form that seems to fit best curves of incidence of some forms of cancer.

There is a fundamental difficulty in obtaining general agreement in dose effectiveness factors. The data on radiation-induced human carcinogenesis are still too sparse and too contradictory to provide a basis in themselves for these values. As a result of the indefinite character of the data base, there is a spectrum of opinion among radiobiologists as to the use of dose reduction factors. Some radiobiologists believe it probable that the factor of unity is the best approximation to current understanding. Some very few radiobiologists even argue for dose effectiveness factors exceeding unity at low doses. We have found that the more common situation is agreement on the basis of general radiobiology that dose reduction factors should be used, but lack of agreement on values and circumstances.

c. Conclusions on Carcinogenesis

It is our summary viewpoint that an upper limit of the full range of uncertainty on delayed dose-effect estimates of the order of a factor of 5, but certainly no more than a factor of 10, above the RSS values would encompass most of the range of scholarly disagreement in the biological community relative to carcinogenesis at low dose rates.

3. Genetic Effects

For genetic effects it is necessary to use animal and in vitro data, because there are no good genetic effects data in humans. There is reason to place greater confidence in the use of animal or in vitro information than is true for somatic effects, but nevertheless, the overall confidence in the human risk is less because of the lack of direct human observations. This area is still in a great state of flux, with much research actively under way. It seems to us useful to consider updating the RSS estimates of genetic effects, and the uncertainties, from time to time. Nevertheless, there remains the key problem of understanding how to allow for future medical progress in prevention, therapy, or cure of many of these diseases: when discussing diseases many generations hence, this factor escapes estimation quite completely.

4. Thyroid Dose-Effect Coefficient

The calculations in WASH-1400 used a dose effectiveness reduction factor of 0.1 for iodine-131 dose relative to X-rays. The documentation cited in support of this choice of dose reduction factor for humans (A. M. Safa, et al.) gives an alternate interpretation for low incidence of thyroid nodules, i.e., that thyroid destruction doses occurred. As a result this factor has been questioned and further review is necessary.

E. Conclusions and Overview

In this discussion, we have attempted to describe the complexity of the issues involved in estimating disease and mortality effects from radiation exposure. This complexity prevents the radiobiological community from making precise estimates. The issues are difficult, and many years will pass before we attain resolution, even in the best of circumstances.

Another little-discussed but widely understood point is that our understanding of the effects of radiation-induced diseases, while meager and empirical on the whole, far exceeds our understanding of almost all other comparable disease effects, such as those from chemical or biological agents, under either acute or chronic exposure conditions. The radiation biologists have nothing to be ashamed of in their ability to make the estimates of radiation effects found in RSS, and in the luxury they enjoy to argue over detailed interpretation so vigorously and with so much and varied data. This sobering point underscores the need to consider the entire injury/disease controversy in its larger context of environmental impact rather than in the isolation of effects of radiation alone.

It is inevitable that further work will be done, and revisions made, on the RSS estimates. There is much disagreement about the details of the estimates made by the RSS team charged with making the disease and mortality estimates. For example, although all members of the RSS team contributed their full and honest efforts to the task, the spectrum represented by that team was not broad enough to encompass the full range of scholarly opinion on the subject. This led the RSS team to make estimates with a narrower range of stated "uncertainty" that would otherwise have been the case.

DISJOINT ITEMS

1. BROWNS FERRY

Background

The fire at the Browns Ferry reactor site, a TVA BWR at Browns Ferry, Alabama, on March 22, 1975 is of particular importance for a number of reasons. It was the most serious accident to befall a commercial light water reactor in the United States, it produced common cause failures in a variety of otherwise independent and redundant systems, and it compromised the reactor safety systems more than has ever occurred in other events. During the fire various means were used to make up water lost by evaporation, but after 5-1/2 hours the situation was at its worst. At that time, with control of all eleven relief valves lost, the ability to maintain the system at low pressure was also lost, and the only source of makeup water was a control rod drive pump not intended for this purpose and of inadequate capacity for the long term. In the end the valves were repaired in time, and the reactor was safely shutdown, but it is important to analyze the lessons that have been learned from this event.

Was it a "near miss"?

This term has been widely used, and one should be quite clear about its meaning, since it implies the assignment of other possible outcomes to a sequence of events that actually occurred. Clearly the exact sequence of events at Browns Ferry could only have had the outcome it actually had since that is history, but it is possible to do alternate sequencing by searching within the actual events for those which may have had a non-deterministic component, and studying the alternate futures provided thereby. This is not a trivial point, in that one assigns statistical character to causal events only at some risk, and the analysis of Browns Ferry has lent itself to this error.

An example will suffice. If a meteorite falls from the sky and lands within a hundred yards, it is proper to call it a near miss in the sense that the places on which meteorites fall can be expected to be randomly chosen by the meteorites, so that it is possible to assign a random process to the impact point. (Incidentally, another trap into which some people fall is to infer, in this case, that the occurrence of an unlikely event makes it a probable event. It is, for example, not more likely that a second meteorite will fall near one, just because the first one did.) On the other hand, hundreds of cars that pass within five feet of one's own on a two-way street each day do not constitute near misses, because their trajectories are causally determined and it would be improper to introduce random variables without some understanding of a mechanism for such randomization. In this case, of course, some do cross the center-line occasionally, but it requires more knowledge than simply the five-foot miss distance to determine the frequency of such events.

The reason for these comments is that Appendix XI of WASH-1400 contains a probabilistic analysis of the Browns Ferry event, computed by assuming the initiation of the fire and assigning non-deterministic variables (in our view uncritically) to a number of the key parameters of the accident. For example, seven of the eleven available pressure relief valves failed early in the fire, and the remaining four after 5-1/2 hours. Since, as was mentioned above, the failure of the last four posed the most serious threat to the reactor, it would be interesting to know what the "likelihood" was of this latter failure having occurred much earlier when the reactor was boiling off coolant at a greater rate. This was calculated in WASH-1400 by assuming an exponential probability law for the failure of these last four valves, and using as the parameter for this exponential the observed 5-1/2-hour time for their failure. Thus the probability that these valves failed earlier than 5-1/2 hours is calculated to be 63%, using a model which ignores the mechanism of such failure. In actual fact, these last four relief valves failed because they were activated pneumatically, and their common accumulator simply ran out of air during their frequent activation to maintain the depressurization during the previous period. Clearly, in such a case, the mechanism is causal, and a blind probabilistic calculation is incorrect. (The same argument using an exponential model can be used to show that an automobile which has been observed to get 250 miles on a full tank of gas has a 15% probability of getting 500 the next tankful.)

Similar comments can be made about probabilistic calculations of repair times, without reference to underlying models, etc. For this reason, we believe that the question of whether the Browns Ferry incident was a "near miss" depends upon the extent to which one can in fact identify random elements in the sequence of events that actually occurred. These are few, and not dominant. We, therefore, prefer not to comment upon the question of whether the probability of core melt at Browns Ferry, given the initiating event, calculated in Appendix XI of the RSS (1/300) is more or less correct than values produced by alternative calculations of other people. We would rather ask what lessons can be learned from Browns Ferry, a very serious accident.

Was it Predicted by RSS?

Though this is not an entirely useful question, the answer is no, in that fire does not emerge from a reading of WASH-1400 as an important accident initiator.

There is a point which is not widely understood by the public, and requires some discussion. When a great many rare events each have a small probability of occurrence, the chance that at least one of them will occur can be rather high. One often finds people saying that some event, which has extremely low probability, has occurred, so that there has been magic somewhere. This is especially true if the event has some unusual interest. For example, in the game of bridge, all hands are equally probable in an honest game, and there are in fact 6.35×10^{11} possible hands. Even so, a person being dealt 13 spades will find himself the subject of a newspaper story, even though the event could have been entirely fortuitous. We do not know enough about the relative roles of causal and statistical factors in the development of the Browns Ferry fire to

comment upon whether this discussion is relevant, and this is in any case largely a matter of perception. However, whether or not the fire was or should have been predicted by RSS, we believe that it has emphasized the importance of common cause initiators such as fires, earthquakes, etc. Such initiators will be discussed in the next section.

Human Behavior

This accident was initiated by human error, thus emphasizing the importance of human initiators. On the other hand, after it was under way, the people present at Browns Ferry behaved well under trying circumstances, improvised where necessary, and provided that element of causal human intervention upon which we are often dependent. There were mistakes which are identifiable in retrospect, such as failure to use water on the fire in the early stages, but overall it was human intervention which saved the reactor. This is an important point, as we will discuss in our section on human factors, in that a substantial element of conservatism in the body of WASH-1400 (other than in the analysis of Browns Ferry) lies in failing to take credit for the fact that well-trained humans provide adaptive coping capability. It is probably not possible to quantify this fact, but it is gratifying that it is there.

Quality Assurance

There were serious failures of quality assurance at Browns Ferry, which contributed in a substantial way to the seriousness of the accident. Among these were violations of cable separation requirements whose function is to prevent the failure of redundant safety systems by fires in a single cable tray. This raises a very serious question of the extent to which we can be confident that the actual construction and materials in a reactor conform to the design specifications. This problem might become more serious if a long accident-free period results in an increasing level of complacency.

Conclusion

In summary we feel that the Browns Ferry fire was a very important event in the history of commercial nuclear power in this country, which should be milked for all the lessons which it can provide. Too much of the discussion has centered around proving, often by dubious means, that it was or was not a near miss, and to placing the blame. Too little time and energy have been expended in using the event to help quantify some of the inputs to a reactor safety analysis which are most difficult to quantify: quality assurance, human behavior, common cause failure, etc. It is even now not too late for the NRC to have some impartial body do these things.

