Uncertainty Estimation in Nuclear Power Plant Probabilistic Safety Assessment: State of the Art and Use in Decision Making

Sergio B. Guarro, Garth E. Cummings
Lawrence Livermore National Laboratory, Livermore, CA USA

INTRODUCTION

Probabilistic Risk Assessment (PRA) was introduced in the nuclear industry and the nuclear regulatory process in 1975 with the publication of the Reactor Safety Study by the U.S. Nuclear Regulatory Commission [1]. Almost fifteen years later, the state of the art in this field has been expanded and sharpened in many areas, and about thirty-five plant-specific PRAs (Probabilistic Risk Assessments) have been performed by the nuclear utility companies or by the U.S. Nuclear Regulatory Commission. Among the areas where the most evident progress has been made in PRA and PSA (Probabilistic Safety Assessment, as these studies are more commonly referred to in the international community outside the U.S.) is the development of a consistent framework for the identification of sources of uncertainty and the estimation of their magnitude as it impacts various risk measures. Techniques to propagate uncertainty in reliability data through the risk models and display its effect on the top level risk estimates were developed in the early PRAs. The Seismic Safety Margin Research Program (SSMRP) study was the first major risk study to develop an approach to deal explicitly with uncertainty in risk estimates introduced not only by uncertainty in component reliability data, but by the incomplete state of knowledge of the assessor(s) with regard to basic phenomena that may trigger and drive a severe accident [2,3]. More recently NUREG-1150, another major study of reactor risk sponsored by the NRC, has expanded risk uncertainty estimation and analysis into the realm of "model uncertainty" related to the relatively poorly known post-core-melt phenomena which determine the behavior of the molten core and of the reactor containment structures [4].

The question of how the risk/safety assessor and the decision maker should view the issue of uncertainty in risk estimates is central to the more general question of how to effectively use PSA results in the safety and regulatory decision making process. The current approach adopted in the industry and regulatory practice uses only a portion of the information that is gathered in a modern PRA/PSA. The thesis of this paper is that decision criteria can be adjusted to respond to the different levels of uncertainty found in PRA/PSA results, and that this adjustment can reflect practical concerns if a distinction between "reducible uncertainty" (i.e., uncertainty that can be reduced in the near future by further investigation) and "irreducible uncertainty" (i.e., uncertainty that cannot be readily reduced in a foreseeable time-frame) is made in classifying and displaying uncertainty. Such a distinction would for example allow decisions to be placed in a time-order perspective. The SSMRP study was a partial precursor of this concept, as it introduced a distinction between "random uncertainty" and "model uncertainty". Our thesis is discussed in light of the approaches to uncertainty analysis taken in the SSMRP, in NUREG-1150 and in other PSA studies, and by discussing for a simple, "common life" situation, the information value contained in probabilistic knowledge with different levels of uncertainty associated with it. We also discuss how decisions based on central tendency estimates may be inadequate when large uncertainties are present.

REVIEW OF UNCERTAINTY ANALYSIS APPROACHES IN PAST PSA STUDIES

Uncertainty in probabilistic risk assessment has been recognized as an important issue from the very beginning, which is not too surprising, as the use of probability-based estimates is in itself one of the few means we have to quantify imprecise or incomplete knowledge. Within the context of PRA, however, "uncertainty" signifies "uncertainty of a second order", that is, the estimated range of variability of the probabilistic estimates themselves (this is why some refer half-seriously to PRA uncertainty as the "probability of a probability").

