

**USING SOFT DATA TO MAKE ‘PROBABILISTIC  
RISK ASSESSMENTS’ REALISTIC**

Rex Brown

THE INSTITUTE OF PUBLIC POLICY  
GEORGE MASON UNIVERSITY

MS 3C6, Fairfax, Va. 22030-4444  
rbrown@gmu.edu

April 1999

## Abstract

A regulator or risk manager needs to assess the safety of an activity's realistically, in a way that takes into logical account all knowledge that he believes to be relevant. He will wish to build on all considerations, however inconclusive, that he *already* uses in his decisions (and does not wish to ignore). He will also want to use any additional data he can develop or has access to, for example, using Probabilistic Risk Assessment. However, PRA, as commonly understood, analyzes only those sources of risk on which there is hard data, and only takes hard data into account in assessing those risks. Responsible executives are not satisfied to simply equate the uncertainty they will use in their decisions to a PRA, though they should certainly take that into account. We propose a methodology, grounded in statistical decision theory, which accommodates soft as well as hard data. It is illustrated in the context of nuclear and other environmental regulation.

## Problem

A regulator or facility operator is commonly called upon to evaluate how safe a facility is, and whether that level of safety is acceptable. The analytic support he has to draw on is usually limited to a choice of: either introducing some objectivity into his safety evaluation, by relying on what is commonly known as PRA, or not doing any analysis at all. Senior risk managers and regulators I have worked with (e.g. in reactor safety and nuclear waste isolation) often opt for the latter option—no formal analysis. They do not accept the PRA as a realistic substitute for their professional judgment on a facility's safety. They are concerned in particular, that human behavior is inadequately considered or that important knowledge they have has been neglected.

On the other hand, they may feel quite comfortable with PRA for purposes *other than realistic risk assessment*. Murley (1994) notes: "One should view a PRA estimate of the core damage frequency for a plant more as figure of merit for that plant's design, rather than a prediction of the actual likelihood of a core damage accident at that plant".

Performance requirements may take the form of an acceptable hazard probability, like 1/10,000 of core damage. A regulator or other evaluator needs a way to assess the probability that he considers complete and realistic, and open to review or challenge. The approach needs to include all the knowledge and judgment the regulator would normally use when; for example, he decides whether to close down a nuclear plant--and hopefully to add to it. How can he enhance and make more transparent what he does now? Usual procedure would rely on his professional judgment, paying cautious attention to available PRAs. Above all, the regulator must be assured that the resulting assessment is no worse (i.e. less realistic) than the assessor's unaided judgment.

The current state of PRA art, however valuable for many purposes, does not meet that test. NRC regulators are often reluctant to replace prescription by probabilistic performance requirements. This is usually because they are not prepared to use PRA as a realistic "best" probability that will serve as the sole test of compliance. For example, PRA may omit intangible sources of risk that are known to be important. Apostolakis (1993) has noted that PRAs do not always distinguish between plants with good and poor organization.

It is said that such problems will be ironed out in time as PRA matures, with better data and methods, but the regulator cannot wait. Moreover, some problems may be essentially beyond the reach of the PRA paradigm, because it asks different questions and limits itself to certain contributions to the answer.

### What is missing in conventional PRA?

*Current PRA paradigm.* Conventional PRA provides a promising core for realistic safety performance assessment, but it is incomplete. I believe we can characterize PRA in its purest form as follows. It addresses only those sources of risk on which "hard" data is available (however soft or suspect their assumed relation to reality). In addressing these risks it disregards other knowledge. As in traditional reliability engineering, it treats risk as a process that generates hazardous events with some frequency to be estimated (as in "core damage frequency"). Although leading exponents of PRA claim to have been striving to avoid these limitations for more than a decade (Garrick and Kaplan, 1988), I contend that they still characterize most current practice, including five PRAs which have been presented as state-of-the-art in NRC regulatory guidelines (NUREG-1150).

This process is the result of contributing risk processes whose frequencies therefore are to be estimated. Some processes (like component failures) have relevant experimental and other data on which to base a frequency estimate. Other risk frequencies do not yet have enough data to estimate, but we expect them to yield eventually to empirical research. PRA only estimates those frequencies it has data on<sup>1</sup>. Others are omitted because "we don't have the data." Practitioners strive to improve data of the risks covered, and develop data on the others, but in the meantime acknowledge that the frequency estimate does not cover all sources of risk.