2. COMMON CAUSE FAILURES

Much has been said and written about the question of whether common cause failures, which can bridge event-trees and cause simultaneous uncorrelated breakdown in different systems, can be or have been properly accounted for in RSS. (Indeed, the issue is regarded as so important by RSS that some thirty pages of Appendix XI on the subject are lifted verbatim from the corresponding thirty pages in the Main Report, without attribution. Appendix IV then refers us to both the Main Report and to Appendix XI without warning us that they are identical.) To understand the importance of common cause failures, it is important to review the philosophies of "defense-in-depth" and redundancy.

If an accident requires the concatenation of (say) three separate and independent events, each of which has a probability of occurrence of one in a hundred, the overall probability of the accident is clearly the product, which is 10^{-6} . In that event, the subevents occur often enough to be studied statistically and quantified. Further, in this hypothetical example, there will be some 10,000 occurrences of the first two events in the sequence before there is even a 1% probability of the final accident, so that the potential for learning by experience is great.

On the other hand, the common cause failure may invalidate some of the sequential features of an event-tree, while at the same time causing failures on adjacent and apparently disparate event-trees.*

It is argued in RSS that common cause failures cannot and ought not to be treated as items in themselves, but should be taken into account throughout the analysis, in that the probability of component failures must be understood in every case to be in part independent and in part commonly caused. There are places in the report in which this is done (through a so-called bounding approach) but only for those event-trees which have been selected for other reasons. We are unconvinced by the arguments for completeness of the report in connection with common cause failures.

Since the case for completeness cannot be made convincingly, one must look elsewhere to quantify one's confidence in the safety of reactors against common cause failures.

*It has been argued that a common cause failure can only lead to core melt if it appears upon one of the event-trees which in turn leads to core melt, and that, in that sense, the invention of new sequences (the so-called "what if's") is irrelevant. Of course this is true, but RSS does not deal with all event-trees, but only those which by normal analysis have appreciable probability of leading to core melt. The potential for a common cause failure to activate event-trees which would otherwise not pass through the filter afforded by RSS cannot be so easily discounted.

This is helped by a somewhat sharper delineation, since common cause can do several things - it can form part of the failure rate probability of individual components, it can compromise redundancy and defense-in-depth, and it can activate low probability accident sequences that would otherwise be overlooked. Though there are many sophisticated arguments about the extent to which such effects can be categorized, and it is indeed possible to some unquantifiable extent, and it is of course true that even for common cause failures all the normal defenses of diversity and redundancy are at least to some extent active, we believe that the principal assurance against common cause failures must lie in dealing with the initiating events. These are, for the most part, known, and are denumerable. Among these are:

1. Fire
2. Earthquakes
3. Acts of violence
4. Explosions and missiles, internal and external
5. Massive electrical failure
6. Human error
7. Floods, tsunamis
8. Tornadoes, hurricanes

If an adequate defense at the beginning is mounted against these initiating events, and an analysis can be performed that bounds them suitably, much of the philosophical discussion of the adequacy of the treatment of common cause failure becomes irrelevant. In that event the protection against common cause failures would lie in efforts to prevent or cope with the initiators, while safety with respect to normal accident sequences would lie in redundancy, defense-in-depth, and meticulous engineering, construction, and quality assurance.

3. HUMAN FACTORS

The role of human intervention is significant in many of the important event-trees and fault-trees in the RSS. The overall RSS risk assessment depends significantly upon the numerical values of certain human-error rates; also, the role of operators and others in mitigating the consequences of equipment failures should be important.

For this reason, we have attempted to review the way the human factor was incorporated into the RSS, and to examine the methodological approach more generally. As an opening comment, our Group has been surprised at the paucity of peer comment on this aspect of the RSS, considering its central role in the overall risk assessment. We must also indicate that none of our Review Group's members possess professional expertise in this area: our review has, therefore, relied heavily upon the advice and comments of others. Finally, although this section as written emphasizes the RSS analysis of the human role in accidents, we feel it important to note here our general high regard for the quality of power reactor operators and other staff.

The human factor enters the RSS analysis in several different ways:

1. In the course of an accident sequence, a human can operate certain safety equipment to assist in preventing an accident from developing fully. This can be by operating equipment either according to previously developed procedures, or according to improvised procedures as necessary.
2. An operator or other employee, attempting to follow normal procedures during the course of an incident; can make a mistake that aggravates the situation.
3. An accident can be initiated by an error on the part of a plant employee.
4. A plant employee (operator, maintenance crew, etc.) can inadvertently disable equipment during normal operation or when the plant is down (say, for calibration or maintenance purposes), so that its availability is compromised if an accident occurs.

The RSS event-tree/fault-tree methodology takes human intervention into account explicitly, although it can only do so insofar as its level of detail allows. The difficulties that our Group perceives in this aspect of the RSS are as follows:

1. The human performance data base is weak in many aspects. This is partly due to insufficient reactor-relevant experience, partly due to inadequate data collection, and partly because there has not been a strong motivation to collect certain types of data.

2. Much of the human-factor data used in the RSS came from related nonnuclear experience. This was necessary in order to complement insufficient nuclear power plant data.
3. It is not easy to construct event trees and fault trees realistic enough to incorporate human-error aspects in the proper way. The required level of detail is great, and the role of operators and other plant employees in bringing about equipment unavailability is in some cases difficult to incorporate.
4. Some human-intervention phenomena are time-dependent: for example, there may be only a narrow time window during which an operator must perform a certain task. The time dependence is in some cases particularly difficult to build into the methodology, although it is by no means impossible.
5. Where the data base was weak, or the event-tree too difficult to construct in enough detail, the RSS seems to have used the approach of letting expert opinion determine the quantitative contribution to the sequence, so as to complete the full evaluation. Uncertainties were then attributed to this expert opinion.

Despite the difficulties, the RSS recorded some important achievements in the human factors area. Among them were:

1. The role of operators and other plant employees in contributing to the major accident sequences has been delineated better than had been the case prior to the RSS.
2. The need for better human factors data has been explicitly recognized, including identification of those areas that probably contribute significantly to the overall risk. Although some have criticized RSS for its reliance on nonnuclear data, it seems to us that, to the contrary, experience with other highly complex systems operated by well-trained personnel is entirely relevant, and the RSS team is correct in using it.
3. Although our Risk Assessment Review Group has been unable to do enough detailed work on our own, we have surveyed several experts in this field. There seems to be a consensus among these experts that, despite the handicaps of a weak data base and difficult event-trees, the RSS estimates are well done. However, a substantive review has not occurred.

The limitation of the RSS in its human factors evaluation include the following:

1. For many of the event-trees in which human errors enter, there are insufficient relevant human data, and the RSS estimates relied on expert opinion.

2. There is the intrinsic methodological weakness due to the difficulty of incorporating time-window information into the event-trees.
3. Operators and other employees might make matters much worse during an accident, in ways that are intrinsically hard to analyze.
4. Operators and other employees might make matters much better during an accident, in ways that are intrinsically hard to analyze.
5. Due to points (3) and (4), there may be much uncertainty in the actual role of operators and other plant employees in accident sequences.

Our main overall comment is that although we don't know of any specific way in which the RSS methodology in this difficult area could have been improved, it has not been possible for our Group to determine how closely this part of the RSS work has achieved the accuracy quoted for the overall RSS risk assessment. We do feel that this area is one of the major contributors to the general problem faced by the RSS in making quantitative risk estimates.

Although we have no quantitative basis for defending our view on the subject, our Group does feel that the RSS study quite likely has significantly under-rated the role of operators and other employees in mitigating or controlling some potential accident sequences; this under-rating is due to the intrinsic difficulty in subjecting such intervention to event-tree analysis. It also seems to us, again without any definitive demonstration, that there might be some event-trees in which human error with important consequences might have been overlooked. We see no easy way at this time to balance out these two possible types of effects so as to arrive at a more confident overall assessment of the accident risk.

We recommend that the NRC undertake a systematic program to evaluate the need for better human data, so as to improve this aspect of any future RSS-type analysis. We know of one important study along these lines, now under way under Electric Power Research Institute sponsorship. We recommend that this study's results be considered by NRC staff along with other possible approaches to an improved data base.

4. FORMAT AND SCRUTABILITY

It must be said at the outset that RSS represents the fruits of a praiseworthy and pioneering effort to deal in depth with an extremely complex subject. It is a monumental report, written in three years by a large number of people and (although it is difficult to count the pages because they are not consecutively numbered) is nearly a foot thick. Clearly no one person can, or did, write all of this, and the report suffers thereby from incoherence. The report is divided into parts: a Main Report of some 150 pages (they are numbered), a short and separate Executive Summary, and eleven appendices which deal in greater detail with much of the material. This format is not unusual in major reports, in that the three parts cater to three different audiences. The Main Report caters to the serious reader, who wishes to know just what was accomplished, and how. The appendices are normally intended for people who are not so much readers as workers in the subject, who wish to go beyond the findings, conclusions, and comments of the main report into the detailed technical arguments and supporting data. The Executive Summary is normally reserved to the decisionmaker, who is neither a reader (for lack of time or, often, background) nor a worker, but still needs to know, succinctly, the main conclusions, findings, and recommendations of the report.

