SUMMARY OF THE STATE OF THE ART:
Risk Communication to Decision-Makers

Prepared by

Vicki M. Bier, Director
Center for Human Performance in Complex Systems
University of Wisconsin-Madison
Madison, Wisconsin

December 1999

NRC Contract No. 04-97-064
Project Manager, Isabelle Schoenfeld
Office of Nuclear Regulatory Research

Attachment 2

1. Introduction[1]

Risk communication to regulators and other decision-makers has not received nearly as much attention as risk communication to the general public, despite the importance of risk analysis to regulatory decisions in a number of arenas. However, Balch and Sutton (1995) note that "putting the...audience first" is just as important when communicating to regulatory decision-makers as to the general public. The reliance on risk assessment in government is growing. For example, the U.S. Nuclear Regulatory Commission (NRC) is making increased use of risk-informed regulation; the U.S. Environmental Protection Agency (EPA) is also increasingly relying on probabilistic methods of risk assessment, including analyses of uncertainties (Thompson and Graham, 1996). As the information that government decision-makers must absorb becomes more technical, effective risk communication to decision-makers becomes more important, and several organizations (including the EPA) have now begun studying risk communication to decision-makers.

2. Aims and objectives of risk communication to decision-makers

A group of individuals responsible for risk management was convened by the American Industrial Health Council (AIHC) and other organizations (AIHC et al., 1989) to identify their needs with regard to risk assessment information as follows. In addition to th eAIHC, the group membership also included the EPA, the U.S. Food and Drug Administration, the U.S. National Institute of Occupational Safety and Health, several major corporations or trade organizations, leading universities, and other research organizations. The needs identified by the group are as follows:

> "The risk assessment presentation must be comprehensive and understandable."
>
> "The applicability and usefulness of the assessment for public policy decision-making should be clearly stated."
>
> "The presentation must be credible and fully defensible (so it will not result in after-the-fact surprises)."
>
> "The risk assessment report should contain a clear and relatively brief summary that includes balanced treatment of all relevant contentious issues."
>
> "The basis for the choice of critical scientific assumptions should be described along with discussion and resolution of science issues, so far as possible."

---

[1] 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 employee, makes any warranty, express or implied, or assumes any legal liability or responsibility for any third party's use, or the results of such use, of any information, apparatus, product or process disclosed in this report, or represents that its use by such third party would not infringe privately owned rights. The views expressed in this report are not necessarily those of the U.S. Nuclear Regulatory Commission. This report has not been published as a NUREG series publication.

"Conclusions should be drawn so as to be relevant to the specific risk management policy framework."

Note that these features are generally desirable in communicating any type of information, but are likely to be particularly important in communicating about risks, since risk assessments are often difficult to understand, laden with assumptions, and controversial. The summary of a later workshop sponsored by the AIHC et al. (1992) similarly states that risk analyses should be "*relevant..., timely and comprehensible*" (emphasis in original), and should provide a variety of risk measures (e.g., both societal and individual risk estimates) as well as a clear statement of uncertainties.

These guidelines describe a state of affairs that many would agree is desirable, but do not provide much concrete guidance about how to achieve those objectives. More specifically, based on input from a focus group of managers at the U.S. Environmental Protection Agency (EPA), Bloom et al. (1993) summarize the key issues they typically consider when making risk-related decisions, including:

Legal requirements
Possible adverse effects of the particular hazard being regulated
Available options for reducing the risk
Extent of concern about the issue on the part of various groups
Reliability of the information on which the decision will be based

The final report of the Presidential/Congressional Commission on Risk Assessment and Risk Management (1997) provides a similar list of issues to be considered. Bloom et al. also recommend that risk communication briefings should address both quantitative and qualitative descriptions of risks (since qualitative aspects of risk are often critically important to members of the public), and should include discussions of any key uncertainties (e.g., significant gaps in the available information base). The National Research Council (1994) similarly states that: "Risk managers should be given characterizations of risk that are both qualitative and quantitative... When EPA reports estimates of risk to decision-makers..., it should present not only point estimates of risk, but also the sources and magnitudes of uncertainty associated with these estimates."