The frequency of the risks that *are* addressed is normally estimated by treating the data as representative of the process; or related to it by assumptions (e.g. that failures in the field equate to failures in the laboratory). Other knowledge, for example on any difference between assessed and similar reference situations, is not taken into account. Exclusions may include informal experience, common-sense arguments or "soft" data, such as regulatory inspector reports.

---

<sup>1</sup>But see discussion below on efforts to include expert opinion.

“Probabilistic Safety Assessment” (PSA) becoming the term of choice for quantitative risk assessments in general in the international community. I suggest, therefore, letting the letters PRA stand for “Partial Reliability Assessment. I believe that expression characterizes most of what has traditionally passed for PRA (in spite of moves to extend it I noted).

This PRA is successfully used in many applications, for example to prioritize hardware and other risks on which there is substantial empirical data that swamps other softer data. It should be useful in evaluating the *relative* safety of generic reactor designs, which are largely unaffected by licensee-specific risks (say from deficient management practice), or by other sources of risk (say external incidents) which are likely to affect contending designs equally. However, regulators, risk manager and many PRA practitioners recognize that it will not do, as it stands, to evaluate overall plant safety realistically.

*Risks beyond PRA paradigm.* Some risks cannot usefully be viewed as stable, repeatable processes (e.g. the response of a particular plant management to an accident). In any case, the risk may imperfectly correspond to any frequency that can be observed (e.g. as generated by a simulator). The linking assumptions used may not be substantiated or realistic (e.g. the assumption that management "goes by the book").

It is also true that there are significant current efforts to extend the scope of PRA, to address the gaps, for example using expert judgment to assess model parameters, Bayesian statistics and good old "engineering judgment". However, I believe a more radical, full-blooded paradigm shift is called for if a realistic probabilistic assessment of risk is called for.

#### What is needed to complete PRA?

There is a growing recognition in the field that a more comprehensive form of PSA is needed which produces a probability (not frequency) that the assessor can accept as realistic. It should address all sources of risk, including unique events, which do not have frequencies to be observed. It needs to take into account the assessor's best judgment, together with whatever information he has or has access to (however tentative or controversial). This will include but rarely be limited to available PRAs. Time and resource pressures oblige the model to be very coarse grained.

This version of PSA might be called *Comprehensive Uncertainty Evaluation*, abbreviated to CUE (since it does, indeed depend critically on *cues* from whatever source). The underlying theory is well-known personalist<sup>2</sup> (rather than frequentist) statistics. As with PRA, the target event (like core damage) in CUE is the result of contributing events (like component failure and management error), and these can be modeled in the

---

<sup>2</sup> Also known as "Bayesian" statistics, somewhat misleadingly (since it may not involve Bayes theorem).

same structure, but at high level. But now, the assessor, in order to be as realistic as possible, assesses probabilities for all sources of risk, however tenuous the data, and brings to bear whatever he knows or believes. Another assessor, with different beliefs and knowledge, will come up with different target assessment, just as he would in normal regulatory or risk management practice.

*Implementing CUE.* Within the personalist paradigm, CUE can be variously implemented. Incompleteness in the knowledge used in PRA can be addressed informally by simply having the assessor adjust raw PRA values in the light of his best judgment. (This is useful when PRA is used as a private decision aid, rather than, say, public testimony in legal proceedings--though even here there may be room for realistic judgmental adjustments). It can be addressed formally by pooling it with alternative assessments. For example, precursor history can provide an independent assessment of core damage probability (Murley, 1994), subject to its distinctive errors (such as typically small samples).

Realistic uncertainty is normally greater than that derived in PRA that only addresses certain uncertainty parameters. CUE would address all sources of risk, and also the "shiftability" of assessments—i.e. the extent to which assessments are liable to revision in the face of new evidence. Frequency may be replaced by some interpretation of "true" probability (Brown, 1990).

[FIGURE 1]

Figure 1 shows how CUE relates to and builds upon PRA. The domain of PRA is to the left of the vertical dashed line, though commonly the only grounds for risks that PRA addresses are hard data and assumptions, shown in the two bottom left boxes.