In the case of RSS, the functions of the Main Report and of the appendices (apart from Appendix XI, which was a response to comments on the draft report) were more or less as outlined above, but the Executive Summary was assigned an entirely different function. It was written, not for the decisionmaker with a shortage of time, but for the public, and is in a generally soothing tone. It begs all sorts of questions by comparing the risks of early fatalities from nuclear power plant accidents (favorably) with other man-made and natural risks, such as motor vehicles, air travel, electrocution, hurricanes, etc., each of which has an entirely different risk and benefit profile from nuclear power. Thus, the question has become a hot issue as to whether only prompt fatalities from nuclear accidents should be compared with prompt fatalities from (say) drowning as was done, or whether some differently weighted allowance should be made for the larger long-term effect. It is our view that it is unfortunate that the Executive Summary, which is in fact an interpretation of RSS for the public, was issued as part of the report in its present form. It adds little, but has served to divert attention from the more serious issues of reactor safety. It would, in our view, have been more proper (though separately arguable) for NRC to have issued a separate document comparing nuclear risks with others to which the public is subjected, and not to have associated it with RSS.

Another important subject is the issue of scrutability, as it relates to the whole issue of peer review. The latter depends heavily upon a sufficiently coherent and complete organization of the material in a report, that a serious coworker seeking to understand the technical basis for material in the report, or seeking to reproduce the

calculations for himself, can do so without undue* hardship. Unfortunately, we find that this is all too often not true of RSS. While recognizing that perfect coherence is only to be found in one-word poems, each of us has had the experience of trying to understand, in responsible detail, just how a particular subject was treated in the report. It is extremely difficult (perhaps an index would have helped), and one finds oneself shunted about from main report to appendix to appendix, always seeking the final word, but (like the search for the verb in a German sentence) never quite getting there. This would be a minor cavil, were it not that the confidence of peers in the integrity of the report rests on their ability to search out a subject of interest to them, and to assure themselves that it has been treated responsibly and well. Not to be able to do so is not to conclude that the treatment is faulty, but it does nothing to engender confidence in the results. Examples are legion.

While the previous paragraph may seem unduly harsh, the lack of scrutability is a major failing of the report; observed by nearly everyone who has worked with it. This is not to criticize in any way the clarity of mind of those who wrote the various parts of the report, but rather to say that its usefulness to the technical working community could have been much improved. Whether more time would have helped is not clear to us.

*Just what hardship is due a reader is debatable.

5. THE PEER REVIEW PROCESS AND THE ADEQUACY OF TREATMENT OF PEER COMMENTS

Since no technical person or group is infallible, the peer review process plays an essential role in quality control of technical output. Its function is to expose the data, processes, techniques, and conclusions of any scientific work to the community of peers. Criticism and study coupled with responsiveness do not assure that the work is right, but only that it conforms to the normal quality standards current in the trade. Though the RSS staff set out to expose itself to a peer review process, at least in part for these reasons, the undertaking was substantially more difficult than would have been the case with a simple scientific study. The Reactor Safety Study was a complex undertaking, with contributions from many individuals and disciplines. The data were often not available, and the techniques were often innovative. Further, peer review occurred at a time of great public concern about nuclear power.

The issues of the peer review process will be discussed below, and recommendations outlined. It should be noted here that the purpose of this discussion is not to review specific comments. In particular, we are, in this section, neither endorsing nor contradicting any specific peer comments which we use as illustrations of the process. Our concern is with the process.

a. Technical Peer Review Versus Other Levels of Peer Review

The primary objective of the peer review process, as contemplated by the Reactor Safety Study staff, was to obtain peer review by technically competent people of the validity of the estimates of probability and consequences of reactor accidents. Given the public temperament about nuclear power, it was inevitable that the comments received covered a wide range of other subjects bearing upon the acceptability of nuclear power. Though this would doubtless have occurred in any case, the Executive Summary and Chapters 6 and 7 of the Main Report all presented conclusions in a manner suggesting the acceptability of the risks estimated. Thus, the essential technical peer review process was corrupted, on both sides, by social commentary. Since the capability in the outside world for technical peer review was sparse to begin with, attention to these other areas drew meaningful effort away from the technical review.

b. Capability for Technical Review

The Reactor Safety Study took three years and millions of dollars to perform. Not only is the review of such a monumental undertaking a very difficult process, but it requires a large investment of resources. Further, and not surprisingly, a very large fraction of the national competence to perform review is employed in the reactor industry, where there were surely some inhibitions. In the end, there were few organizations which had resources for an adequate review, but many

responded, with varying levels of expertise and motivation. These included the Environmental Protection Agency, the American Physical Society Reactor Safety Study Group, the Atomic Industrial Forum, the Union of Concerned Scientists, the Electric Power Research Institute, and others.

The time for review between draft and final versions was short, considering the magnitude of the necessary effort. It has only been in the three-year period since the end of the study that many of the reviews have been completed, and a number of comments were only engendered and received during the testimony developed by this Review Group.

So far as we know, all reviewers had a major difficulty with the availability of information in WASH-1400, and the difficulty involved in tracing a line of thought through the study. We have alluded to that in our discussion of scrutability above, but it was a pervasive and difficult problem, crippling many efforts to accomplish responsible peer review.

c. Staff Response to Peer Review

The RSS staff took comments from many sources, and grouped similar comments together. While there are many critics who are aggrieved that their comments have been grouped with others, and there is indeed always a danger of loss of specificity in such a grouping, it is in general a valid approach, and the RSS staff developed generic answers to generic forms of criticism.

As mentioned in our introduction, the staff response to peer comments was often defensive. Thus, there are cases in which cogent comments from critics were either not acknowledged or evaded, leading to an adversary posture rather than the cooperative posture sought. (An egregious case in point is the lack of response to the frequent suggestion that summary figures in the Executive Summary be labeled "early fatalities" as opposed to "fatalities." While latent fatalities should surely not be simply added to early fatalities, it is equally incorrect to omit them). There are many other cases of inadequate staff response to responsible peer review*, and one peer reviewer has made the following statement to us in connection with his company's critique of WASH-1400:

*Examples are:

- a. In Figures 1-1 and 1-2 of the Executive Summary comparing reactor accident rates to other types of risks, a number of reviewers requested that the indication of the transition from experienced to extrapolated risks on the comparison risk curves be corrected. This was not done.
- b. In defense of the "adequacy of fault-tree methodology" in Appendix XI ("Analysis of Comments on the Draft WASH-1400 Report"), a figure was offered (p. 3-4) showing 50 data points in which "the ratio of observed failure rates to predicted failure

"A total of 45 deficiencies were found in WASH-1400 Draft. Of these 23 were considered in WASH-1400 Final in Appendix XI. Fourteen of these 23 were considered to be adequately covered in Appendix XI, and nine were judged not to be covered adequately. In addition to these 23, eight revisions occurred in WASH-1400 Final of which four are adequate and one is inadequate, while three resulted in an improvement. Thus, 34 of the original 45 deficiencies remain in WASH-1400 Final."

We do not argue that the RSS staff ought to have been automatically obedient to all comments received from the outside, and indeed many of these were beneath notice, but rather to say that the record of response to valid criticism was weaker than we would have liked.

When dealing with highly politicized subjects like reactor risk assessment, a major effort should be devoted in the beginning to separate technical review from social commentary, and adequate time should be allowed to permit technical review. In cases like RSS, where the accumulation of the data base and documentation takes so much of the effort, peer review will be facilitated where supporting documentation is maintained, and made available to responsible reviewers. While there is no question that this takes resources and staff time, it is essential to the value of peer review, and thereby to the credibility of the product.

rates (using fault-tree methodology) was within a factor of 4." It was not noted that the successful predictions shown were for simple systems such as meters and switches where the failure rate is determined by the failure rate of a single component whose failure rate is known.

- c. In response to comments concerning the potential seriousness of groundwater contamination, it was stated in WASH-1400 that, "If, for example, the receiving water body [of the contaminated groundwater flow] is a relatively small river with a flow rate of 13,000 cfs, the peak strontium-90 concentration in the river for the depressurization release will be ... well below maximum permissible concentration." (XI, 10-1) In fact, we find using the numbers in WASH-1400 (a release rate of more than 100 Curies of Sr⁹⁰ per day for a year according to Figure VII 3-6) and the then-maximum permissible concentration for public drinking water of 0.3×10^{-6} Ci/liter a contamination level exceeding the then-current maximum permissible concentration by a factor of 10.

6. EARTHQUAKES

An earthquake can be the cause of simultaneous failure of several unrelated parts of a nuclear power plant, thereby qualifying as a potential initiator of a common cause failure. RSS found earthquakes to be small contributors to the total risks of core melt in a nuclear power plant, partly because of the low probability that a severe earthquake will occur at the location of the plant, and partly from an estimated low probability of plant failure under earthquake stress. The latter resulted from the design and construction process that is intended to protect nuclear power plants from damage from earthquakes.

The peak acceleration associated with a large earthquake varies surprisingly slowly with the magnitude of the earthquake, while the peak horizontal displacement at a standard distance (the definition of the Richter magnitude) increases by approximately a factor of 10 per magnitude. Thus, large magnitude earthquakes are associated with large displacements and long periods, but the acceleration does not increase as rapidly. They do, however, make their effects known over a much wider region.