The authors of WASH-1400, the first nuclear power plant probabilistic risk assessment study [1], propagated uncertainty in basic data (e.g., the ranges between the 5th and 95th percentiles of plant component expected failure frequency distributions) to the top of the system model fault trees and from there through the event trees to express coremelt frequency and radiation release doses. This was accomplished by Monte Carlo simulation techniques and yielded uncertainty ranges for the expected coremelt frequency, release doses, etc. The major limitation of WASH-1400 was that it did not explicitly address the issue of "model uncertainty", that is the uncertainty related to assumptions not supported by complete knowledge on the course and severity of accident sequence phenomenology and on the effectiveness of mitigative safety devices. Generally speaking, the WASH-1400 methodology tended to underestimate uncertainty, by limiting its analysis to the propagation of uncertainty related to failure rate estimates (i.e., "data uncertainty"), although an embryonic element of model uncertainty had been introduced in the study by assigning event tree branch "decision points" for ECCS intervention and estimating a probability for the failure of this safety function, as possibly resulting from unassessed phenomena).

SSMRP attacked the problem of model uncertainty by explicitly accounting for it in a two-tier analysis approach , which included both the use of expert opinion and an experimental approach to construct uncertainty bounds and account for the impact of different assumptions and system models on risk measures [2]. Specific techniques were developed to carry through this approach in the analysis of seismic events, because under these conditions component failure modes become correlated, as they are driven by a common cause initiator (i.e., the earthquake). The SSMRP basic methodology was developed for the analysis of seismic risk, but its fundamental techniques for treatment of correlation and model uncertainty could also be applied in analyses of risk related to other "external event initiators" (e.g., tornadoes, floods, etc.). A demonstration of the SSMRP approach was done by applying the method to an analysis of the Zion Nuclear Power Plant [3]. One of the more dramatic effects of accounting for model uncertainty in such an explicit way was that the span of the 90% confidence interval for risk parameters appeared to grow (from previous less formal assessments) to a few orders of magnitude. In the Zion study, the median coremelt frequency was assessed at 3E-5 per year, with upper (90%) and lower (10%) bounds of 8E-4 and 6E-7 per year. Interestingly, the predominant reaction of the nuclear industry community to these findings of large assessed uncertainty was one of skepticism.

The Indian Point PRA was the first industry-sponsored PRA to propose an explicitly formulated approach for dealing with expert opinion and model uncertainty. It allowed assessors to express their estimates in the form of distributions for risk parameters of interest rather than point estimates, and formalized the uncertainty propagation techniques by propagating the entire distributions in discretized form, which in turn allowed the representation of data or model uncertainty by distributions of arbitrary shape, not just lognormals as was customary in previously produced PRAs.

The recent NUREG-1150 study has specifically tackled the problem of generalizing and improving the approach first proposed in the SSMRP, by identifying major model uncertainty "issues" in the system and phenomenology models adopted in the risk study [4]. In this approach multi-outcome nodes in event trees which model the course of an accident sequence are defined and expert opinion polling sessions are used to assign weights to the various hypothesized outcomes. In other words, a discretized distribution of "outcomes" is generated by expert opinion for every issue. This uncertainty, along with the "traditional" data uncertainty is propagated through the event tree risk models by Latin hypercube sampling techniques (a more efficient variety of the Monte Carlo simulation method). Just like in the SSMRP the uncertainty range of the NUREG-1150 results spans typically three orders of magnitude, and again many in the industry have reacted by questioning the validity of such an assessment (i.e., they feel that uncertainty is being overestimated).

THEORETICAL AND ACTUAL USE OF PSA IN DECISION MAKING

A consistent approach to the handling of uncertainties within a probabilistic safety assessment study cannot be taken if the objectives of the study have not been clearly established by the decision-making party to whom the study is addressed. In turn, the decision maker cannot establish decision criteria that fully take advantage of the information contained in PRA/PSA studies if he/she is not fully aware of the value of this information (including the presentation of the uncertainty in the assessment), and knowledgeable about the interpretation (or interpretations) that may be given to the information received.

