September 30, 1997 SECY-97-221

FOR: The Commissioners

FROM: L. Joseph Callan /s/
Executive Director for Operations

SUBJECT: ACCEPTANCE GUIDELINES AND CONSENSUS STANDARDS FOR USE IN RISK-INFORMED REGULATION

PURPOSE:

To respond to Commission staff requirements memoranda related to risk-informed regulation in the following areas:

- Framework and acceptance guidelines for plant-specific risk-informed decisionmaking; and
- Standards for probabilistic risk assessments (PRAs) to be used in risk-informed decisionmaking.

BACKGROUND:

The Commission has provided guidance to the staff on its applications of risk-informed regulation in several staff requirements memoranda (SRMs). Specifically, guidance was provided in the SRM dated January 22, 1997, for SECY-96-218, the SRM dated April 15, 1997, for Direction Setting Issue (DSI) 12, and the SRM dated March 7, 1997, for DSI 13.

Contacts:
Mark A. Caruso, RES
415-6436
Gareth Parry, NRR
415-1464

DISCUSSION:

The staff's responses to this Commission guidance are discussed below. Staff responses to guidance contained in these SRMs related to performance-based regulation, and other staff interactions on codes and standards, are contained in separate Commission papers.

The staff's basic framework for risk-informed regulation is described in draft regulatory guide DG-1061[1]. This guide provides a set of basic principles for the staff's risk-informed regulatory practices, as well as acceptance guidelines. As discussed below, the staff will implement the guidance in SECY-96-218 and DSI 12 SRMs as part of finalizing the DG-1061 framework and acceptance guidelines.

Evaluating safety impacts in an integrated manner

In the SRM on SECY-96-218, the Commission directed that staff risk-informed regulatory guidance should evaluate all safety impacts of proposed changes in an integrated manner. The staff has done so, by defining in DG-1061 a framework for analyzing and evaluating proposed changes to the current licensing basis (CLB) of licensed nuclear power plants. This framework is based on the Commission's policy to use PRA technology in a manner that complements traditional engineering approaches and the defense-in-depth philosophy. The guidance in DG-1061, as well as the application-specific regulatory guides (RGs) and associated standard review plans (SRPs), indicates that the staff expects risk-informed decisions to be made in an integrated manner, considering traditional engineering and risk information, and that decisions may be based on qualitative factors as well as quantitative analyses and information. The staff's set of key principles defined in DG-1061, which blend together defense in depth, safety margin, and risk assessment concepts, were established to ensure this integration.

In the SRM on SECY-96-218, the Commission also directed that staff risk-informed regulatory guidance should use risk insights to identify areas where requirements should be increased or improvements could/should be implemented. The staff's draft DG-1061 and SRP Chapter 19 discuss the issue of using risk assessment both to improve safety and reduce burdens. DG-1061 describes a set of staff expectations regarding use of PRA in risk-informed regulation, one part of which indicates that the DG-1061 approach should be used "to identify areas where requirements should be increased as well as where they could be reduced." The staff's application-specific guidance[2] make this general guidance more explicit; for example, the in-service testing (IST) guidance states that systems, structures, and components which presently are outside of the IST program should be included in the program if they are found to have high safety significance, based on the licensee's risk-informed engineering evaluation. Another feature of the IST guidance that can lead to

---

[1] "An Approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant Specific Changes to the Current Licensing Basis," NRC Draft Regulatory Guide, DG-1061, June 1997.

[2] In June 1997 the NRC issued draft Regulatory Guides and Standard Review Plan sections for comment by the public which address Inservice Testing of Pumps and Valves (DG-1062 & draft SRP Chapter 3.9.7), Graded Quality Assurance (DG-1064) and Technical Specifications (DG-1065 & SRP Chapter 16.1).

improved safety is that if a licensee proposes to change only test intervals, and not methods, then that licensee must give consideration to components that are potential candidates for more frequent testing as well as to candidates for less frequent testing.

The staff is now considering the need for additional guidance that encourages licensees to search for regulatory burden reductions which reduce risk or are risk neutral prior to proposing CLB changes involving risk increases. If developed, this guidance would be documented in the final version of DG-1061.