Earthquakes of magnitude five occur approximately 1,000 times per year, worldwide, but do not disturb large areas. For example, the Santa Barbara earthquake of August 13, 1978, of Richter magnitude 5.1, produced estimated horizontal accelerations of more than 0.2g at the home of the Chairman of this panel, approximately 10 kilometers from the epicenter. 20 kilometers further away, it was hardly felt. Such earthquakes are not a threat to a nuclear reactor as long as the reactor is placed far enough from an active fault to reduce the peak acceleration to a reasonable value, and is in fact built to withstand an acceleration of that value. Substantial safety factors are necessary (and are included in the licensing standards) because, in fact, the response of a structure is not entirely a function of the peak acceleration, but depends upon the whole spectrum of displacement to which it is subjected. This varies widely according to the magnitude of the earthquake, the location and structure of the fault, the local geology, etc., and it is only in recent years that detailed spectrum analyses of the displacement in the near field of substantial earthquakes has been possible.

The most detailed study of some of these issues as related to nuclear power plants has been made by Hsieh and Okrent*, who not only considered the detailed response of a plant to various levels of acceleration, but also considered the results of design and construction errors in a plant. These can greatly reduce the safety factors believed to be present, and must be considered.

*Hsieh, T. and Okrent, D., Some Probabilistic Aspects of the Seismic Risk of Nuclear Reactors, UCLA-ENG-76113, University of California, Los Angeles, Department of Energy and Kinetics. December 1976.

The RSS estimate which led to the conclusion that earthquakes do not contribute significantly to reactor risk is based on an analysis which led to the conclusion that two or more safety systems must fail in order to lead to core melt. (In this analysis an error occurred in the neglect of a combinatorial factor involving the number of such pairs.) Further, the probability of such simultaneous failure of independent systems was calculated by the incorrect square root method to which we have previously referred. In fact, an earthquake subjects the entire plant to stress, and can cause simultaneous failures in uncorrelated systems.

We believe that the entire subject of earthquake risk in nuclear plants deserves much more attention than it has received. It may be that not enough is truly known about the source mechanisms and spectra to make possible truly sophisticated calculations of the threat to a given plant from a given fault, but we believe the treatment can be improved. In particular, it is surely wrong to characterize a large earthquake as a point source at the location of the nearest fault, since it is well known that really large earthquakes involve the release of energy along large sections of a fault. WASH-1400 concludes that earthquakes do not contribute significantly to the risk attached to operation of nuclear power plants. Other calculations (by Hsieh and Okrent), based upon different assumptions, lead to the conclusion that they might contribute significantly. Investigation at a level of sophistication necessary to resolve this issue to our satisfaction has not yet been performed, and it is important that it be done.

7. RISK PERCEPTION

Public perception of risk in nuclear power plants is affected by many forces, only one of which is the process of risk estimation and evaluation that is responsive to a regulatory program. The latter has as its objective the identification and mitigation of problems, both short- and long-term, but is nonetheless part of the environment that leads to public awareness of risk.

The purpose of WASH-1400 was to make a "realistic" determination of the probability and consequences of reactor accidents, especially where there might be public consequences. In this context, the term "realistic" meant the best estimate of probabilities and consequences, without the introduction of systematic biases to provide either conservatism or optimism. It was further intended to be equally objective in the presentation of the results.

There were a number of places in which the study fell short of its goal of objectivity. Some were understandable, in that a good part of the analysis involved situations for which there was no adequate data base and required value judgments on the part of the RSS staff. (Some were not properly identified.) Further, the Executive Summary and Chapters 6 and 7 of the Main Report touched on a broader issue, the acceptability of risk. Thus there were problems associated with presentation and interpretation of the results.

a. Definition of Risk

The RSS team did not choose a specific definition of risk, but displayed its results through graphs of the probability of occurrence of an event against the consequences of that event. Other means of displaying the results could produce different risk perceptions. The perception of risk is a very complex subject, though in common usage, risk means simply the potential of an undesirable event. A quote from Keynes highlights this problem:

"In estimating the risk, 'moral' or 'physical,' it must be remembered that we cannot necessarily apply to individual cases results drawn from the observation of a long series resembling them in some particular. I am thinking of such arguments as Buffon's when he names 1/10,000 as the limit, beyond which probability is negligible, on the ground that, being the chance that a man of fifty-six taken at random will die within a day, it is practically disregarded by a man of fifty-six who knows his health to be good. If a public lottery, Gibbon truly pointed out, were drawn for the choice of an immediate victim, and if our name were inscribed on one of the ten thousand tickets, should we be perfectly easy?"

b. Influence of Judgments on Risk Estimates

Objective probabilities are determined by measurement in repeated trials, and objective consequences are those that are directly measurable. Neither were available for the large consequence, low probability events to be evaluated by the Reactor Safety Study team, so recourse to synthesized probabilities, synthesized objective consequences or extrapolation was necessary. The validity of such efforts is often difficult to verify. For example, experience in failure of nonnuclear pressure vessels could be different from cases in which high pressure and high radiation flux occur simultaneously. The estimates of scram failure probability in Anticipated Transients Without Scram is another case in point. (See our later discussion in Section 11.)

The use of synthesized probability and extrapolation is justifiable. However, the estimates are no longer 'objective' according to our use of the term, and all of the value judgments must then be made visible and the scope of the uncertainties discussed and estimated.

c. Acceptable Levels of Risk for Nuclear Energy

After estimating the probability and consequences, the Reactor Safety Study, in the Executive Summary and Chapters 6 and 7 of the Main Report, compared estimated risks with the probability and consequences of other societal risks.

Determination of acceptability of risk is a separate matter from objective estimation of risk. In this case, behavior toward risk must be measured or the perception of risk ascertained. People act on their subjective perception of risk, which may be altered by providing clarifying information if such information is presented in a manner deemed credible.

The estimation of the probability and consequences of reactor accidents exceeding the design basis is only one of the risks in the reactor systems and the total fuel cycle. To judge acceptability of nuclear reactors solely on the risk of early fatalities, and latent health effects, and property damage for Class 9 accidents is inappropriate, though this is an important ingredient. All the issues associated with the nuclear fuel cycle are important, including economic and environmental matters, and weapons proliferation. Further, public perception hangs heavily upon the perceived credibility of the Nuclear Regulatory Commission.

8. ALLEGATIONS BY UCS CONCERNING RSS TREATMENT OF QUALITY ASSURANCE/
QUALITY CONTROL*

The RSS critique by the Union of Concerned Scientists focusses considerable attention on the way the RSS treated the issue of quality assurance/ quality control (QA/QC). This issue is discussed by UCS partially "in an effort to see the internal workings of the RSS and, in particular, to learn how the study handled sensitive problems that could, if unfavorably resolved, call into question the overall safety of commercial nuclear power plants."† Thus, UCS says that its study of the QA/QC issue is a "case study."

The essence of the UCS critique centers around two issues: first there are two memoranda sent by Mr. E. Gilbert of the study team, in which Gilbert outlines for the consideration of the team leaders different approaches to the RSS treatment of the QA/QC problem. The second is an allegation that the RSS suppressed the inclusion of an entire section on the subject because, presumably, it would have been damaging to the case the study team is accused of trying to make.

Basically, Mr. Gilbert in his memoranda outlines for the team leaders the importance of the QA/QC issue in the context of the larger RSS study. It appears that, at a meeting of several RSS team members, two possible approaches to the issue were considered. Gilbert discusses these approaches.

In the first postulated approach ('Approach A'), the study would attempt only to find those data supportive of the general conclusion that QA/QC was performed adequately by the reactor vendors, licensees, and the AEC itself. In this approach, a critical look at the QA/QC issue would not be taken, but instead documentation would be sought to engender confidence in the ongoing QA/QC program. In the second approach ('Approach B'), the QA/QC findings of RSS would not be predetermined, but rather the RSS team would consider the deficiencies and nonconformities actually found, on an audit basis, in the field. Gilbert notes his preference for using Approach B "to reach a finding (supportable by facts gleaned from the AEC records) that all major QA/QC deficiencies were uncovered by the AEC audit and that they were repaired and/or resolved in a manner which assures plant safety system reliability."

Gilbert then goes on to list the advantages and disadvantages of the two approaches. High among the disadvantages of Approach B is the problem that the field data from

*F.v.H. dissociates himself from this section.

†Union of Concerned Scientists, The Risks of Nuclear Power Reactors: A Review of the NRC Reactor Safety Study WASH-1400 (NUREG-75/014) (Cambridge, MA: August 1977), pp 159-160.

audits show "generally only negative findings" because that is the nature of any audit report, and hence "the objective reader can question whether the whole story is not actually much worse...."

The allegation made by the UCS report does not draw its own conclusion explicitly, leaving the conclusion mainly to the reader. However, Mr. Daniel Ford, Executive Director of UCS, in a presentation to our Risk Assessment Review Group, directly alleged that this incident represented deliberate dishonesty on the part of the RSS team and urged an investigation.

Mr. Ford indicated "that the study group deliberately sought to exclude the collection of information that it believed had the potential to show either the unreliability of reactor safety systems or the indeterminacy of that question." What is objectionable in this matter, Mr. Ford said, is "the fact that it is completely contrary to any objective scientific analysis of how you analyze quality assurance and how you fit it into the state of the art report on reactor safety." Dr. Lewis said he believed that Mr. Ford was saying that they dishonestly closed their minds to available data that might have brought to question the safety of nuclear reactors. Mr. Ford said that was correct.