In the commercial nuclear arena, the most important decision making process is the one related to plant licensing, which sees the Nuclear Regulatory Commission as the party required by law to be the decision-maker. After the publication of WASH-1400, the NRC began to consider how decision criteria could be developed and applied, for licensing and plant retrofitting purposes, which would use the information obtained from nuclear power plant probabilistic risk assessments. In essence, it was suggested that the NRC's decision criteria concerning plant licensing and operation could be based on the plant compliance with a set of "Safety Goals", established by the Commission and stating various maximum limits for risk measures expressed as probabilistic expectations of core-melt frequency and health-effects on the public.

The NRC's objective of establishing a "safety goal policy" generated much debate (and criticism) in the nuclear community. Comments ranged from support for the idea of establishing such goals [6] to skepticism about the political and technical viability of the safety goal policy objective [7]. Also, industry and research community suggestions for the type of goals to be established varied widely, both in the type of approach suggested and in the quantitative measures proposed. In 1983 the Commission published a policy document containing a set of proposed safety goals for trial evaluation [8]. The current policy was established in 1986 as the outcome of this trial period [9].

A recent paper reviews and discusses both the "trial" and the current safety goal policy [10]. We will comment here only on the aspects of the safety goal policy which relate to uncertainty estimation and its use in decision making. Interestingly enough, neither the trial goals nor the final ones directly address the question of uncertainty. This is per se a major deficiency, and it is interesting to observe that the discussion on the treatment of uncertainty in the safety goal policy statement is viewed by most as one that has been handled indirectly by the choice of "distribution parameter" used for the decision making process. Bier, in the reference just given, comments on the switch from the use of median values (e.g. the median of the distribution representing the core melt frequency of a particular plant) in the trial goals, to the mean in the final goals. She notes that the statement by the NRC staff commenting on this issue [Stello 1986: "The choice of which value {to use} is of course judgmental and without a clear cut theoretical basis."] "is simply not true" and presents an argument in favor of the use of mean values as opposed to medians or other percentiles, based on the observation that very broad distributions may have mean values higher even than the 95th percentile. Although this is mathematically correct, we should observe that the distributions adopted in PRA/PSA studies are mathematical models which should be carefully validated at the two tails to verify that physical impossibilities are not exceeded. Even more important, we do not think it is conceptually and practically correct to approach the issue of uncertainty by attempting to reduce an essentially two-dimensional problem (e.g., "location" and "width" of the distribution) into a one-dimensional one (e.g., location of the mean, median, or any other single parameter). Such an oversimplification can be expected to fail in any situation where the nature of the risk is such as to produce distributions with shapes strongly stretched and skewed in one direction or the other.

In the following we try to show that the assessment and display of uncertainty may remain void of application and significance if the decision maker does not actually intend to declare and maintain a "risk bias" based on the magnitude of risk uncertainty and on the severity of the expected loss consequences. In other words, the decision maker can make good use of uncertainty information only if his/her decision criteria are formulated as explicit functions of both conservative or non-conservative bias, and of the degree of uncertainty in the estimated decision parameter itself. Conservatism or non-conservatism is important because society makes judgments on "perceived", rather than "objective", risk. For the purpose of clarification we illustrate in Figure 1 the difference between an objective measure of risk (the distribution of core melt frequency for a hypothetical plant given in Figure 1a), and the same measure as perceived by a viewer who is unfavorably impressed by the high-end tail of the distribution (i.e., a "conservative" or "risk averse" viewer) and who will tend to "see" the distribution as shifted towards higher values as shown in Figure 1b. Similarly, Figure 1c represents the perception of a viewer who is relatively unconcerned about that same high-end tail (i.e., a "non-conservative" or "risk prone" viewer) and who will tend to see the distribution as shifted in the opposite direction. The effect of conservatism is essentially that of weighting the "objective" distribution with a positive-gradient weight function (as shown in Figure 2a), whereas that of non-conservatism is of weighting it with a negative-gradient weight function (as shown in Figure 2b).