Methodology for assessing changes

In the SRM on SECY-96-218, the Commission also directed that the staff develop a methodology for assessing changes in risk that uses statistical concepts and gives consideration to uncertainties. The staff has proposed an approach in DG-1061 and draft Standard Review Plan (SRP) Chapter 19 for assessment of changes in risk resulting from proposed changes to a plant's CLB. In this approach, the results of the PRA are compared with acceptance guidelines on core damage frequency (CDF) and large early release frequency (LERF) which have been based on the Commission's Safety Goals and subsidiary objectives. The guidelines include assessments of the absolute values of CDF and LERF and the amount of change in these measures associated with the proposed CLB change. Mean values are used for these comparisons, consistent with the guidance in the Safety Goal Policy Statement. Guidelines governing changes in CDF and LERF are consistent with the current guidelines the staff uses when conducting regulatory analysis in support of a backfit. DG-1061 also includes general guidance on how to analyze and consider uncertainties in the evaluation. The guidance provides considerable flexibility, permitting qualitative and quantitative uncertainty analysis, reflecting present limited capabilities to quantitatively address important sources of uncertainty (e.g., modeling and completeness uncertainty).

Since issuing DG-1061 for comment, the staff has continued to discuss the topic of uncertainty analysis with the Advisory Committee on Reactor Safeguards, Subcommittee on PRA, received some initial public comment on the treatment of uncertainty, and studied several approaches for expanding the guidance on treatment of uncertainty. The attachment to this paper discusses these approaches, including the pros and cons of each, and suggests that a characterization of uncertainties using statistical methods may, in some cases, need to be supplemental with other information and, in other cases, may not be necessary to support risk-informed decisionmaking. For example, current PRAs often do not include quantitative, probabilistic analyses of certain accidents (e.g., those initiated during shutdown operations) which can be important contributors to a facility's total core damage frequency and risk and the associated uncertainties. Such "completeness uncertainty" may be more appropriately addressed qualitatively.

In addition, for very small risk increases, a rigorous assessment of uncertainty or strict limits on baseline CDF and LERF may not be necessary because of the insignificance of the proposed risk increase. In fact, it would be inconsistent with the working definition of *risk-informed* (i.e., to focus licensee and NRC staff attention on those systems and activities of most risk significance) to treat very small changes as important. The staff is continuing to assess the various approaches discussed in the attachment while developing final guidance on the treatment of uncertainties in DG-1061. In developing this guidance, the staff will use the perspectives gained from this assessment, in conjunction with public comment received on the draft guide, and further discussions with the Advisory Committee on Reactor Safeguards. Any major policy issues arising in the development of this final guidance (due to the Commission on December 31, 1997) will be brought promptly to the Commission for decision.

Procedure for monitoring cumulative changes in risk

In its January 22, 1997, SRM, the Commission indicated that the staff should establish procedures to monitor the cumulative changes in risk for a given nuclear facility as the result of license amendments that are conducive to quantitative risk assessments.

Guidance in Section 3.3 of DG-1061 states that the staff expects the licensees to track and consider the cumulative impact of all plant changes on risk, including those not submitted for NRC review and approval. Draft SRP Chapter 19, which covers risk-informed changes to the current licensing basis, instructs reviewers to verify that: each application is carried out with reference to a PRA model that already reflects previous applications; that cumulative changes from license amendments are being monitored; and the accumulation of applications has not created dominant risk contributors. Furthermore, the “increased management review” cited in DG-1061 for applications which result in risk metrics that are close to exceeding acceptance guidelines includes an evaluation of the cumulative impact.

The guidance contained in DG-1061 and draft SRP Chapter 19 is intended to be sufficient to enable the staff to track and monitor cumulative changes in risk associated with risk-informed license amendment requests. This guidance will be retained in the final issuance of these documents. The staff notes, however, that this guidance may not be sufficient to capture changes made in accordance with 10 CFR 50.59 that are not submitted for NRC review and approval. Accordingly, the staff is giving consideration to amending its guidance, as necessary, to address relevant facility changes that are not submitted for review and approval.