Further inquiries by members of our Review Group, including interviews and correspondence with Mr. Gilbert, have revealed nothing additional in the files except a memo by Gilbert that appears to carry out the recommendations of his two earlier memoranda. It is important to note that, after having raised the two possible approaches, Gilbert himself recommends Approach B as the appropriate strategy. Approach B is intrinsically capable of arriving at an honest conclusion, whereas Approach A is not. Thus, the outcome of the incident is acceptable on its face; the dishonest Approach A was rejected. This does not completely close the issue, however, because there is still the problem of how the RSS study team could have come to the point of even giving serious consideration to Approach A, or to anything like it! Such consideration is difficult to comprehend and deserves strong disapprobation.

There is a distinction between "suppression" and "omission" of material. On the issue of the alleged "suppression" of an extensive QA/QC section in the RSS final report, our group finds that its omission was one among numerous examples in which material was omitted for a variety of reasons. Mr. Gilbert in a memorandum to our panel, offered his explanation, as follows: "...pages 180-212 are possibly the most interesting parts of Document 1 [the complete (but undated) write-up of the PWR/BWR QA/QC procedures] because they summarize both the competence and attention to detail of the government inspectors and their technical consultants. Their performance should engender public confidence in the effectiveness of government Q-A audits and corrective procedures. Unfortunately, their findings can also be used to support the worst fears and insecurities of the professional intervenors. The facts are indisputable. But we can also see how easily these same facts could have been misused in damaging ways. I think Rasmussen/ Levine made the right decision (a tough one) in not including this material in the WASH-1400 report."

The RSS study team has explained this omission as due to the proposed section's being not germane to the particular RSS task for which it was developed. On the general issue of "omissions," we refer the reader to the Section of this report on scrutability.

We consider allegations of dishonesty to be deserving of the most serious and careful study; each must be backed up by evidence commensurate with the seriousness of the charge. It is understandable that an incident such as this should raise doubts among the skeptical as to the overall honesty of the RSS effort. However, this allegation of dishonesty is without merit in our view. Furthermore, we do not accept the argument that this incident can be used as a 'case study,' generalizable to reveal lessons as to the intellectual integrity of the larger RSS effort.

9. CURRENT ROLE OF PROBABILISTIC METHODS IN THE REGULATORY PROCESS

The Reactor Safety Study is used in several ways in the regulatory process. Use is made of its methodology, of whole parts of its results, and of generic applications of the methodology.

Use of the techniques was anticipated in several applications made before the Study was completed. A probabilistic study of aircraft crashes has led to a criterion limiting acceptability of sites for nuclear power plants to areas far enough from airports. The earlier reviews (1972-73) of anticipated transients without scram (ATWS) started from a requirement that these events have a probability of less than 10^{-6} per reactor year, and assessed the ability of reactor designs to meet this. Other uses of probabilistic methods have been made with varying success in increasing numbers since issuance of the Reactor Safety Study. A few examples are: more refined methods of seismic analysis based on probability of occurrence of earthquakes and probability that structures can withstand assumed earthquake stresses, performance of backup power systems, the effects of overpressure and other transients on reactor pressure vessels, the probability of piping failure, and the effects of missiles generated by breakup of a turbine.

The radiological dispersal and effects models have been used in such ways as for further investigation of liquid pathways, both through streams and groundwater, for analysis of relative suitability of alternate reactor sites, and for some tentative studies of the relative hazards of electrical power generation from coal and fission.

Generic studies have been conducted using the event tree methodology of the Reactor Safety Study. In these cases, where specific designs are not evaluated, the probabilities associated with branches of the event trees are either estimated from general considerations or borrowed from the results of fault tree analysis in the Reactor Safety Study. Generic analysis has been done for fast breeders, for floating nuclear power plants, for safety of spent fuel shipping containers, and for spent fuel processing plants, among others.

There has been consideration of use of the Reactor Safety Study itself, its refinements, and conclusions based on these in revising the basis of part of the regulatory process. One example considered in this connection is the analysis required by 10 CFR 100 (Reactor Site Criteria), whose philosophical basis is partly a view that reactor core meltdown should be no more probable than 10^{-6} per reactor year, because any meltdown would be catastrophic. The conclusions of the Reactor Safety Study disagree with both the probability of such an event and its effects, and so a revision of the site criteria themselves is suggested. No specific steps in this direction have yet been taken.

10. ACTS OF VIOLENCE: SABOTAGE AND WAR

Acts of violence that could damage nuclear power plants and as a result endanger people in the area include sabotage and war. In either case, there would be a deliberate plan to produce damage and perhaps subsequent danger to people. The risk from sabotage was not calculated in the Reactor Safety Study. The omission was deliberate, and proper, because it was recognized that the probability of sabotage of a nuclear power plant cannot be estimated with any confidence.

The Study draws on other analyses to conclude that features of nuclear power plants that protect against risks from accidents also make these plants more difficult to sabotage successfully. These features include heavy shielding, thick-walled pipes and vessels, features that cause shutdown of the chain reaction and continuance of core cooling if necessary, restraints on piping and components to provide resistance to earthquakes, and other engineered safety features.

The same analyses are cited as basis for a conclusion that the worst consequences from acts of sabotage should not be expected to exceed the worst consequences from accidents. It was also concluded that nuclear power plants appear less susceptible to sabotage than most other civil or industrial targets. Nonetheless the NRC has, in recent years, implemented stricter security standards, which should further reduce the risk.

It is also worth noting that some features that make the plants difficult to sabotage successfully would also help to make most sabotage scenarios benign from the standpoint of public injury. It would be much easier to sabotage the plant so as to lead to shutdown and a need for expensive repair than to produce a core melt.

We do not mean to imply by the above that sabotage is not still an important issue

We turn to acts of violence in war. Of course, the direct consequences of a nuclear war would dwarf any secondary effects from damaged nuclear power plants. The possibility of conventional war varies over the globe, probably lower in nuclear-armed countries.

Over much of the world, the possibility of conventional wars cannot be neglected. In these areas, damage to nuclear power plants that may be present could add substantially to the public risk. In fact, it appears that nuclear power plant design and siting practices should be different for those parts of the world in which conventional wars cannot be ruled out. In these cases, different studies using the methodology of the Reactor Safety Study could be helpful in rational design and siting criteria.

11. ATWS (ANTICIPATED TRANSIENTS WITHOUT SCRAM)

The analysis in WASH-1400 of failure of a reactor to scram on demand has caused considerable controversy. It was in considering this event that the practice of averaging two very disparate estimates by taking their geometric mean was most deplorable, because the estimated upper and lower bounds on probability of common cause failure were a factor 10^6 apart, and the lower of the two has no relevance to the probability of scram failure.

The ATWS events were included among all events started by transients. WASH-1400 concluded that all transient events combined were a less significant source of risk for PWRs and were the dominant source of risk for BWRs. These conclusions depended very much on the estimated probabilities of failure to scram.

The Division of Systems Safety of NRC has since issued NUREG-0460,* which provides a position on ATWS which it is proposed that NRC adopt in future safety applications. NUREG-0460 develops its argument through dependence on purely actuarial information on common mode failure probability to scram. All new relays in the scram circuit of the Kahl reactor in Germany were found at one point to have a scram defect due to inability to open on spring action when the current is cut off, because the protective plastic coating had not been cured properly. This one statistic, along with an estimate of the probability that a severe transient requiring scram might take place while the inability to scram persisted, led to a calculated frequency of anticipated transients without scram of about 2×10^{-4} per year for a water reactor.

Without further information it might be argued that the one occurrence might be an accidental example of a member of a large set of very rare events, as we discuss in our comments on the Browns Ferry fire. However, NUREG-0460 points out that a similar problem appeared (and was corrected) in preoperational testing at the Monticello nuclear plant in the United States.

The method used in NUREG-0460 is clearly an improvement over that used in WASH-1400. But there are some additional points that must be recognized.

As NUREG-0460 points out, the two plants analyzed in WASH-1400 were not typical regarding ATWS. The PWR was of a class that is substantially protected against ATWS, and it is likely that the risk of core melt for this PWR is even less than WASH-1400 concludes it

*U.S. Nuclear Regulatory Commission, Anticipated Transients Without Scram for Light Water Reactors: Staff Report by the Division of Systems Safety of the U.S. NRC, NUREG-0460, Volumes 1 and 2, April 1978. Available from the National Technical Information Service, Springfield, VA 22161.

to be. Many other PWRs are believed to be more subject to damage from ATWS. The BWR that was analyzed was instrumented to undergo recirculation pump trip if ATWS occurs, and this reduces the severity of the first pulse of energy in the transient. Not all BWRs have this feature. A more fundamental problem is found for the period subsequent to the first pulse. The liquid boron poison system designed to stop the chain reaction might not act fast enough, and the effectiveness assumed by WASH-1400 might not be assured.

Therefore NUREG-0460 recommends measures to be taken for all BWRs and some PWRs, to mitigate the effects of an ATWS event should they occur.

These points are relevant to the question of how far the results of WASH-1400, even if amended to account more correctly for the contribution of ATWS to risk, can be extrapolated to a full nuclear industry. The authors of WASH-1400 stated concern about this extrapolation, but made it nonetheless. To understand the question of risk from this source, it is necessary to update and broaden this risk assessment. It would appear from the conclusions of NUREG-0460 that the risk from ATWS is higher than was estimated in WASH-1400.