*Resources required.* A common, but understandable, misconception is that, because CUE extends the scope of PRA, it is more costly. Paradoxically, the opposite is usually the case. CUE can be aggregated at whatever level available resources and timeliness require. Default inputs are judgmental, based on whatever the assessor knows, including available PRAs. When a pressing regulatory decision has to be made, the supporting CUE will typically be done with a few man-weeks of effort, whereas PRAs normally take at least several months.

In the extreme, CUE can be performed in a few *days*, by modeling directly the knowledge of an executive team charged with deciding if a reactor is acceptably safe<sup>3</sup>. A "quick-and-dirty", which might still be better than nothing, could even be for an experienced manager to make global assessment of the *proportion* of total risk omitted, and adjust a conventional PRA accordingly. ("PRAs for plants of this type typically understate total risk by 50-200%, but this PRA seems more complete than most, so let's say 30-150% for performance evaluation purposes".

---

<sup>3</sup>For example, management teams in non-nuclear fields are increasingly using interactive computerized "decision conferences" to model their judgment and produce a major decision in 2-3 days (Porter, 1987).

*Case study: the need for a backfit.* A backfit evaluation was needed for a pressing NRC decision on whether to close down a certain reactor or require some safety backfit. CUE was conducted for senior NRC staff for under \$50,000 (Brown and Ulvila, 1988), building on an earlier large-scale PRA, whose action implications were not clear to the regulator. That PRA had produced an expected core damage "frequency"<sup>4</sup> of 5E-5 (i.e. 5/100,000) that met a published safety goal of E-4. It addressed only internally initiated events and assumed (among other things) that operator responses followed approved procedures. A 90% "uncertainty range" of 2E-5 to E-4 was reported (still meeting the safety goal), but it reflected uncertainty only about parameter values within the PRA models used.

The key regulator had not judged that this exercise was very useful. Not only was he reluctant to accept its optimistic output as it stood, but he was not clear how to "splice" in other knowledge and opinion he had access, for example about missing considerations. He was planning to recommend closing down the unit pending acceptable backfits, based largely on inspection reports and on his and others' judgment.

A CUE analysis was performed which extended the PRA model to address all sources of risk. In particular it allowed for probabilistic assessments of operator error and externally initiated events and allowed for the possibility of modeling error. The assessments were elicited from NRC staff familiar with the plant and from research contractors who had worked on the issues in question. The regulator's own judgment was introduced by means of an interactive computer program, which allowed him to second-guess and override specific inputs.

At the end of the exercise, *his probability* of core damage was 3E-5, with a margin of uncertainty of 3E-6 to 2E-4. Since the upper end failed the E-4 safety goal he decided (as he had been inclined) the plant should be closed down pending either an acceptable backfit or additional evidence, which would bring his uncertainty range entirely within the safety goal.

#### More detailed hypothetical illustration

[FIGURE 2]

Figure 2 shows more fully a possible implementation of a top-down CUE case, and its relation to the PRA it builds on. In each box the lower dotted bar gives a mean and uncertainty range from PRA, and the upper solid bar from CUE. For concreteness, scales (exponential) are shown and the bars are intended to have quantitative meaning, but are quite imaginary. *[Logical consistency of bars as quantified needs checking]*

---

<sup>4</sup>The details of the real case have been changed to preserve confidentiality.

Thus, the final box at the top shows a frequency distribution from PRA of E-6 to E-3 and a probability distribution of  $5E-5$  to  $E-2$  from CUE. I.e. CUE has raised the mean risk assessment by two orders of magnitude, but CUE uncertainty is still about 2.5 orders of magnitude (mainly because several independent estimates have been pooled, which offsets the additional sources of risk and therefore uncertainty).

Both CUE and PRA assessments are derived from lower order assessments in boxes below. Here are some possible considerations that may have entered into these assessments.

At the lowest level, internal operating risk is shown in the bottom left hand box. The PRA assessment there (mean  $2E-6$ ) is assumed to be the output of an "Individual Plant Evaluation" required from each plant, though it is based heavily on generic data, in the absence of plant-specific data. The uncertainty (two orders of magnitude) derives solely from error in model parameters (not, say, in model structure).