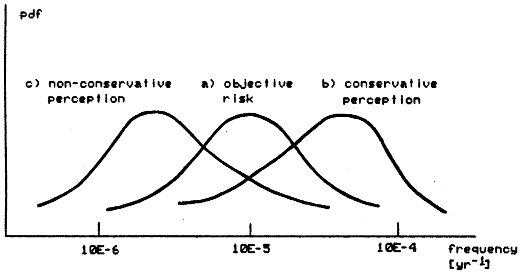

Figure 1: Risk and risk perception

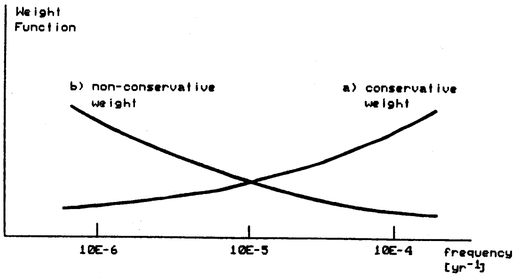

Figure 2: Risk weighting functions

In our day-to-day living our degree of conservatism is affected both by the location and by the width of a "risk distribution"; we become more conservative both as the consequence dimension of risk increases and as the uncertainty spread of the risk measure becomes larger. A simple example may serve to illustrate the point. Let us assume that we are invited to gamble an amount of money X on a head vs. tail outcome for a coin toss: if we know that the coin is a fair one, that is, that it will always give us a 50/50 chance of winning, our decision on whether to accept the gamble will be normally based on what ratio exists between X and the amount of money Y we have at our free disposal for gambling. Now let's assume that we are told that the coin is biased so that the odds are no longer 50/50 but 70/30 one way or the other. With the same ratio X/Y of bet to available money as before, we should behave exactly in the same way if we are neither conservative nor non-conservative. However, if we are conservative, this new dimension of uncertainty will tend to make us view our odds more pessimistically, as we would be unfavorably impressed by the possibility of the odds being 70 to 30 against our choice of coin toss outcome. Conversely,if we tend to be optimistic we will perceive the deal as giving us the opportunity of betting at odds of 70 to 30 in our favor. This example can be illustrated mathematically using a discrete probability distribution for the coin toss odds , a decision criterion for accepting or refusing the gamble based on the X/Y ratio, and conservative/non-conservative bias weight functions similar to those depicted in Figure 2. The main point, however, is that uncertainty is totally irrelevant to our decision process unless we establish: a) how much we can afford to lose, and b) whether in our decision we want to reflect a conservative, or non conservative attitude. Note again that if we decide to apply no bias and are able to maintain "objectivity", the existence of uncertainty regarding the direction of the coin bias should have no effect on our decision to accept or refuse the gamble, as the overall odds for or against choosing the winning side remain the same (i.e., 50/50) as in the situation with the unbiased coin. It is perhaps important to notice that human decision makers are very rarely objective. This is not a surprise to social scientist and psychologists, and shouldn't be to mathematicians or statisticians either, as we saw in our example that the mathematical representation of an "objective" decision maker is that of a straight horizontal-line weight function, a very special case indeed among all the possible shapes and inclinations of weight function representation curves.

Clearly, addressing some of the issues raised in the preceding discussion in a practical and politically feasible manner, within a policy for application of quantitative nuclear power plant safety goals, presents several difficulties. More specifically, it may be difficult to express such complex and delicate issues as risk and uncertainty aversion within the formulation of a few simple "rules". Not addressing the issues, however, may be far worse, as it leaves the safety goal policy open to criticism on these important questions, and may make, in the long term, its practical application impossible.

DISCUSSION OF DIRECTIONS FOR A DECISION MAKING FRAMEWORK

The distinction between "data uncertainty" and "model uncertainty" has been conceptually useful to the PRA analyst because it separates uncertainty principally related to lack or poor quality of data (mostly component failure rates) from uncertainty resulting from (presently) unverifiable assumptions and models. However it is not unreasonable to contend that the end user of a PRA analysis, e.g. the decision maker, would be perhaps more interested in the distinction between the portions of uncertainty that can be reduced (e.g., in the case of data uncertainty, by means of better data collection efforts and by the growing with time of the recorded plant observation periods, and, in the case of model uncertainty, by the setting-up of phenomenology research efforts and accident simulation experiments). In some cases the distinction between "reducible" and "irreducible" uncertainty may turn out to coincide in practice with the one between data and model uncertainty, as it may be indeed very difficult to substantially reduce the uncertainty associated with many of the complex phenomena that are relevant to the potential progression of severe reactor accident sequences.