One key information need identified by Balch and Sutton (1995) is the desire of senior decision-makers "to have more control over the public problems they encounter...They would rather avoid crises than manage them." Thus, Balch and Sutton recommend that risk communication should be "proactive," helping decision-makers identify the most likely public responses to the various decision options under consideration. Just as Rowan (1991) suggests that one purpose of risk communication to the general public can be to raise awareness of a particular hazard or problem, the same purpose can be served by risk communication within government agencies. Useful risk analysis and risk communication efforts need not be limited to the results of formal risk assessments "commissioned" by senior decision-makers, but can also include informal "heads-up" notification of possible problems.

Balch and Sutton note that managers "need to show success," so that even small efforts can be worthwhile if they result in "early victories" (e.g., demonstrations of public support for a particular decision or program, analyses showing that a particular decision has reduced risk or saved money). They also point out that useful information is available not only at the time of an initial decision (e.g., during program planning) and after program completion, but also while a program is underway. Thus, they recommend that risk communicators do "formative research" to provide input to improve a program during the program, not just at the start; and do "process research" on how the program is working during the program, not just at the end. Manager-friendly research includes diagnostic feedback and mechanisms for early warnings and quick adjustments in the program followed by feedback on how the adjustments are working.

Finally, the AIHC et al. (1992) note that "two-way communication between assessors and users" of risk assessment information is just as important in dealing with government managers as with the general public. Therefore, Balch and Sutton (1995) recommend conferring with senior managers before a risk analysis is actually performed, to determine how the results will be used and which questions most need to be answered. This is echoed by the AIHC et al. (1992), who recommend that "potential users" of risk analyses (including both government decision-makers and other stakeholders) and their concerns should be identified early, so that the risk analysis can be designed to address those concerns. Similarly, just as in communication with the public, Balch and Sutton (1995) emphasize that risk analysts should "*Promise only what you can deliver and deliver what you promise*" (emphasis in original), in order to avoid any loss of credibility for risk assessment as a discipline.

3. Treatment of uncertainty, variability, and correlation

While information on uncertainties is an essential input to regulatory decision-making, it is often especially difficult to communicate. Moreover, the theory of decision-making under uncertainty is a highly technical subject, and often not transparent to decision-makers not trained in this methodology. For example, Thompson and Graham (1998) state, "the lack of experience that risk managers and their staffs have with probabilistic risk results means that understanding such results requires more effort than they would expend on point-estimate results. Decision-making is also more difficult because risk managers cannot simply compare the results to regulatory 'bright lines'" (i.e., definitive "decision criteria" for determining whether a given probability distribution represents an acceptable level of risk). Therefore, this section presents a brief discussion of several ways in which uncertainties can affect decision-making, and some suggestions on how to effectively address these issues in risk communication.

The problem of decision-making under uncertainty is addressed by standard decision analysis theory (see for example Raiffa, 1968; Clemen, 1991), which is theoretically rigorous and is by now well accepted. Decision-making based on explicit statements of uncertainty (i.e., using probability distributions to characterize the uncertainty about important inputs to the decision) has numerous advantages over decision-making based purely on point estimates or "deterministic" analyses (Thompson and Graham, 1998). Morgan and Henrion (1990) describe these benefits as "the value of knowing how little you know." Morgan and Henrion point out that in the vast majority of cases (i.e., for most utility functions or loss functions), taking uncertainty into account will lead to a better decision than ignoring the uncertainty and relying on a single point estimate (such as the mean value of the uncertain quantity).

The methodology of decision theory provides guidance not only on how to select the best from among several uncertain options, but also on "the value of information"—i.e., whether it is preferable to make a final decision, or gather more information and defer the decision until that information becomes available. In the words of Thompson and Graham (1998), "statistical decision theory provides...insight about the sources of uncertainty that the decision-maker may wish to resolve or reduce prior to making a decision." Thus, one important use of risk analyses (especially analyses that include an explicit statement of uncertainties) is to assist in determining whether additional research is needed.