The CUE assessment of that internal risk shows a lower probability ( $5E-6$ ), because the assessor allows for conservatism in the PRA, judges the plant to be better than average and takes account of a recent safety backfit after the PRA was done. Uncertainty is larger (three orders), because it includes model error (i.e. the difference between reality and the model even if ideally parameterized model). The CUE is assumed to be a direct adjustment to the PRA output, i.e., highly aggregated. It could equally well be the output of a more disaggregated model, adjusting intermediate PRA assessments, such as emergency core cooling performance. External operating risk in the next box over is based on a PRA for external events that was reported in the plant operator's required "Internal Plant Evaluation" (mean E-5, uncertainty E-6 to E-4). It shows CUE with the higher expected risk ( $5E-4$ ) and larger uncertainty (E-6 to E-4), because it includes sabotage omitted from the PRA due to "lack of data."

Moving up one level, internal and external operating risks are routinely combined with a standard statistical algorithm to give total full power risk, in both cases. Moving right, less-than-full power (mainly shutdown) has not been addressed in PRA (yet), so there is no PRA bar. The CUE uncertainty ( $5E-7$  to  $2E-2$ ) is large, having no PRA to build on.

Going up another level, routine combination produces two core damage risks. So far we have assumed standard management performance in all previous assessments. CUE shows an upward adjustment to risk (from none to a hundred-fold). It reflects an assessment--very uncertain--that the management of this plant is below par, based on the most recent Systematic Appraisal of Licensee Performance (SALP), itself based largely on inspection reports. The implied CDP adjusted for management deficiencies (E-3, and 7 orders uncertainty) is not shown, but incorporated in the final CUE stage.

In addition to the PRA based assessment above, CUE incorporates two independent assessments of the same core damage probability. One is based on an analysis of past accident precursors. It happens to show the same expected risk (E-3), but the assessor

has relative high confidence in this approach, hence the shorter uncertainty bar (2.5 orders). The third assessment (Exp E-4, uncertainty of 2.5 orders) reflects the regulator's firmly held judgment, based on recent visits and 20 years of experience, that this plant is safer than most.

The final step (top box) leaves the PRA unchanged (from two levels below). The CUE assessment combines the three separate assessments below, using existing methods (Brown and Lindley, 1986). It is basically a confidence-weighted average of the three approaches (Brown and Pratt, 1994). It produces a significantly higher core damage risk than the PRA (E-3 vs. 5E-5), but, interestingly, no greater uncertainty (two orders). Even though each assessment being pooled had large uncertainty, they tended to confirm each other, which reduces uncertainty.

Of course, these notional assessments have no substantive significance, since they came arbitrarily out of my ill-informed head. But they may show the kinds of reasoning that would produce a CUE assessment, adapted from a conventional PRA.

### Discussion

*Logical foundations.* The general methodology of CUE is well established, with the statistical decision theory paradigm (Raiffa and Schlaifer, 1962; Lindley, 1972). No great innovation is needed to adapt it to the special needs of risk assessment, but I have introduced some non-standard wrinkles, which others may take issue with, such as the construct of impersonal probability (Brown, 1993).

The analytic paradigm proposed is basically the formalization of common sense. Any reasonable informal argument can be accommodated. For example, "radioactive release through human intrusion depends on whether an economic use for nuclear waste is found" can be represented by a conditioned assessment model. "I used to think no unsophisticated intruder could break into a reactor, until the TMI van incident" could be represented by a Bayesian updating model.

*Handling uncertainty.* The most audible complaint against incorporating questionable professional judgment into a model of safety performance is that it "*introduces uncertainty*". It does no such thing: An assessor's uncertainty about risk is irreducibly there, due to limited data and ambiguity about its significance, whether or not he takes it into account quantitatively. Modeling only those parts of the problem which are relatively "certain" does not make the remaining uncertainty go away; it simply conceals it from view. The same is true about disagreements among assessors (e.g. licensees and interveners). Little is gained by "killing the bearer of ill tidings"!

CUE can be further extended to *model the assessment uncertainty* itself, involving even greater subjective judgment (Brown, 1989). This can pinpoint the sources of uncertainty and point the way to reducing it. Thus, we derive, say, a probability of core damage and

an assessment uncertainty about that probability. Uncertainty is usually (but not invariably) reduced due by pooling several approaches, including those anchored to PRA (as in the above example).