Somewhat disappointing to many in the PRA/PSA technical community is the fact that, in most past decision making situations related to reactor safety and operation, uncertainty information has not been explicitly used, and the decisions made have been justified principally on the basis of the value of a central tendency measure (such as the median or the mean). An example of this is the inquiry conducted by the Atomic Safety and Licensing Board (ASLB) [11] on the Indian Point nuclear power plant. The ASLB received several days of depositions by members of the NRC Staff and their consultants on the magnitude of the uncertainty in the Indian Point PRA and in the estimates of risk made by the NRC Staff itself in relation to that plant. In the end, however, the recommendation made by the NRC Staff (which was essentially accepted without variations by the ASLB) was based merely on mean risk estimates. At the time of the inquiry, a framework for taking into account uncertainty information in decision making had not been developed. Worse yet, it still hasn't.

The most useful consideration on uncertainty, which we are confident could be put to almost immediate use, is based on the concept of "reducible" versus "irreducible" uncertainty. The coin toss example is still useful to illustrate this concept. If we knew that we could in some way analyze or observe the biased coin, our "conservative" decision would probably be that of postponing the decision on whether to gamble until we were in fact able to discover something more about the coin. Conversely, a "non-conservative" decision could be that of starting to gamble and keep the "adversary" engaged until the new information became available. Finally, if we had no confidence of gaining any more knowledge about the coin, postponing or not postponing the decision to gamble would offer no potential benefit whatsoever.

Analogous situations exist in nuclear licensing decision making, although they have obviously little to do with gambling and the factors affecting there the decision process are infinitely more complicated than in our simple coin-toss example. In the presence of large "reducible" uncertainty in the risk posed by a specific plant, one which would also include the possibility of exceeding acceptable safety limits, the conservative decision could be that of not granting a plant license until the uncertainty could be reduced and compliance with the safety limits could be demonstrated at a high level of confidence. The non-conservative decision would be, under the same circumstances, that of granting the license and assume that additional knowledge would later become available to support licensability.

In the past it has been difficult to agree on whether the NRC was following conservative or non-conservative criteria in plant licensing, mainly because not much agreement could be reached on what quantitative safety goal values could be really considered "acceptable". In practice the NRC has often licensed plants in the light of uncertainty on bounding assumptions and sought to justify these assumptions through confirmatory research. In other words the NRC has often tried to identify "worst case scenarios" and verify that under such extreme scenarios a plant doesn't pose undue risk to the public. Unfortunately this strategy reverses the problem back into the deterministic decision domain from which the regulatory process was supposedly set to move away. The keeping of a deterministic decision criteria strategy is supported by many in the nuclear community in light of the high variability and "lack of robustness" of PRA/PSA results. Variability of PRA/PSA results is undeniable. For example, it appears that the results of the still unpublished second draft of the NUREG-1150 study show in many instances major accident sequence contributors to overall risk which are significantly different from those identified in the first draft. The question is whether this variability is really surprising, given that these results proceed from sampling of expert opinion taken to account for accident sequence model uncertainty. In the application of classical statistics, nobody is surprised when different samples of a distributed population or experimental setup show large variability. If one takes the view that a PRA or PSA is not a periscope into the future but a sample, taken on our current state of knowledge about a set of very complex systems and phenomena, then the variability of results is simply a confirmation that our state of knowledge is indeed uncertain, and that well set deterministic assessments on these systems and phenomena are not presently credible. If one accepts such a view the next question is then whether we can make valid decisions in the presence of large uncertainties. We think the answer is yes, as human society has historically adapted to uncertainty in the information available for decision making by using "adjustable biases". As we have argued in our example, these may increase or decrease the amount of conservatism which is applied in making the much needed decisions, in light of the severity of the expected consequences of negative events and of the magnitude of the uncertainties on such consequences -- especially those uncertainties which do not appear to be easily reducible in the short term future.