To help clarify the concept of value of information, Brown and Ulvila (1987) distinguish between "outcome uncertainty" ("what might actually happen and with what probability?") and "assessment uncertainty" (how much the results of the analysis might change with additional information). The authors note that outcome uncertainty is the concern of decision-makers who must reach a final decision immediately, while both outcome uncertainty and assessment uncertainty are of interest to decision-makers who have the option of collecting more information before deciding. In discussing assessment uncertainty, they also distinguish between "unlimited [or perfect] information" (i.e., if it were possible to ascertain the "true risk") versus amounts of "new information" that might result from a realistic research effort (i.e., information or analysis that might actually become available, say, by waiting a few years). They point out that new information can not only reduce the outcome uncertainty but also shift the central estimate of that outcome. Finally, they suggest several possible graphical representations of both outcome uncertainty and assessment uncertainty. However, no empirical results about the effectiveness of these various presentation formats are provided, so it is difficult to determine whether the formats suggested by the authors are an improvement over other methods of displaying uncertainties. Moreover, some risk analyses compute and present only the total outcome uncertainty, not the assessment uncertainty. In these cases, decision theory cannot offer a rigorous answer to the question of whether additional research is justified, but only an upper bound to the value of information (obtained under the unrealistic bounding assumption that research could eliminate all of the uncertainty).

The distinction between "assessment uncertainty" and "outcome uncertainty" established by Brown and Ulvila (1987) is closely related to the distinction between "state of knowledge" uncertainty and "population variability" originated by Kaplan (1983). In this choice of terminology, "state of knowledge" uncertainty is generally taken to include those uncertainties that are (at least in principle) reducible by further research, and thus corresponds to the use of the term "assessment uncertainty" by Brown and Ulvila. Those aspects of the total outcome uncertainty that can *not* be readily reduced through further research are then considered to reflect inherent variability in the population at hand. For example, there may be variability among different nuclear power plants, and also state of knowledge uncertainty about the risk at any particular power plant (or about the average risk in a population of power plants).

State of knowledge uncertainty and population variability have different implications for decision-making. This is well explained by the National Research Council (1994): "[U]ncertainty forces decision makers to judge how *probable* it is that risks will be overestimated or underestimated for every member of the exposed population, whereas variability forces them to cope with the *certainty* that different individuals will be subjected to risks both above and below any reference point one chooses" (emphasis in original). Morgan and Henrion (1990) and Cullen and Frey (1999) both provide useful discussions of the distinctions between state of knowledge uncertainty and population variability (sometimes referred to simply as "uncertainty" and "variability"). For example, Morgan and Henrion note:

> A common mistake is failure to distinguish between variability due to sampling from a frequency distribution and empirical uncertainty that arises from incomplete scientific or technical knowledge...the scientific uncertainty can be reduced only by further research... But this sort of uncertainty analysis is impossible unless the two kinds of uncertainty are clearly distinguished.