If a performance requirement in regulation calls for simple compliance with a probabilistic standard, as proposed for nuclear waste (Code of Federal Regulations 10 CFR part 191) only the expected or first order probability needs to be assessed. However, an additional layer of "reasonable assurance" might reasonably be required, by also specifying some acceptable level of uncertainty (second order probability).

In any case, PRA is no less subject to uncertainty. Where CUE uses assessments by safety specialists, PRA uses analyst's assumptions which certainly have no greater guarantee of realism, and indeed are chosen primarily for analytic tractability (Apostolakis, 1991). The argument that insistence on empirical documentation assures objectivity is a delusion, as witnessed by the common observation that PRAs for the same case have been known to differ by two or more orders of magnitude. This is because informed judgement is often replaced by arbitrary assumption, e.g. that this plant is representative of all plants, that the prescribed procedures will be followed or, conversely, that no recovery from an accident will occur.

*Role of CUE to aid private judgment* . An assessor can look to CUE to structure his thinking, make best use of PRA and any other data, and guide the search for new data. It can impose logical structure on how he uses his own knowledge and judgment, including previously unused considerations at the back of his mind. The old adage "Garbage-in-garbage-out" may be quite true. Doubtful inputs assure doubtful outputs. However, processing that (unavoidable) garbage properly-i.e. using the logical structure of CUE to infer output from input--can make it more usable, by ensuring that no garbage is added in the middle. You can also see what garbage is in there and whether it is causing problems (which you can't do readily with traditional PRA, even if judgmentally adjusted). CUE can also ensure that all currently available outside knowledge, including others' expertise is absorbed; and can progressively incorporate new data, by directing the search into the most productive directions (Brown, 1994).

I am not suggesting that a risk assessor should automatically substitute CUE for his own direct best judgment -- only that he consider it as an alternative perspective on risk. It will be a number of years before CUE technique is mature enough, and its practitioners skilled enough, to guarantee that any given application adequately captures everything the assessor knows or believes.

*Dangers of limiting PSA to PRA*. Most senior nuclear regulators appear to feel that PRA is not to be used (at least not yet), for bottom line risk evaluations to verify performance requirements (NUREG 1050, INSAG-3). However, though there is some controversy on this score and there are some exceptions, notably in nuclear waste management (40 CFR 191). There are normally too many elusive sources of risk, where little that can be said

comes from hard data. A risk assessment of Chernobyl that left out organizational factors, due to lack of hard data, would not have been very realistic or useful.

A danger of overlooking this incompleteness, whether temporary or permanent, is that total risk is understated (which may be a motive for some to keep them excluded). More subtly, it will encourage licensees to focus attention on those sources that are modeled, e.g. hardware rather than people. If utilities believe the regulator will hold them to stricter account on these measures, this may distort the allocation of safety resources. This is particularly dangerous if it focuses utility attention on reducing accident risk rather than its public consequence.

*Role of CUE in public enquiry.* Stephen Watson<sup>5</sup> makes a telling point. "CUE as a *private* aid to take the regulator from PRA to decision makes sense, but I don't think it can be a *public* requirement, because of the lack of objectivity, and therefore potential for disagreement in the argument." On the other hand some senior NRC staff do not exclude the possibility of using CUE as the basis for Individual Plant Evaluations (IPE), possibly alongside traditional PRA's.

The way to address the contentious aspects of CUE may be through a set and reviewable procedure for selecting and eliciting expert opinion along the lines suggested by Ralph Keeney and others, e.g. with carefully chosen panels (see NUREG-1150 section 2.7). Though not exactly objective, the reviewability of such a procedure reduces its personal nature and therefore its susceptibility to manipulation by interested parties. It thus may be more politically acceptable in controversial regulatory decision making, such as verifying compliance with performance requirements.

*Developing the field.* In the long run PRA and CUE approaches may converge, as enough empirical data is developed for most sources of risk to be addressed in PRA, and for most personal judgment to be swamped in CUE. But, as Keynes said, in the long run we are all dead. Given real problems in the here and now, each school has its uses.