12. THE INFLUENCE OF DESIGN DEFECTS AND QUALITY ASSURANCE FAILURES

The history of failures of other complex and presumably safe systems is heavy with instances of design defect and quality assurance failures. (A spectacular case in point is the collapse of the Tacoma Narrows Bridge.) Such defects manifest themselves outside the framework of the RSS methodology when they do not contribute to the component failure data base. Thus, for example, the Tacoma Narrows Bridge failed in a torsional mode due to the generation of Von Karman vortices in a reasonably low wind, but there are many other suspension bridges that would fail in this mode if the winds were to get much higher. Similarly, it was the failure of the baggage door on the DC-10 airplane which buckled the interface between the passenger and baggage compartments, revealing that the hydraulic lines were routed along just that bulkhead. (By contrast, a 747 which hit the runway lights at San Francisco had two of its three redundant hydraulic systems disabled by the impact, but adequate separation criteria preserved the third and led to a safe landing.)

We do not know how important a contributor of this sort would be to the probability of core melt, but do know that it is a fact of life for a large complex system. For example, any of the following could become important, and they are listed here as eclectic samples of the sort of thing we have in mind:

- (1) The net positive suction head for the Containment Recirculation Spray and Low Head Safety Inspection System pumps at the PWR analyzed in WASH-1400 (Surry) was recently found to be not as great as specified by the pump manufacturer. (NUREG-0090-10, Report to Congress on Abnormal Occurrences, Oct-Dec. 1977, p. 5.)
- (2) Some safety related electrical equipment in reactor containments has been found not to be adequately resistant to the environmental conditions (steam, heat) which would prevail in the containment following a loss of coolant accident. (Ibid, p. 9.)
- (3) Steam generator tubes have been failing at Surry #2 and other reactors as a result of intergranular stress corrosion cracking. (The failure of a number of steam generator tubes as a result of the pressure transients during a loss of coolant accident could render the emergency core cooling system ineffective. NUREG-0900-5, Report to Congress on Abnormal Occurrences, July-September 1976, p. 6.)
- (4) There was a violation at Browns Ferry of NRC's separation requirements which increased the likelihood of the common cause failure of redundant systems. (NRC, Notice of Violation, July 28, 1975.)

- (5) The electrical interconnection of safety-related and nonsafety-related systems, which has become an issue at a number of power plants.

13. CALCULATION OF POPULATION DOSES FROM GIVEN RELEASES OF RADIONUCLIDES

The ranges of uncertainty given in the WASH-1400 consequence calculations are given as a factor of 4 for early fatalities and illnesses, and factors of 1/6 and 3 for cancer deaths and genetic effects. These ranges are supposed to include uncertainties in atmospheric dispersion calculations, the effectiveness of evacuation, exposure-dose coefficients, and dose-biological effect coefficients. The uncertainties associated with most of these factors were not estimated in WASH-1400, so that sensitivity to them should be explored. Examples are:

a. Atmospheric Dispersion Calculations

Some exploratory analysis* that has been reported to us indicates that the calculated number of early fatalities as a function of distance is quite sensitive to the model for the plume rise and the value of the deposition velocity.

b. Effectiveness of Evacuation

Most of the short term fatalities in WASH-1400 are associated with "ineffective evacuation." This corresponds to the assumption that evacuation does not occur before the arrival of the radioactive cloud but that most of the population is removed within 4 hours thereafter (VI - 13-35). It is assumed in WASH-1400 that such ineffective evacuation will occur one-third of the time. It appears quite possible that the WASH-1400 ineffective evacuation scenario represents an average rather than a pessimistic expectation, in which case the early fatalities would be increased by a factor of about three.

c. Exposure-Dose Coefficients

In Draft WASH-1400, the inhalation dose was the major contribution to early fatalities (VI - 84). In the final report, both the inhalation exposure-dose coefficients and the dose-consequence coefficients have been reduced to the point where the inhalation dose is a minor contributor to early fatalities. The Draft used a 1959 ICRP report (#2) as its source for exposure-dose coefficients. The final report includes a new set of coefficients which was apparently developed specifically for the WASH-1400 group.

To our knowledge, the new exposure-dose coefficients have not been reviewed by either the NCRP or the ICRP. It is not yet clear therefore what their level of acceptance is.

*Jan Beyea, A Study of Some of the Consequences of Hypothetical Reactor Accidents at Barseback, Swedish Energy Commission, 1978.

THE RISK ASSESSMENT REVIEW GROUP

Since the publication of the Reactor Safety Study in 1975, the Study has been the subject of much debate. Congressman Morris K. Udall, Chairman of the Committee on Interior and Insular Affairs of the U.S. House of Representatives, held hearings on the Study in June 1976. Subsequently, the Staff of the Subcommittee on Energy and the Environment published in January 1977 a report entitled Observations on the Reactor Safety Study. On March 14, 1977, on the basis of that report, Congressman Udall requested that the Nuclear Regulatory Commission convene a review group reporting directly to the Commission, including persons not employed by NRC and who have been critical of the Study, for the purpose of preparing a new Summary of the Report.

The Commission responded to Congressman Udall on April 4, 1977, agreeing that a further undertaking was in order but of a somewhat different nature. The Commission expressed the view that what was called for was an independent review of peer comments on the WASH-1400 final report and of recent developments in risk assessment methodology. The Commission commented that such a review would assist the Commission in establishing future policy regarding the use of risk assessment in the regulatory process.

The Commission then began efforts to convene such a Panel. With Commission approval, Mr. Saul Levine, Director of Nuclear Regulatory Research, contacted Dr. Harold W. Lewis to seek his availability to serve as Chairman of the Review Group. Upon Dr. Lewis' acceptance, he and Mr. Levine redrafted the charter for the Review Group and prepared a recommended membership list which included some who have been critical of the Study, some who have praised it, and some who have not taken a specific position on it. On June 30, 1977, after consultation with the Director of OMB, as required by the Federal Advisory Committee Act, the Commission approved the proposed charter (see Annex II) and the membership of the Review Group. The members and a brief description of their backgrounds are as follows:

- Harold W. Lewis, Professor, Department of Physics, University of California, Santa Barbara, California. Chairman of the American Physical Society study group on light-water reactor safety, 1975.
- Robert J. Budnitz, Director, Energy and Environment Division Lawrence Berkeley Laboratory, University of California, Berkeley, California. Member of American Physical Society study group on light-water reactor safety, 1975.
- Herbert J.C. Kouts, Chairman, Department of Nuclear Energy, Brookhaven National Laboratory, Upton, New York. Former Director of the NRC Office of Nuclear Regulatory Research (1975-1976), former Director of AEC's Reactor Safety Research Division (1973-1975).
- Walter B. Loewenstein, Director, Safety and Analysis Department, Nuclear Division, Electric Power Research Institute, Palo Alto, California.
- William D. Rowe, Deputy Assistant Administrator for Radiation Programs, U.S. Environmental Protection Agency, Washington, D.C.

- Frank von Hippel, Research Scientist, Center for Environmental Studies, Princeton University, Princeton, New Jersey. Member of American Physical Society study group on light-water reactor safety, 1975.
- Fredrik Zachariasen, Professor, Department of Physics, California Institute of Technology, Pasadena, California.

On August 7, 1978, Dr. Budnitz assumed a position as Deputy Director of the Office of Nuclear Regulatory Research, but Dr. Lewis recommended to the Commission that Dr. Budnitz complete his work with the panel.

The group reported its findings and recommendations to the Commission on September 7, 1978.

THE CHARTER

The Review Group will provide advice and information to the Commission regarding the final report of the Reactor Safety Study, WASH-1400, and the peer comments on the Study, advice and recommendations on developments in the field of risk assessment methodology and on future courses of action which should be taken to improve this methodology and its application. This advice and information will assist the Commission in establishing policy regarding the use of risk assessment in the regulatory process, in improving the base for the use of such assessments. It will also clarify the achievements and limitations of the Reactor Safety Study.

CORRESPONDENCE REGARDING FORMATION
OF THE REVIEW GROUP

MORRIS K. UDALL, ARIZ., CHAIRMAN

- 55 -

COMMITTEE ON INTERIOR AND INSULAR AFFAIRS
U.S. HOUSE OF REPRESENTATIVES
WASHINGTON, D.C. 20515

CHARLES CONKLIN
STAFF DIRECTOR

LEE MC ELVAIN
GENERAL COUNSEL

STANLEY E. SCOVILLE
SPECIAL COUNSEL

MICHAEL C. HARDEN
MINORITY COUNSEL

March 14, 1977

PHILIP BURTON, CALIF.
WALTER W. MASTENBERGER, WIS.
BOB STAFFORD, WASH.
JAMES H. KAUFMAN, JR., TEX.
TERRY INNIS ALD, WYO.
JOHNATHAN B. BISHOP, N.Y.
JOHN P. SEIBERLING, OHIO
HAROLD RUNNELS, N. MEX.
ANTONIO BORJA WON PAT, GUAM
RON DE LUGO, V.I.
BOB ECKHART, TEX.
GOODLOE E. BYRON, MD.
JIM SANTINI, NEV.
PAUL E. TSONCAS, MASS.
JAMES WEAVER, OREG.
BOB CARR, MICH.
GEORGE MILLER, CALIF.
THEODORE M. (TED) RISENHOOVER,
OKLA.
JAMES J. FLORIO, N.J.
DAWSON MATHIS, GA.
PHILIP R. SHARP, IND.
MATTHEW P. MC HUGH, N.Y.
JOHN KRESS, CALIF.
EDWARD J. MARKEY, MASS.
PETER H. KOSTMAYER, PA.
BALTASAR CORRADA, P.R.
AUSFIN J. MURPHY, PA.
NICK JOE RAHALL II, W.VA.
DICKIE F. VENTO, MINN.
JERRY HUCKABY, LA.
LAMAR GUDGER, N.C.