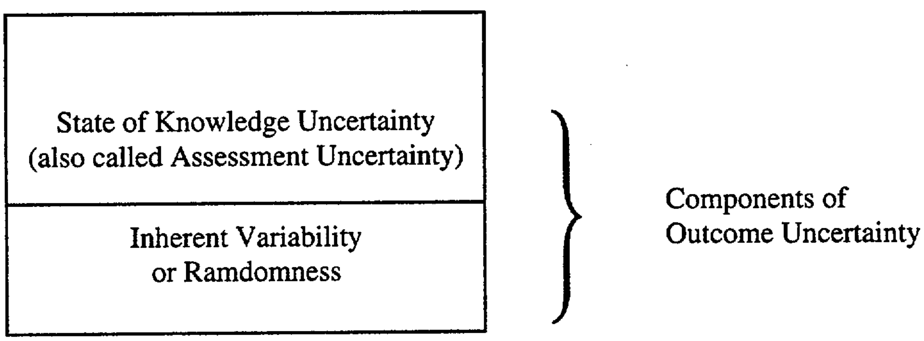

A clear understanding of the different implications of uncertainty and variability for decision-making is useful in risk communication, since it allows the risk communicator to highlight these implications for the decision-maker. To illustrate these implications, consider the case where the various plants or facilities in a population differ a great deal from each other (i.e., high population variability), but the risk at each one is accurately known (i.e., low state of knowledge uncertainty). In this case, it makes sense to target regulatory efforts at the facilities with the highest estimated risks. Thus, risk analysis results can be useful in prioritizing regulatory efforts, especially when the state of knowledge uncertainties in those results are relatively small. By contrast, when the plants or facilities in a population have relatively similar risk estimates (i.e., low variability), but those estimates are all quite uncertain, then the facility with the highest risk *estimate* may not actually have the highest *risk*. In this case, it may be better to spread one's regulatory efforts more or less evenly across the entire population, to protect against the possibility that risks have been mis-estimated and a high-risk facility has been neglected.

Large state of knowledge uncertainty also indicates that the average risk level of the entire population is not known very well, so further research may be desirable (e.g., to determine whether the risks are large enough to warrant increased ~~in~~ regulatory effort), as noted by Morgan and Henrion (1990). This suggests that risk communication messages should clearly distinguish state of knowledge uncertainty from population variability, and that risk analyses (particularly "generic" analyses addressing the risks from an entire population of similar facilities) should provide the information needed to make that distinction). In the absence of such information, it will often be unclear whether the results of a risk analysis indicate a need for carefully targeted regulatory efforts to address the highest-risk facilities in the population, or further research to clarify the actual risk levels.

Similar logic applies to the risk contributors within a given system. Targeting regulatory efforts at the largest risk contributors (as estimated based on PRA) will be useful when the state of knowledge uncertainties in the PRA are relatively small, and when the various contributors differ a great deal in their importance to risk. By contrast, when uncertainties are large relative to the differences between risk contributors, then the use of PRA results to prioritize regulatory effort (e.g., inspection effort at a particular plant) will be relatively unreliable, since the PRA may fail to identify the most important risk contributors. Risk communication messages that clearly summarize the uncertainty about each risk contributor in a PRA, and highlight the implications of that uncertainty, can help decision-makers such as inspectors make effective use of the PRA results in improving the targeting of their efforts. By contrast, ranking of risk contributors with no indication of uncertainties may occasionally suggest that the most important contributors to risk are known, when in fact there is still substantial uncertainty about that question.

Bier (1989) discusses additional distinctions between state of knowledge uncertainty and population variability. For example, if the total outcome uncertainty represented by a given generic PRA consists primarily of state of knowledge uncertainty, then the plants or facilities in that generic population will tend to have highly correlated risk levels. In this case, the range of possible societal consequences will be much broader than if the total uncertainty represented primarily population variability. Thus, the chance of having two or more accidents over a period of several years will tend to be greater with large state of knowledge uncertainty than with large population variability (since with large variability, not all facilities in the population will have high risks). Similarly, back-fit costs following an accident will also tend to be larger in cases with large state of knowledge uncertainties. To see this, consider the different impacts of an accident caused by a generic industry-wide design flaw versus an accident caused by a site-specific feature. An accident that is caused by a generic design flaw is likely to result in back-fits at all similar plants or facilities in the population, at a large total cost. By contrast, if the cause of the accident is truly site-specific and unlikely to occur at other plants, then safety improvements will be needed at only one (or perhaps a few) facilities, at a much lower total societal cost. Therefore, risks that are characterized by large state of knowledge uncertainty and small population variability may merit greater risk reduction expenditures than risks characterized by large population variability and small state of knowledge uncertainty, if it is not possible to cost-effectively reduce the state of knowledge uncertainty through further research. One way to highlight this for decision-makers is by presenting not only the PRA results for a single plant or facility (e.g., a hypothetical "average" plant), but also the range of overall consequences (e.g., total nationwide consequences over a period of several decades) associated with the entire population of facilities.