In essence, realistic safety goals should address both real and weighted risk. The decision makers should declare which bias they intend to apply (in view of the perception of nuclear risk by public opinion, however, it is not likely that the Nuclear Regulatory Commission may consider to declare and take a "neutral" or "non-conservative" stance). PRA/PSA research, on the other hand, should concentrate on developing and perfecting techniques to "separate", in final result displays, uncertainties which are deemed reducible from those that are not likely to be. The distinction between reducible and irreducible uncertainty is more important in a practical sense than the typical distinction made between "data uncertainty" and "model uncertainty". Decision criteria should thus include cost/benefit assessments to assign the allocation and prioritization of backfits and of research and data collection resources, into areas of plant design and analysis where the greater impact in terms of both risk reduction and risk uncertainty reduction is possible. Some recent NRC studies, including NUREG-1150 have indeed sought to follow this principle [4,12], and one must hope that these efforts will continue in the future and be able to be conducted in close coordination and integration with PRA/PSA studies conducted by NRC and industry.

CONCLUSIONS

Although the safety-related regulatory decision-making process is affected by many non technical factors, we believe that the relation between risk uncertainty and the triggering of decision bias factors can be recognized and investigated in its "technical" aspects -- i.e., those which can be identified "objectively" through the observation of common societal attitudes both towards risk and the uncertainty in the information available to assess risk. A set of effective decision making criteria cannot be formulated independently of a technical and objective understanding of the relations between risk, uncertainty and the degree of conservatism that society and the decision maker(s) may be calling for in the presence of different combinations of these factors. A better understanding of these technical aspects is crucial to support the use and credibility of probabilistic safety assessment as a decision support tool.

REFERENCES

1. 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
2. G.E. Cummings: "Summary Report on the Seismic Safety Margins Research Program," UCID-20549 (NUREG/CR-4431), January 1986
3. M.P. Bohn, et al.: "Application of the SSMRP Methodology to the Seismic Risk at the Zion Nuclear Power Plant," UCRL-52746 (NUREG/CR-1489), May 1980
4. U.S. Nuclear Regulatory Commission: "Reactor Risk Reference Study," NUREG-1150 (Main Report, Draft for Comment), February 1987
5. Power Authority of the State of New York; Consolidated Edison Company of New York: "Indian Point Probabilistic Safety Study," 1982
6. "Chairman Parker to President Carter: Approval, but 'Major Disagreements,'" Nuclear Industry 24, 4-5, 1977
7. J.M. Griesmeyer, D. Okrent: "Risk Management and Decision Rules for Light Water Reactors," Risk Analysis 1, 11-27, 1981
8. U.S. Nuclear Regulatory Commission: "Safety Goals for Nuclear Power Plant Operation," NUREG-0880, Rev. 1, May 1983
9. U.S. Nuclear Regulatory Commission: "Safety Goals for the Operation of Nuclear Power Plants; Policy Statement," Federal Register 51, FR 28044, August 4, 1986
10. V.M. Bier: "The U.S. Nuclear Regulatory Commission Safety Goal Policy: A Critical Review," Risk Analysis 8, 563-568, 1988
11. Atomic Safety and Licensing Board: "Special Hearing on Indian Point Unit No. 2 (Consolidated Edison Company, Docket No. 50-247-SP) and Indian Point Unit No. 3 (Power Authority of the State of New York, Docket No. 50-286-SP), 1983
12. J.W. Pittman: "A Prioritization of Research Activities," NUREG-1319, December 1988