Clarification of terms

In the January 22, 1997, SRM, the Commission expressed concern that the fourth element of the staff’s proposed approach for risk-informed, performance-based regulation was vague as written and that the staff should consider more definitive language. The fourth element addresses the need to establish performance criteria that capture precursors to failures that put the public at risk and not the failures themselves, which the staff termed *intolerable outcomes*.

Since the SRM was issued, the staff has removed the phrase *intolerable outcomes* from its draft guidance and replaced it with the following language:

> “The [licensee's] monitoring plan should be structured such that performance degradation is detected and corrected before plant safety can be compromised.”

For example, in the risk-informed IST application, performance characteristics will be as required by the ASME Code or Code Case, with the intended goal that degradation is not significant for components that are placed on an extended test interval and that the failure rate assumptions for these components remain valid.

In the January 22, 1997, SRM, the Commission also approved the staff’s recommended policy of allowing small increases in risk under certain conditions, for proposed changes to a plant’s licensing basis, but requested that the staff define the terms *small* and *under certain conditions* more clearly. As discussed above, the acceptance guidelines described in DG-1061 include assessments of the absolute values of CDF and LERF and the amount of change in these measures associated with the proposed CLB change. The guide specifies the maximum allowable increases in CDF and LERF ($1x10^{-5}$ per reactor year and $1x10^{-6}$ per reactor year, respectively), which are considered to be *small* compared with Commission’s quantitative health objectives and subsidiary objectives.

The *certain conditions* which control the amount of allowed increases are specified in DG-1061 and include whether the plant’s mean CDF or LERF are above the acceptance guidelines (in which case only risk neutral or risk-reduction changes would be permitted), or below the guidelines, (in which case increases would be limited to an amount less than or equal to the available margin). Proposed changes for plants with mean CDF or LERF within a factor of ten of the guidelines would receive additional NRC management attention, considering such issues as the robustness of the PRA, cumulative impact of previous changes in risk, impact of change on operational complexity, and the current level of plant performance (as measured by such factors as inspection findings and performance indicators). It should be noted that, as discussed above, the staff is considering a change in this area that would allow very small risk increases, regardless of baseline CDF or LERF values.

Objective standards

In the SRM on DSI 12, the Commission indicated that “the staff should develop objective standard(s) for the application of risk-informed, performance-based and risk-informed less prescriptive approaches to regulations on an expedited basis. Such standard(s) could be in the form of individual plant safety goals and subsidiary objective performance criteria as discussed in the issue paper.”

The staff has defined in DG-1061 acceptance guidelines which serve as “objective standards.” As discussed above, these guidelines include assessments of the absolute values of CDF and LERF and the amount of change in these measures associated with the proposed CLB change. The numerical values used in these are based on the Commission’s Safety Goal Policy quantitative health objectives and subsidiary objectives. The objective standards would also include assuring that the defense-in-depth philosophy and engineering safety margins are appropriately retained.

Experience from Implementing the Maintenance Rule

In the SRM on DSI 12, the Commission also indicated that "the staff should also describe how any relevant knowledge developed in the implementation of the maintenance rule will be utilized in the development of risk-informed, performance-based regulation."

In SECY-97-055, dated March 4, 1997, the staff informed the Commission about lessons learned from the implementation and initial baseline inspections of the maintenance rule and insights gained from the NRC staff's experience with the maintenance rule that will be considered when developing other risk-informed, performance-based rules. The following key insights have implications for future development of risk-informed, performance-based regulations:

- The most significant observation relative to the implementation of the maintenance rule is that, in general, licensees have developed the knowledge and ability to use a risk-informed approach to maintenance planning and to equipment reliability and unavailability. In developing their programs, licensees have frequently exceeded the minimum guidance in the NRC Regulatory Guide 1.160 and NUMARC 93-01.

- The maintenance rule gives licensees considerable flexibility for implementing requirements (which is a key attribute of a performance-based rule). This has lead to a wide variety of site-specific implementations of the maintenance rule which has made the inspection process resource intensive.

- Inspection and enforcement of the maintenance rule has been challenging because the rule does not specify standards for minimal acceptance and includes a provision which is a suggestion, not a requirement. Preparation and training of the staff, over and above that associated with more prescriptive rules, has been necessary to ensure consistent inspection and enforcement.