Bier (1989) also notes that uncertainty creates additional complexities when comparing two risks (e.g., the risks of two alternative methods for performing the same function), especially if they are correlated. When the technologies are so different that their risks are unlikely to be correlated (e.g., when comparing coal to nuclear power), then the difference in risk from using one technology rather than another can be determined by simply subtracting the two risk estimates under the assumption that they are independent. However, in some cases the alternatives being compared might be quite similar (e.g., a particular plant with and without an extra safety system, a particular system with and without the addition of a redundant train of equipment). In such cases, even if the overall risk levels of both alternatives are quite uncertain (e.g., because of uncertainty about dose-response curves or some other feature unrelated to the improvement being considered), we may know with high confidence that adding an extra system or train will actually decrease the risk. Subtracting the estimated risks of the two alternatives under the assumption of independence can overestimate the uncertainty about the extent of risk reduction provided by the safer option (and hence lead decision-makers to understate its effectiveness). To avoid such misinterpretations, it may be helpful to present not only risk estimates for the two alternatives individually, but also a probability distribution for the difference between the two risks (i.e., the extent of risk reduction that would be achieved by selecting the safer alternative). This can help decision-makers distinguish between uncertainty about the effectiveness of a given risk reduction measure, and uncertainty about the magnitude of the risk that it is designed to address.

4. Format of risk communication

Bloom et al. (1993) provide several guidelines on how to structure briefings, based on input from senior EPA managers. In particular, briefings should ideally begin with "an overview of why the action (e.g., regulation) under consideration is important," "who... cares about this issue," and "what the major stakeholders [including other government offices] are saying." The technical portion of the risk communication should then discuss "the severity of the adverse health effect," "the level of confidence...in the data," "where the data gaps are," and the nature and size of the affected population. With regard to risk management, managers indicated that briefings should provide information on the costs and benefits of the various "potential risk management options" being considered, as well as "the consequences of doing nothing" and any "studies in progress which could yield new and important information."

Similarly, focusing specifically on risk analysis of carcinogens, the AIHC et al. (1989) identify several desirable attributes for risk assessments:

> "The scope and objectives of the report are explicitly stated."
>
> "The report's content is laid out impartially, with a balanced treatment of the evidence bearing on the conclusions."
>
> "The risk assessment presentation includes a description of any review process that was employed, acknowledging specific review commentary."
>
> "The key findings of the report are highlighted in a concise executive summary."
>
> "The report explains clearly how and why its findings differ from other risk assessment reports on the same topic."
>
> "The report explicitly and fairly conveys scientific uncertainty, including a discussion of research that might clarify the degree of uncertainty."

The AIHC study group report then goes on to outline specific attributes considered desirable for the various stages of a risk analysis, including hazard assessment, dose-response analysis, exposure analysis, and overall "risk characterization" or summary of results. Many of these attributes are not specific to cancer risk assessment, and could be usefully applied to other types of risk analyses. For example, in the area of hazard assessment, the AIHC report: recommends that all relevant information is presented, it highlights critical aspects of data quality, and identifies research that would permit a more confident statement about human hazards.

Bloom et al. (1993) rated several specific information formats (including verbal, graphical, and tabular presentations). In general, presentation formats of intermediate complexity (e.g., simple tables or bar charts) were viewed as being most effective by the EPA decision-makers that participated in the study. By contrast, overly complex formats "with more detail than needed to make decisions" (e.g., presenting too many percentiles on a single chart) and overly simplistic presentations (e.g., a single probability distribution, with no comparison points and no information on risk contributors) were both found to be problematic.