JOE SKUNTZ, KANS.
DON H. CLAUSEN, CALIF.
PHILIP E. RUPPE, MICH.
MANUEL LUIJAN, JR., N. MEX.
KEITH G. BERREBUS, KANS.
DON YOUNG, ALABAMA
ROBERT E. BAUMAN, MD.
STEVEN D. SYMMS, IDAHO
JAMES P. (JIM) JOHNSON, COLO.
ROBERT J. LAGOMARSINO, CALIF.
DON MARFROTT, UTAH
RON MARLENEE, MONT.
ELDON RUDD, ARIZ.
MICKEY EDWARDS, OKLA.

The Honorable
Marcus A. Rowden, Chairman
Nuclear Regulatory Commission
Washington, D.C. 20555

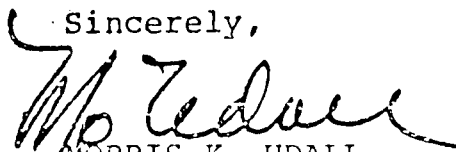
Dear Marc:

I am enclosing herewith a Subcommittee staff report entitled Observations on the Reactor Safety Study. This report is based on information presented at, and subsequent to, the Subcommittee's June 11, 1976, hearing.

You will note from the report that there exists a widespread belief that the Reactor Safety Study was presented in a manner which created a misleading impression of the certainty and comprehensiveness of its conclusions. I believe this misimpression should be corrected by preparation of a new Summary.

Accordingly, I am requesting that you convene a review group for the purpose of preparing a new Summary. I believe that the panel should report directly to you and should include persons not employed by the NRC and who have been critical of the Study. (I would hope that such persons, if eligible, are reimbursed for their services.)

Sincerely,


MORRIS K. UDALL
Chairman



56
UNITED STATES
NUCLEAR REGULATORY COMMISSION
WASHINGTON, D. C. 20555

April 4, 1977

OFFICE OF THE
CHAIRMAN

The Honorable Morris Udall, Chairman
Subcommittee on Energy and the Environment
Committee on Interior and Insular Affairs
United States House of Representatives
Washington, D.C. 20515

Dear Mr. ^{Mo} Chairman:

Thank you for your letter of March 14, enclosing a report, "Observations on the Reactor Safety Study," and requesting that a review group be convened for the purpose of preparing a new summary of the study. The Commission believes that a further undertaking is in order, but of a somewhat different nature.


We recognize that opinions can reasonably differ as to the best means of summarizing so complex and extensive a report as WASH-1400. We do not believe, however, that preparation of a new summary to WASH-1400 would be a productive undertaking. The Executive Summary was written by the group which prepared the Reactor Safety Study, to be a part of that study; the fact that it represented their views will remain unaltered, irrespective of whether a second summary is prepared by a different group. More importantly, in our view, a revision of the Executive Summary at this time would obscure the important point that the Reactor Safety Study as a whole should be seen as one step, already accomplished, in the continuing development of risk assessment methodology. We believe that what is called for at this time is an independent review of peer group comments on the WASH-1400 final report and of developments in risk assessment methodology that have occurred since the report was published. Such a review, together with WASH-1400 itself and the extensive efforts on risk assessment presently underway within NRC, could be valuable to the NRC in establishing future policy regarding the use of risk assessment in the regulatory process.

We intend to take steps to convene a review group for such a purpose. We agree the group should report directly to the Commission, and should include a diversity of views. We believe that the group should, as you suggest, include persons who have been critical of the study. We expect

that the necessary detailed arrangements for this effort will take about a month; therefore, we plan to be back in touch with you about May 1 with the names of the members of the review group as well as a statement of its charter.

We hope that the results of this effort will contribute to the ongoing discussions regarding the usefulness and applicability of the study's findings.

Sincerely,


Marcus A. Rowden
Chairman

NINETY-FIFTH CONGRESS

MORRIS K. UDALL, ARIZ., CHAIRMAN

PHILLIP BURTON, CALIF.
 ROBERT W. KASTENMEIER, WIS.
 LLOYD MELOS, WASH.
 ABRAHAM RAZIN, JR., TEX.
 TENO RONCALIO, WYO.
 JONATHAN B. BINGHAM, N.Y.
 JOHN F. SEIBERLING, OHIO
 HAROLD RUNNELS, N. MEX.
 ANTONIO BORJA WON PAT, GUAM
 RON DE LUGO, V.I.
 BOB ECKHARDT, TEX.
 GOODLOE E. BYRON, MD.
 JIM SANTINI, NEV.
 PAUL E. TSONGAS, MASS.
 JAMES WEAVER, OREG.
 BOB CARR, MICH.
 GEORGE MILLER, CALIF.
 THEODORE M. (TED) RISENHOOVER,

JOE KAUBITZ, KANS.
 DON H. CLAUSEN, CALIF.
 PHILIP E. RUPPE, MICH.
 MANUEL LUJAN, JR., N. MEX.
 KEITH G. SEBELIUS, KANS.
 DON YOUNG, ALASKA
 ROBERT E. BAUMAN, MD.
 STEVEN D. SYMMS, IDAHO
 JAMES P. (JIM) JOHNSON, COLO.
 ROBERT J. LAGOMARSINO, CALIF.
 DAN HARRIOTT, UTAH
 RON MARLENEE, MONT.
 ELDON RUDD, ARIZ.
 MICKEY EDWARDS, OKLA.

OKLA.
 JAMES J. FLORIO, N.J.
 DAWSON MATHIS, GA.
 PHILIP R. SHARP, IND.
 MATTHEW F. MCMUGH, N.Y.
 JOHN KREBS, CALIF.
 EDWARD J. MARKEY, MASS.
 PETER H. KOSTMAYER, PA.
 BALTASAR CORRADA, P.R.
 AUSTIN J. MURPHY, PA.
 NICK JOE RAHALL II, W. VA.
 BRUCE F. VENTO, MINN.
 JERRY MUCKABY, LA.
 LAMAR GUDGER, N.C.

- 58 -

COMMITTEE ON INTERIOR AND INSULAR AFFAIRS
 U.S. HOUSE OF REPRESENTATIVES
 WASHINGTON, D.C. 20515

CHARLES CONKLIN
 STAFF DIRECTOR
 LEE MC ELVAIN
 GENERAL COUNSEL
 STANLEY SCOVILLE
 SPECIAL COUNSEL
 LOUIS STRIEGEL
 MINORITY COUNSEL

June 13, 1977

The Honorable
 Marcus A. Rowden, Chairman
 Nuclear Regulatory Commission
 Washington, D.C. 20555

Dear Marc:

The June 9 issue of Nucleonics Week contains a reference to a report being prepared at the NRC on the matter of the Reactor Safety Study. This reference indicated the possibility of some disagreement concerning the purpose of the review.

As you recall, I requested on March 14 that you convene a group for the purpose of preparing a new summary of the Study. This request was based on the belief that the existing summary had created a misleading impression concerning the certainty and comprehensiveness of the Study's conclusions.

Your response to my requests indicated that you did not believe preparation of a new summary would be a productive undertaking. While I do not agree with your assessment of the utility of a revised summary, I will not contest the matter. However, I do believe it essential that the proposed review focus primarily on the issue of what the Reactor Safety Study does and does not tell us about reactor safety. The review should contain a clear statement of the group's view of the accuracy and limitations of the Study's methodology and conclusions.

Sincerely,



MORRIS K. UDALL
 Chairman



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

- 59 -

June 17, 1977

OFFICE OF THE
CHAIRMAN

The Honorable Morris K. Udall, Chairman
Subcommittee on Energy and the Environment
Committee on Interior and Insular Affairs
United States House of Representatives
Washington, D. C. 20515

Dear Mr. ^{Mo}Chairman:

Thank you for your letter of June 13, concerning the mission of the review group which the Commission will establish to review peer comments on the WASH-1400 final report and developments in risk assessment methodology. The Commission has been engaged, as I indicated in my letter of May 9, in defining the makeup and charter of the group. In addition, we are undertaking consultation with the Director of OMB as required by the Federal Advisory Committee Act, 5 U.S.C. App. I, §9, prior to formal creation of the group. We are pleased to provide herewith that information.

Professor Harold Lewis of the University of California has agreed to serve as Chairman of the group. You will recall that Professor Lewis was Chairman of the American Physical Society's Reactor Safety Study Group in 1975 and rendered distinguished service in that effort. Enclosure 1 is a list of the members of the group. We feel that this membership provides a knowledgeable, balanced group which should produce a useful report.


Enclosure 2 is a draft Review Group charter, which provides the details of the group's mission and lists the areas which are to be covered in the report. We believe that this charter is responsive to your concerns. A more detailed charter will subsequently be prepared in accordance with the provisions of the Federal Advisory Committee Act.

Professor Lewis has planned for the group to be finished in its work approximately six months from its first meeting, and for the report to be submitted directly to the Commission shortly thereafter. This report together with WASH-1400 itself and the extensive efforts on risk assessment presently underway within NRC should be valuable in establishing future policy regarding the use of risk assessment in the regulatory process.

Ref: SECY-77-301

We look forward to a useful and productive report from the Review Group.
If you would like any further information, please let me know.