As maintenance rule baseline inspections continue, the staff will continue to assess the results for insights that are pertinent to the future development of risk-informed, performance-based regulations.

Consensus Standards for PRA

In June, 1997, NRC staff met with representatives of the American Society of Mechanical Engineers (ASME) to discuss cooperation with both industry and professional societies to develop new codes and standards, as directed in the SRM on DSI 13, dated March 7, 1997. The development of PRA standards was one subject of this meeting; ASME indicated their interest and is assembling an ad hoc committee that will have the responsibility to develop such standards. This committee will be comprised of ASME personnel, NRC staff, national laboratory, academic, and other industry personnel.

A charter for this committee is now being drafted, and will describe both programmatic and technical information. Programmatic information will include goals and objectives of the committee, committee membership and associated responsibilities, schedules and milestones, and peer review. Technical information will include anticipated scope of the standards (e.g., Level 1, 2 and 3 PRA, including core damage accidents initiated by internal and external events during full power operation) and the level of PRA modeling and analysis appropriate for different PRA uses.

The committee will use current material as starting points for the PRA standards. Such material could include draft NUREG-1602, and owner's groups comparison and certification documents. Draft NUREG-1602, which was published for public comment in association with the staff's draft regulatory guides and SRPs, describes the tasks associated with a full-scope PRA and the main attributes of each of these tasks in performing a state-of-the-art PRA, as well as associated documentation, and peer review needs. The Commission will be informed of progress on this development work in the quarterly updates of the PRA Implementation Plan.

CONCLUSIONS:

A number of items identified by the Commission in the SRMs on SECY-96-218, DSI 12, and DSI 13 have been addressed by the staff. This has been accomplished primarily in the context of the draft regulatory guidance for using PRA in plant-specific, risk-informed decisionmaking that was published for comment in June 1997.

Several issues identified in the SRMs require continued staff work:

- Several approaches for treating uncertainty in risk-informed regulation have been identified, but require review following the receipt of public comment on the staff's draft guidance and further interactions with ACRS. Any major policy issues arising in the development of this final guidance (due to the Commission on December 31, 1997) will be brought promptly to the Commission for decision.

- The staff is working with ASME to develop PRA standards. A charter, including scope and schedule, for this work is now being developed.

COORDINATION:

The Office of the General Counsel has reviewed this paper and has no legal objections to its issuance.

L. Joseph Callan
Executive Director
for Operations

Attachment:
As stated

Attachment 1

Uncertainty in PRA Results in the Context of Risk-Informed Decisionmaking

The draft Regulatory Guide DG-1061 has proposed an approach to using the results of probabilistic risk assessments (PRAs) in making risk-informed decisions associated with changes to the current licensing basis (CLB). In the version of DG-1061 released for public comment, a set of guidelines was proposed that could be used to demonstrate that the fourth fundamental principle of risk-informed regulation has been satisfied. In applying these guidelines, the licensee must demonstrate that estimates of the risks from the plant, characterized in terms of the risk metrics core damage frequency (CDF) and large early release frequency (LERF), and the impact of the proposed change in the CLB are such that the absolute values do not approach or exceed the acceptance guidelines, and that risk increases, if any, are small, with small being calibrated by numerical values specified in DG-1061. It was recognized, however, that there are significant uncertainties in the numerical predictions of PRA models which must be taken into account when making decisions based, in part, on those results. Methods have been developed for, and applied to, characterizing uncertainty in PRA results. The performance of an uncertainty analysis is an essential element in decisionmaking that provides a decisionmaker with assurance that the best decision has been made, consistent with his state of knowledge. This attachment discusses the nature of uncertainty in PRA models, how it can be characterized, methodological limitations and options for overcoming those limitations, and how the results of an uncertainty analysis may be used in decisionmaking.