Several authors have noted that excessive emphasis on the numerical results of a risk analysis and the statistical methods that were used in the analysis is generally not very helpful. In particular, many decision-makers are not trained in statistical techniques, and are likely to perceive such issues as "not relevant" to the problem at hand (Balch and Sutton, 1995). Thus, for example, Balch and Sutton (1995) recommend avoiding even such seemingly benign terms as "research," "evaluation," or "data," focusing instead on specific "problems or opportunities...for the agency—and solutions." Rather than an emphasis on statistical methods, decision-makers are likely to want more information about the qualitative assumptions underlying the analysis, and the reasons for the results. Therefore, formats that give insight into the causes of risk (e.g., diagrams illustrating particular event sequences that contribute significantly to risk) are likely to be more useful than simple numerical summaries or probability distributions.

More specifically, in communicating risk information to professionals in other disciplines, attention should be paid to ensure that the topics of greatest relevance to those disciplines are emphasized. For example, engineers are likely to be more interested in understanding the specific scenarios that contribute to risk analysis results and assessing whether they are reasonable than in the numerical risk results and/or the statistical methods used in the risk analysis. Engineers are also likely to have greater confidence in a risk analysis if they understand how the results of deterministic engineering analyses are reflected in that analysis (e.g., how detailed thermal hydraulic calculations were used in determining the success criteria for a particular safety system). Similarly, psychologists are likely to be interested in understanding how human performance is reflected in a risk analysis (e.g., which human actions and performance shaping factors were considered, how the probabilities of those actions were quantified, and what allowance was made for the uncertainty and unpredictability associated with human behavior). Finally, those responsible for facility operations and maintenance are likely to be concerned about how their input (e.g., descriptions of typical operating practices, reports of observed incidents, and data on preventive and corrective maintenance actions) is reflected in the risk analysis.

With regard to the needs of the NRC in particular, Coe (1998) has suggested that the results of a probabilistic risk analysis (PRA) should be presented in a manner that:

> "Places risk-important SSCs [structures, systems, and components] into an accident sequence context that reveals their risk relationship to other SSCs."
>
> "Helps to reveal accident sequences...in terms of their penetration through specific layers of 'defense-in-depth'."
>
> "Shows how and where the PRA uses traditional engineering analysis (e.g., success criteria) and its influence on the results."

Bley (1997) has recommended a similar approach. In particular, he suggests the use of event sequence diagrams to graphically describe the steps leading to a particular type of accident sequence. With an appropriate computer implementation, this approach would let users (such as risk managers) click on elements of the overall event sequence diagram to obtain additional relevant information, such as system schematics and descriptions of operator actions. Such an approach would greatly reduce the emphasis on numerical results and statistics in presenting risk analysis results, and instead focus attention on the physical phenomena responsible for the risk. Therefore, it is likely to improve both the understanding of the results and also the reliance on risk analysis by key decision-makers.

However, it is important to note that the best presentation format may vary depending on the disciplinary background of the intended audience (e.g., engineers, psychologists, or toxicologists). Therefore, pilot testing of different presentation formats (see for example Lundgren and McMakin, 1998) can be just as important for an audience of decision-makers as for the general public. In the absence of formal pilot testing (e.g., in more time-critical or budget-limited situations), Balch and Sutton (1995) recommend the use of "informal 'communication checks' from someone on the program team" before important presentations, to help minimize the possibility of miscommunication.

4. Conclusion

Little research has been done on effective methods of communicating risk analysis results to decision-makers. The AIHC et al. (1992) workshop participants noted the need for research on several topics, including:

1) Effectiveness of different risk communication strategies
2) Methods for enhancing the "risk literacy" of decision-makers (e.g., their familiarity with and understanding of probabilistic risk methods)
3) Methods for integrating "technical information about risk with information on other social values" (e.g., through stakeholder participation processes)

With the advent of risk-informed regulation, decision-makers are increasingly being asked to take highly technical risk analysis results into account in their decisions. Therefore, it is important to pilot-test risk communication messages and approaches whenever possible, and to share the results of such evaluations as widely as possible to disseminate knowledge about how to most effectively communicate risk information to support regulatory decision-making. Even informal evaluations or anecdotal information on effective and ineffective approaches to risk communication can be valuable, since the understanding of the field is at a relatively rudimentary stage.