Sincerely,


Marcus A. Rowden
Chairman

Enclosures:

1. Membership of WASH-1400
Review Group
2. Review Group Charter

cc: Review Group Members

MEMBERSHIP OF RISK ASSESSMENT REVIEW GROUP

Dr. Harold W. Lewis (Chairman)
University of California - Santa Barbara

Dr. Walter Lowenstein
Electric Power Research Inst.

Dr. Frank Von Hippel
Princeton University

Mr. Robert J. Budnitz
Lawrence Berkeley Laboratory
University of California

Dr. William D. Rowe
Environmental Protection Agency

Dr. Fredrik Zachariassen
California Tech.

Dr. Herbert J. C. Kouts
Brookhaven National Laboratory

REVIEW GROUP CHARTER

The Commission believes that it would be useful to convene a special Review Group for the purpose of reviewing peer group comments on the final report of the Reactor Safety Study and the developments in risk assessment methodology that have occurred since the report was published.

The group should examine peer comments as they pertain to the scope and uncertainties involved in the final report of the Reactor Safety Study and should make recommendations as to additional work that should be done to improve risk assessment methods and applications. Peer group comments on the final report of the Reactor Safety Study include testimony presented to Congressional Committees and the Advisory Committee on Reactor Safeguards, letters to NRC, technical articles and such other material as may come to the group's attention. The review of developments in risk assessment methodology should cover reports of work being done in the field, and NRC's plans for the further development and application of risk assessment techniques. The report of the group should be directed toward aiding the NRC to establish policy regarding the use of risk assessment in the regulatory process, to improve the base for the use of such assessment, to formulate programs for the reduction of uncertainties in risk assessments, and to clarify the achievements and limitations of the RSS.



DEPARTMENT OF PHYSICS

The Honorable Morris K. Udall, Chairman
Committee on Interior and Insular Affairs
U.S. House of Representatives
Washington, D.C. 20515

Dear Mr. Udall:

I have a copy of your letter of June 13 to Mr. Rowden, referring to NRC's putative review of the Reactor Safety Study. Since I'm scheduled to chair that study on behalf of NRC, and since I assume that we all have the same objective — the pursuit of truth — I hope that I may be allowed a couple of comments.

When I was first asked to do this job, I said that I felt that preparation of a new summary (defective as I believe the present one to be) would be a futile gesture, not worth the time, but that there was a great deal that could be done, in the public interest, to clarify the achievements and limitations of the Reactor Safety Study and the potential role of such methodology in the regulatory process. At that time I communicated my views to Henry Myers, of your staff, but I did want to put these essentials in writing, lest there be any misunderstanding at a later date.

Although this will be a review sponsored by, and reported to, the NRC, I hope the report will be useful to others, in our common interest. However, I make no guarantee that the outcome will be supportive of any views now held by anyone — you, me, NRC, UCS, AIF, etc. So far as I am able to assure this, it will be a responsible study, directed at the objectives mentioned in the last paragraph, and we have tried to assemble a group compatible with that purpose. I do hope we will be able to maintain close communication, and that the report will, in the end, fulfill everyone's needs.

Sincerely yours,

H. W. Lewis

HWL:gs

UNIVERSITY OF CALIFORNIA, SANTA BARBARA

BERKELEY • DAVIS • IRVINE • LOS ANGELES • RIVERSIDE • SAN DIEGO • SAN FRANCISCO



SANTA BARBARA • SANTA CRUZ

DEPARTMENT OF PHYSICS
SANTA BARBARA, CALIFORNIA 93106

May 12, 1978

The Honorable Joseph Hendrie
Chairman
U. S. Nuclear Regulatory Commission
Washington, D.C. 20555

Dear Joe:

It will come as no surprise to you that I have been watching with keen interest the negotiations between the Commission and Bob Budnitz, who is a member of our Risk Assessment Review Group. Now that Bob has accepted a position with the Commission, effective after the proposed termination date of our panel, it seems reasonable to assess the implications for our group, and to decide whether any action on our part is necessary. This note is my personal assessment of the situation, suggestion for handling it, and reasons therefor.

I am aware that there are those who believe Bob has irreparably damaged his credibility as an objective commentator on WASH-1400, in accepting a job as Deputy to Saul Levine, who of course played a central role in the preparation of that report. For such people, even without impugning in the slightest Bob's integrity, the appearances of the situation would be sufficient to suggest that he ought to step down as a member of our panel. I do not agree, and there are several reasons, which I will outline below.

In the first place, I believe in treating all personnel problems ad hominem, because that is the best way to make the best use of the short supply of learned and diligent people. Bob was, as you know, a member of our American Physical Society study, and I have worked closely with him from then to the present time. He is simply an honorable man, will do an honest job, and is required for the balance of the panel. It was never intended that each member of the panel stand squarely in the middle on all issues of reactor safety, but that, looked at as a whole, we have a balanced and credible group. I believe that we have, and it includes members whose prior biases (on both sides) are far more open to criticism than is Bob's putative involvement with the Commission. Bob helps to balance the group, and I personally regard that as an important role.

The Honorable Joseph Hendrie
May 12, 1978
Page 2

In the second place, our charter is not just narrowly to critique the Rasmussen report. We are also doing other things which are, in my view, rather more important in terms of what the future holds. It is no secret that we will be saying some negative things about WASH-1400, but even there the point is not to kick a dead horse but to clarify the achievements and limitations of WASH-1400, and through those the question of what can now be said quantitatively about reactor safety. This is an essential ingredient of our other charge, to make a recommendation on whether and how such methodology can be introduced into the regulatory process.

Finally the appearances. There is no question that people with suspicious minds may reach other conclusions than I have about this event, but I do not believe in surrendering to such people. Questions can be raised about my own earlier chairmanship of the American Physical Society study, about Herb Kouts' earlier tour with AEC, and about Frank von Hippel's widely disseminated criticism of the Rasmussen report. I have not explored it, but would not be surprised if several of the members of the panel are recipients of NRC funding for their professional work. There is no end to such things, and I can not subscribe to the view that only those who have no prior experience in a subject are credible.

I regard it as my responsibility to make certain that our report, which I hope will be unanimous, fairly and honestly reflects the facts as we know them, uncorrupted by extraneous considerations. It is my judgment that Bob Budnitz' new job is not a problem in that regard, and I do not plan to take any further action in this matter.

Sincerely yours,



H. W. Lewis

cc: Dr. John Austin ✓

MEETINGS OF THE REVIEW GROUP

1. August 24-25, 1977, in Washington, D.C.
2. October 5-6, 1977 in Washington, D.C.
3. November 10-11, 1977 in Palo Alto, California.
4. December 12-13, 1977 in Washington, D.C.
5. January 5-6, 1978 in Santa Barbara, California.
6. February 6-7, 1978 in Washington, D.C.
7. March 9-10, 1978 in Washington, D.C.
8. April 3-4, 1978 in Berkeley, California.
9. May 23-24, 1978 in Washington, D.C.
10. June 5-8, 1978 in Berkeley, California.
11. August 29-31, 1978 in Washington, D.C.
12. September 7, 1978 in Washington, D.C.

NRC FORM 335 (7-77)		U.S. NUCLEAR REGULATORY COMMISSION BIBLIOGRAPHIC DATA SHEET		1. REPORT NUMBER (Assigned by DDC) NUREG/CR-0400	
4. TITLE AND SUBTITLE (Add Volume No., if appropriate) Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission				2. (Leave blank)	
7. AUTHOR(S) The Risk Assessment Review Group - H. W. Lewis, Chairman, R. J. Budnitz, H. J. C. Kouts, W. B. Loewenstein, W. D. Rowe, F. von Hippel, F. Zachariasen				5. DATE REPORT COMPLETED MONTH YEAR September 1978	
9. PERFORMING ORGANIZATION NAME AND MAILING ADDRESS (Include Zip Code) The Risk Assessment Review Group; an Ad Hoc Advisory Committee Administered for the Commission through the Office of Policy Evaluation				DATE REPORT ISSUED MONTH YEAR September 1978	
12. SPONSORING ORGANIZATION NAME AND MAILING ADDRESS (Include Zip Code) U.S. Nuclear Regulatory Commission				6. (Leave blank)	
10. PROJECT/TASK/WORK UNIT NO. None				8. (Leave blank)	
11. CONTRACT NO. None				12. SPONSORING ORGANIZATION NAME AND MAILING ADDRESS (Include Zip Code) U.S. Nuclear Regulatory Commission	
13. TYPE OF REPORT Final Report			PERIOD COVERED (Inclusive dates)		
15. SUPPLEMENTARY NOTES				14. (Leave blank)	
16. ABSTRACT (200 words or less) <p>Report of the Risk Assessment Review Group of the Nuclear Regulatory Commission. The Group was charged to clarify the achievements and limitations of the Rasmussen report, WASH-1400, to study the state of such risk assessment methodology, and to recommend whether and how such methodology can be used in the licensing and regulation of reactors. The Group found that there were areas of both conservatism and non-conservatism in WASH-1400, so that the uncertainty in the overall probability of core melt is larger than stated in that report. However, the Group could not judge whether the central value given is high or low. The Group recommended that the Commission make much more extensive use of the methodology in setting priorities for the research program, and in more rationally allocating its own resources. In addition, this report contains discussions of a number of disjoint items, such as the Browns Ferry accident.</p>					
17. KEY WORDS AND DOCUMENT ANALYSIS Risk Assessment methodology, Reactor Safety Study (WASH-1400), Fault-tree/event-tree methodology, consequences of reactor accidents			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 \$