The nature of PRA models, and sources and characterization of uncertainty

A PRA model is fundamentally a logic structure that identifies sets of events (initiating event, function, system, or component failures or unavailabilities, human errors, etc,) that have to occur to result in core damage and ultimately, release of radioactivity to the environment. PRA analysts treat many of these events as random processes, and adopt probabilistic models to describe their occurrences, e.g., the Poisson model for initiating events. To develop the logic structure of the model, the analysts make use of a variety of tools that include both deterministic and probabilistic models of phenomena.

It is useful, as in DG-1061, to categorize uncertainties into three types; parameter, model, and completeness uncertainties, since the approaches used to characterize their impact on the result are different.

Each of the models that is used, either in developing the logic structure, or to represent the basic events of that structure has one or more parameters. Typically, the form of each of these models is assumed to be appropriate. However, the parameter values for these models are often not known perfectly. The (epistemic) uncertainty on parameters of the models is typically characterized by establishing probability distributions on the parameter values that represent the analyst's state of knowledge about the values of the parameters. These distributions can be propagated through the analysis to derive probability distributions that characterize the uncertainty on the estimation of metrics such as core damage frequency and large early release frequency. These distributions can be interpreted as expressing the analyst's degree of belief in the values these metrics could take, conditional on the underlying model being correct.

As discussed above, the development of the PRA model is supported by the use of models for specific events or phenomena. In many cases, the industry's state of knowledge is incomplete, and there may be different opinions on how the models should be formulated. Examples include approaches to modeling human performance, common cause failures, and reactor coolant pump seal behavior upon loss of seal cooling. This gives rise to model uncertainty. In many cases, the appropriateness of the models adopted is not questioned, and these models have become, de facto, the standard models to use. Examples include the use of Poisson and binomial models to characterize the probability of occurrence of component failures. For some issues where alternate models are well formulated, PRAs have addressed model uncertainty by using discrete distributions over the alternate models, with the probability associated with a specific model representing the analyst's degree of belief in that model as being the most appropriate. A good example is the characterization of seismic hazard, where different hypotheses lead to different hazard curves, which can be used to develop a discrete probability distribution of the initiating event frequency for earthquakes. Other examples can be found in the level 2 analysis. An explicit representation of model uncertainty can be propagated through the analysis as for parameter uncertainty. More typically, however, particularly in the level 1 analysis, the use of different models would result in the need for a different structure (e.g., where different thermal hydraulic models are used to determine success criteria). In such cases, uncertainties in the choice of an appropriate model are typically addressed by making assumptions and/or, as in the case of the component failure models discussed above, adopting a specific model.

PRAs model the continuum of possible plant states in a discretized way, and are, by their very nature, approximate models of the world. This results in some aleatory aspects of the 'world' not being addressed except in a bounding way, e.g., different realizations of an accident sequence corresponding to different LOCA sizes (within a category) are treated by assuming a bounding LOCA, time of failure of an operating component assumed to occur at the moment of demand. These approximations introduce biases (uncertainties) into the results.

Not all contributors to risk are addressed in PRA models. For example, certain initiating events or modes of operation are often not included in PRAs. This gives rise to what is sometimes called completeness uncertainty. It is characterized as an uncertainty because it gives rise to a bias of unknown magnitude. The only approach to dealing with this type of uncertainty when evaluating the overall risk from operation of the plant is to perform some type of analysis that addresses the missing contributions.

This paper discusses two issues that arise in relation to the use of the numerical results from PRAs in the context of acceptance guidelines such as those given in DG-1061, given the nature of PRA models and their uncertainties. The first is, whether it is, in practice, possible to generate a mathematical characterization of the uncertainty in the results of a risk assessment. The second is, if this is the case, how should the results of this characterization be used in the process of decisionmaking.

Representation of Uncertainty in PRA results

Issue: Is it, in practice, possible to establish a mathematical characterization of the uncertainty in the results of a PRA?

Historically, parameter uncertainties and some model uncertainties have been characterized mathematically using probability distributions. These distributions have been propagated to generate probability distributions on risk metrics, such as CDF and LERF. However, this characterization of uncertainty has the following limitations:

> the characterization of uncertainty is conditional on the basic structure and scope of the model. It does not capture the biases introduced by specific modeling approximations, and does not address the incompleteness uncertainties.
>
> the probability distributions on the inputs to the model, because they are subjective, are likely to differ from analyst to analyst, and so, therefore, is the characterization of the uncertainty on CDF and LERF.