References

American Industrial Health Council, U.S. Environmental Protection Agency, Department of Health and Human Services, and Society for Risk Analysis, Presentation of Risk

Assessments of Carcinogens: Report of an Ad Hoc Study Group on Risk Assessment Presentation, Washington, D.C.: American Industrial Health Council, 1989.

American Industrial Health Council, Center for Risk Management (Resources for the Future), and U.S. Environmental Protection Agency, Improving Risk Characterization: Summary of a Workshop held in Washington, DC on September 26 & 27 1991, Washington, D.C.: American Industrial Health Council, 1992.

G. I. Balch and S. M. Sutton, "Putting the First Audience First: Conducting Useful Evaluation for a Risk-Related Government Agency," *Risk Analysis*, Vol. 15, No. 2, 1995, pp. 163-168.

V. M. Bier, "On the Treatment of Dependence in Making Decisions About Risk," Transactions of the 10th International Conference on Structural Mechanics in Reactor Technology, Anaheim, California, 1989.

D. C. Bley, "Risk Communication," presented to Japan Power Engineering and Inspection Corporation (JAPEIC), Washington, D.C., 1997.

D. L. Bloom, D. M. Byrne, and J. M. Andresen, Communicating Risk to Senior EPA Policy Makers: A Focus Group Study. Washington, D.C.: U.S. Environmental Protection Agency, 1993.

R. Brown and J. W. Ulvila, "Communicating Uncertainty for Regulatory Decisions," in V. T. Covello, L. B. Lave, A. Moghissi, and V. R. R. Uppuluri (editors), Uncertainty in Risk Assessment, Risk Management, and Decision Making, New York: Plenum Press, 1987.

R. T. Clemen, Making Hard Decisions: An Introduction to Decision Analysis, Boston: PWS-Kent, 1991.

D. Coe, "Risk Communication for Risk Informed Regulation," Washington, D.C.: U.S. Nuclear Regulatory Commission, 1998.

A. C. Cullen and H. C. Frey, Probabilistic Techniques in Exposure Assessment: A Handbook for Dealing with Variability and Uncertainty in Models and Inputs, New York: Plenum Press, 1999.

S. Kaplan, "On a 'Two-Stage' Bayesian Procedure for Determining Failure Rates from Experiential Data," *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS-102, 1983, pp. 195-202.

R. Lundgren and A. McMakin, Risk Communication: A Handbook for Communicating Environmental, Safety, and Health Risks, Columbus (Ohio): Battelle Press, 1998.

M. G. Morgan and M. Henrion, Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis, Cambridge: Cambridge University Press, 1990.

National Research Council, Science and Judgment in Risk Assessment, Washington, D.C.: National Academy Press, 1994.

The Presidential/Congressional Commission on Risk Assessment and Risk Management, Final Report, Volume 1: Framework for Environmental Health Risk Management, and Volume 2: Risk Assessment and Risk Management in Regulatory Decision-Making, http://www.riskworld.com/Nreports/nr5me001.htm. 1997.

H. Raiffa, Decision Analysis: Introductory Lectures on Choices under Uncertainty, Reading (Massachusetts): Addison-Wesley, 1968.

K. E. Rowan, "Goals, Obstacles, and Strategies in Risk Communication: A Problem-Solving Approach to Improving Communication about Risks," *Journal of Applied Communication Research*, Vol. 19, 1991, pp. 300-329.

K. M. Thompson and J. D. Graham, "Going Beyond the Single Number: Using Probabilistic Risk Assessment to Improve Risk Management," *Human and Ecological Risk Assessment*, Vol. 2, No. 4, 1996, pp. 1008-1034.