To make the best use of the PRA results in making decisions, it would be advantageous to have the most complete characterization possible of the state of knowledge about the risk. The following approaches could be considered as means to quantitatively compensate for the incompleteness of the uncertainty characterization. For clarity and convenience, completeness and model uncertainties are discussed separately:

<u>Completeness uncertainty</u>

Approach 1: require that PRAs for external initiating event and low power and shutdown modes be performed to a level of detail commensurate with that of the existing PRAs.

Pro: This approach would result in the most complete characterization of risk consistent with the state of the art.

Cons: 1. The resources required to achieve this could be significant, and would tend to discourage those licensees searching for cost effective successes in the short term.

2. Because not all issues are addressed by the state-of-the-art, this would still not be a complete characterization of the risk. For example, methods for addressing issues such as organizational factors and errors of commission are still not available.

Approach 2: use a simple bounding analysis based on generic features, but modified where possible for plant specific characteristics.

Pro: In principle this would be a simple approach, that is not resource intensive to apply.

Con: 1. The degree of bias introduced is unquantifiable. This is significant because, for many licensees bounding calculations would be inadequate to demonstrate that the acceptance guidelines on CDF and LERF in particular had been met.

2. To be at all realistic in allowing for plant specific features, development of an approach may require a considerable amount of development effort on the part of the licensees or the NRC. This is of particular concern for those contributors, such as fires, which are very plant specific, and which, on the basis of those fire PRAs that have been performed, can be significant contributors to risk.

<u>Model uncertainty</u>

Approach 1: require that alternate assumptions or models for the most significant uncertain elements be included in the model in a probabilistic way.

Pro: This approach would provide the most “complete” description of model uncertainty.

Con: 1. It is a resource intensive approach, and, for some modeling issues, would require considerable modification of PRA models.

2. To be comprehensive, this would require some sort of standard to identify the important modeling issues, and to define a set of potential alternate models for each of these issues.

Approach 2: require that specific modeling approaches be used for significant areas of uncertainty, and accept the degree of conservatism or non-conservatism that results.

Pro: This approach would ensure some degree of uniformity across applications that would remove arbitrary decisions.

Con: 1. This could require significant effort on the part of some licensees if they had to modify their PRA models.

2. This approach requires decisions to be made on what the most significant modeling issues are, and which modeling approaches should be used.

3. If conservative models were to be chosen, this would be inconsistent with the Commission’s PRA policy statement, and would cast into doubt the validity of using the QHOs and subsidiary objectives as bases for acceptance guidelines.

4. For many of the PRAs, adoption of conservative modeling assumptions would lead to results that exceeded the CDF and LERF goals, and thus preclude any applications leading to even very small increases in risk.

Conclusion

While, by adopting approaches 1 above, for both completeness and modeling uncertainties, it can be argued that it is possible to characterize uncertainty more fully and in a mathematical way, there would be significant resource requirement implications, and furthermore, there are several new issues that must be addressed in relation to the use of these distributions in a regulatory environment. The resource requirements would arise because this approach would require that the current PRA models be expanded considerably in scope.

The issues about the use of the results in a regulatory environment are associated with the subjective nature of uncertainty characterization. For the use of PRA to be credible, it is essential that it be possible for the NRC to make consistent decisions. Recognizing this then, an important issue is whether each licensee should be allowed, in his characterization of model uncertainty, to provide his own degrees of belief on the appropriateness of alternate models, and to pick and choose the set of alternate models, or should this be prescribed by a standards body, or by NRC. The same questions arise with the characterization of parameter value uncertainty, although in this area, there is typically less controversy.

Moreover, it has to be realized that the biases associated with approximations will remain unquantified. Thus, when different approximations are used by different analysts, the biases on the results will also differ, in an unquantified way.

It can be concluded from this that, without some form of standardization, the characterization of uncertainty will be subjective and will differ from analyst to analyst, and to fully understand its significance, a reviewer will have to decompose the analysis to understand the basis for the characterization.

However, for many applications it might be possible to eliminate or minimize the concern about model uncertainties and incompleteness, so that the parameter uncertainties can represent the real uncertainty. This could arise when the change impacts only a small part of the PRA model for which there is consensus on how to formulate the model. The question is how best to use that information in the context of comparison with acceptance guidelines.

Comparison of Probability Distributions with Acceptance Guidelines

Issue: How would a probability distribution on CDF, LERF, or increments thereof be used to establish a level of assurance that the acceptance guidelines had been met?

Several different approaches to the comparison with acceptance guidelines can be identified. However, it has to be stressed that these methods provide the most useful input to a decision maker when it can be argued that the probability distribution addresses or bounds all the significant sources of uncertainty. As discussed at length above, there are uncertainties that are not addressed explicitly in the uncertainty analysis, e.g., those which have been dealt with by making particular assumptions, and there are biases caused by incompleteness or modeling approximations.

Approach 1: DG-1061 approach.

In DG-1061, it is proposed to use the mean values of the metrics and their increments in the comparison with the acceptance guidelines. The mean values incorporate consideration of those uncertainties explicitly captured in the model, and follow the traditional decisionmaking approach of using the expected value. If the distribution were believed to capture all the significant sources of uncertainty, and if the model were full scope, this would be a sufficient test of acceptability under this approach.

Pro: The use of mean values is conceptually simple, and is consistent with classical decisionmaking for a decision maker who is not risk averse. Furthermore use of a single point estimate, such as the mean value makes it relatively easy to determine risk neutrality for example.

Con: This use of the guidelines could be interpreted as being a speed limit approach. However, the Staff's intent was not to treat the guidelines as such, but to allow qualitative arguments to be used to allow some tolerance when the guidelines are exceeded by some amount. This is by its very nature a subjective process. Although the process is subjective, the Staff's intention is that any rationale should be technically sound and internally consistent within the decisionmaking process.

Approach 2: Use of percentile measures

If all the uncertainty were quantifiable in the sense of allowing the analyst to generate a probability distribution on the numerical result, then an intuitively reasonable approach would be to overlay the distribution on the goal, and determine at what level of confidence the goal is met. This would require a policy decision concerning what level of confidence would be acceptable.

Pro: This is an intuitively appealing approach.

Con: There are, however, several concerns with this approach:

1. The forms of the distributions for characterizing the input uncertainties (particularly parameter uncertainties) are arbitrary. In particular, the tails of the distributions can be strongly affected, since they have typically been considered as not being very significant; it is the central 90% of the distribution that generally receives attention. Their use might give a false sense of assurance particularly at high percentile values.

2. As identified above, this would require a policy decision on what level of assurance was acceptable. Historically, while no particular assurance level has been proven to be more acceptable than others, it is typical to see assurance levels of .95 as being characteristic of acceptability. However, it is not just a question of how much of the distribution lies above the goal, but how it is distributed.

For example, consider the following simplified examples, which might correspond to a representation of a modeling uncertainty:

1. the distribution is bimodal, with 5% at 10^{-3}/RY, and the remaining 95% at 5.3×10^{-5}/RY, which gives a mean of 10^{-4}/RY, which would meet the 10^{-4}/RY goal,

2. the distribution is bimodal, with 20% at 2×10^{-4}/RY, and the remaining 80% at 7.5×10^{-5}/RY, which also gives a mean of 10^{-4}/RY.

Would one of these distributions represent a case that would be more acceptable than the other? The first has a small chance of exceeding the guideline by a significant amount, the second a larger chance of exceeding the guideline by a smaller amount. To determine which is more acceptable would require that the decision maker has some measure of the consequences of exceeding the goal in a particular manner.

This particular example can also be used to illustrate a potential problem with using the mean value as a point estimate.

Approach 3

Another alternative is the approach adopted by NMSS. In this approach, the idea is that there is some tolerance associated with meeting the goal. A performance objective is defined with which the mean value is to be compared, and a higher goal, with which the 95th percentile is compared. To adopt an equivalent approach in the reactor case would require a policy decision setting a higher CDF goal, say $10^{-}3$/RY, with which to compare the 95th percentile.

Given that the DG-1061 acceptance guidelines were established with the Commission’s Safety Goals in mind, and that those goals were meant to be compared with mean values, this approach seems reasonable. The upper guideline would be more directly related to a perceived acceptable level of risk, and would not have the increased conservatism embedded within the current goals on CDF and LERF. It could, for example, be based on the QHOs directly.

Pro: This has the direct recognition that the guideline is not a simple go/no-go speed limit, and that there is a tolerance band rather than a simple speed limit.

Cons: 1. This approach requires a change in policy to determine the form of the acceptance goals.

2. The approach suffers from the same concern as the other two options in that the shape of the distribution is still not being used.

The Role of Uncertainty Analysis in Decisionmaking

Following on from the above discussions, this section discusses how information on the characterization of risk and its uncertainty obtained from state of the art PRAs can best be used in comparison with acceptance guidelines of the type given in DG-1061.

The purpose of an uncertainty analysis is to allow the licensee and the reviewer to feel comfortable that the most appropriate decision has been made, consistent with the state-of-the-art, and the decisionmakers' state of knowledge.

As discussed above, even if it were considered to be possible to characterize uncertainty in a mathematical way, the absence of a standard approach to both the development of PRA models and the characterization of uncertainty, would mean that an integrated probability distribution generated by propagating probability distributions on the elements of the PRA will be subjective, and influenced by the choice of assumptions and approximations adopted by the analyst. This conclusion is independent of the choice of acceptance criteria. For the NRC reviewers of an application that used such a distribution to demonstrate acceptability of safety principle four, it is, as pointed out by Paté-Cornell[1], important to separate the facts and the analysis from the value judgements. Thus it is essential for the reviewer to have the results presented in such a way that the significance of making specific assumptions to the decisions can be determined. In this way, the decisionmaker and reviewer can reduce the problem to the level of the evidence that is being used to make a decision, and demonstrate how that evidence is used; the evidence can usually be agreed upon, it is the interpretation that may be at variance from analyst to analyst.

Therefore, the PRA results must be presented in such a way that it is clear what contributes to risk and how the uncertainties change that picture. For example, are there contributors to risk, such as a particular cut set, or a set of cutsets, that are both affected by the proposed change and have elements, or result from modeling assumptions, that are subject to significant uncertainty?

To elaborate a little, in the same way that the CDF or LERF measures in themselves do not provide an adequate determination of the risk profile of a plant, the probability distributions on these metrics do not provide an adequate description of the impact of uncertainty for a decision maker. To understand the risk profile, it is essential to identify the contributors to CDF and LERF by, for example, interpreting the cutsets. To understand the impact of uncertainty, it is important to the decision maker a) to identify the sources of the uncertainty, b) to decide on a reasonable set of alternates that characterize that uncertainty, and c) to understand how the risk changes given those alternates. It is by considering all these pieces of evidence, and by assessing the relative degrees of belief in the validity of those alternates that an analyst can assess whether the change is acceptable or not.

The issue of completeness is of considerable concern, all the more so because the current acceptance guidelines require a demonstration that the CDF and LERF are less than specified limits whenever the change leads to a risk increase, however small. The requirement to address incompleteness issues when the changes in risk are so small that they are essentially zero, or at least small enough that they are barely detectable, may be detrimental to the implementation of risk-informed regulation. While qualitative approaches to dealing with incompleteness are allowed, it is not clear what form this would take to be convincing.

An alternative to dealing with uncertainty is, rather than trying to analyze it, to design it out of the decisionmaking process. For example, one approach, which would not be in conflict with the safety principles, would be to modify the acceptance guidelines to allow very small increases regardless of the value of CDF or LERF. The precise meaning of very small would need to be defined carefully. If this were the case, then, if the proposed change to the CLB were designed to be such that the missing modes of operation or missing initiating events would not be affected, then a major study of incompleteness would not be required. Similarly, in such cases, it might not be necessary to address all the model uncertainties, but only those that impact the evaluation of the change.

<u>Reference</u>

1. M. E. Paté-Cornell, "Uncertainties in Risk Analyses: Six Levels of Treatment", Reliability Engineering and System Safety, Vol. 54, pages 95 - 112, 1996