However, a sure tip-off that the bridge is incomplete is if some sources of risk are not covered on the grounds that "we don't have the data". The distinctions should perhaps be kept clear, lest we be tempted to interpret PRA as CUE, before the bridge is completed.

There much to be said for continuing to develop PRA as Partial Reliability Assessment along existing lines. Its simplifying assumptions can be sharply stated and certain kinds of conclusion can be drawn from them (like prioritizing well-documented hardware sources of risk), undistracted by certain kinds of controversial judgment and it is more credible to some. There are certain significant technical developments to be made within the PRA paradigm, such as generalizing fault trees beyond binary options.

---

<sup>5</sup>Personal communication, April 1994.

CUE can distill what it needs, possibly in an aggregated form, with clear understanding of what it does and does not cover. A mixed PRA/CUE strategy, whereby *some* personal judgment is incorporated into PRAs for example in the form of selected "expert judgments" (Apostolakis, 1991) may only confuse the issue by being arbitrarily incomplete.

Moreover, the PRA enterprise has a highly developed technology, with an experienced community of specialists trained primarily in reliability engineering. It would be disruptive and unnecessary to re-orient it in directions that call for quite different expertise (e.g. subjective probability elicitation typically found among decision analysts). In time, no doubt, the exponents of PRA and CUE will gain enough mastery of each other's skills that PSA will become a unified field; much as statistics is coming to encompass classical frequentists and personal probabilists, after several decades of separate development.

For now, both can be used to support managerial safety evaluation: PRA presents a relatively researchable piece of the evaluation; CUE attempts to model the whole evaluation.

## REFERENCES

*[Some detail pending]*

Apostolakis, George. The concept of probability safety assessments of technological systems. *Science*. Dec 1990, vol 250, 1359-1364.

Apostolakis, 1993 [PSA Conference Proceedings, Clearwater]

Brown, R.V., and Ulvila, J.W. Does a reactor need a safety backfit? Case study on communicating decision and risk analysis information to managers. *Risk Analysis*, 1988, 8(2), 271-282.

Brown, R.V. Assessment uncertainty technology for making and defending risky decisions. *Journal of Behavioral Decision Making*, 1990, 3(3).

Brown, R.V. Impersonal probability as an ideal assessment based on accessible evidence: a viable construct? *Journal of Risk and Uncertainty*, 7:215-235, 1993.

Brown, R.V. Managing research to establish acceptable risk: Probabilistic goal allocation with application to nuclear waste siting. Submitted to *Reliability Engineering and Systems Safety*. 1994.

. Brown, R.V. and Lindley, D.V. Plural analysis: multiple approaches to quantitative research, *Theory and Decision*, 1986 20, 133-154.

Brown, R.V. and Pratt, J.W. Normative validity of graphical aids for designing and using estimation studies. Zeckhauser R, Keeney RL, Sebenius J (Eds.). In *Wise Choices* Symposium in honor of Howard Raiffa. Wiley. 1996.

Kaplan, S., and Garrick, B.J. On the quantitative definition of risk, *Risk Analysis*, 1981, 1(1), 11-27.

[INSAG-3?]

Lindley, D.V. The 1988 Wald memorial lectures: The present position in Bayesian statistics. *Statistical Science*, 1990, 5(1), 44-89.

Murley, T.E. 1 Perspective on reactor safety. Lecture to MIT Safety Course, July 1994.

NUREG-0880 *Safety goals for nuclear power plants: A discussion paper*. Washington, DC: U.S. Nuclear Regulatory Commission, February 1982.

[NUREG-1050]

[NUREG-1150]

Pratt, J.W., Raiffa, H., and Schlaifer, R. *Introduction to statistical decision theory*. MIT Press. 1995.

Raiffa, H. and Schlaifer, R. *Applied statistical decision theory*. Boston, MA: Harvard University Graduate School of Business Administration, Division of Research. 1962

Ulvila JW and Thompson B. A computerized risk analysis aid. *Comput. Operations Research*. vol 15, 2, 123-134. 1988.

Code of Federal Regulations (1992) *Title 10*, Part 50, Government Printing Office.

Code of Federal Regulations (1985) *Title 40*, Part 191, Government Printing Office.

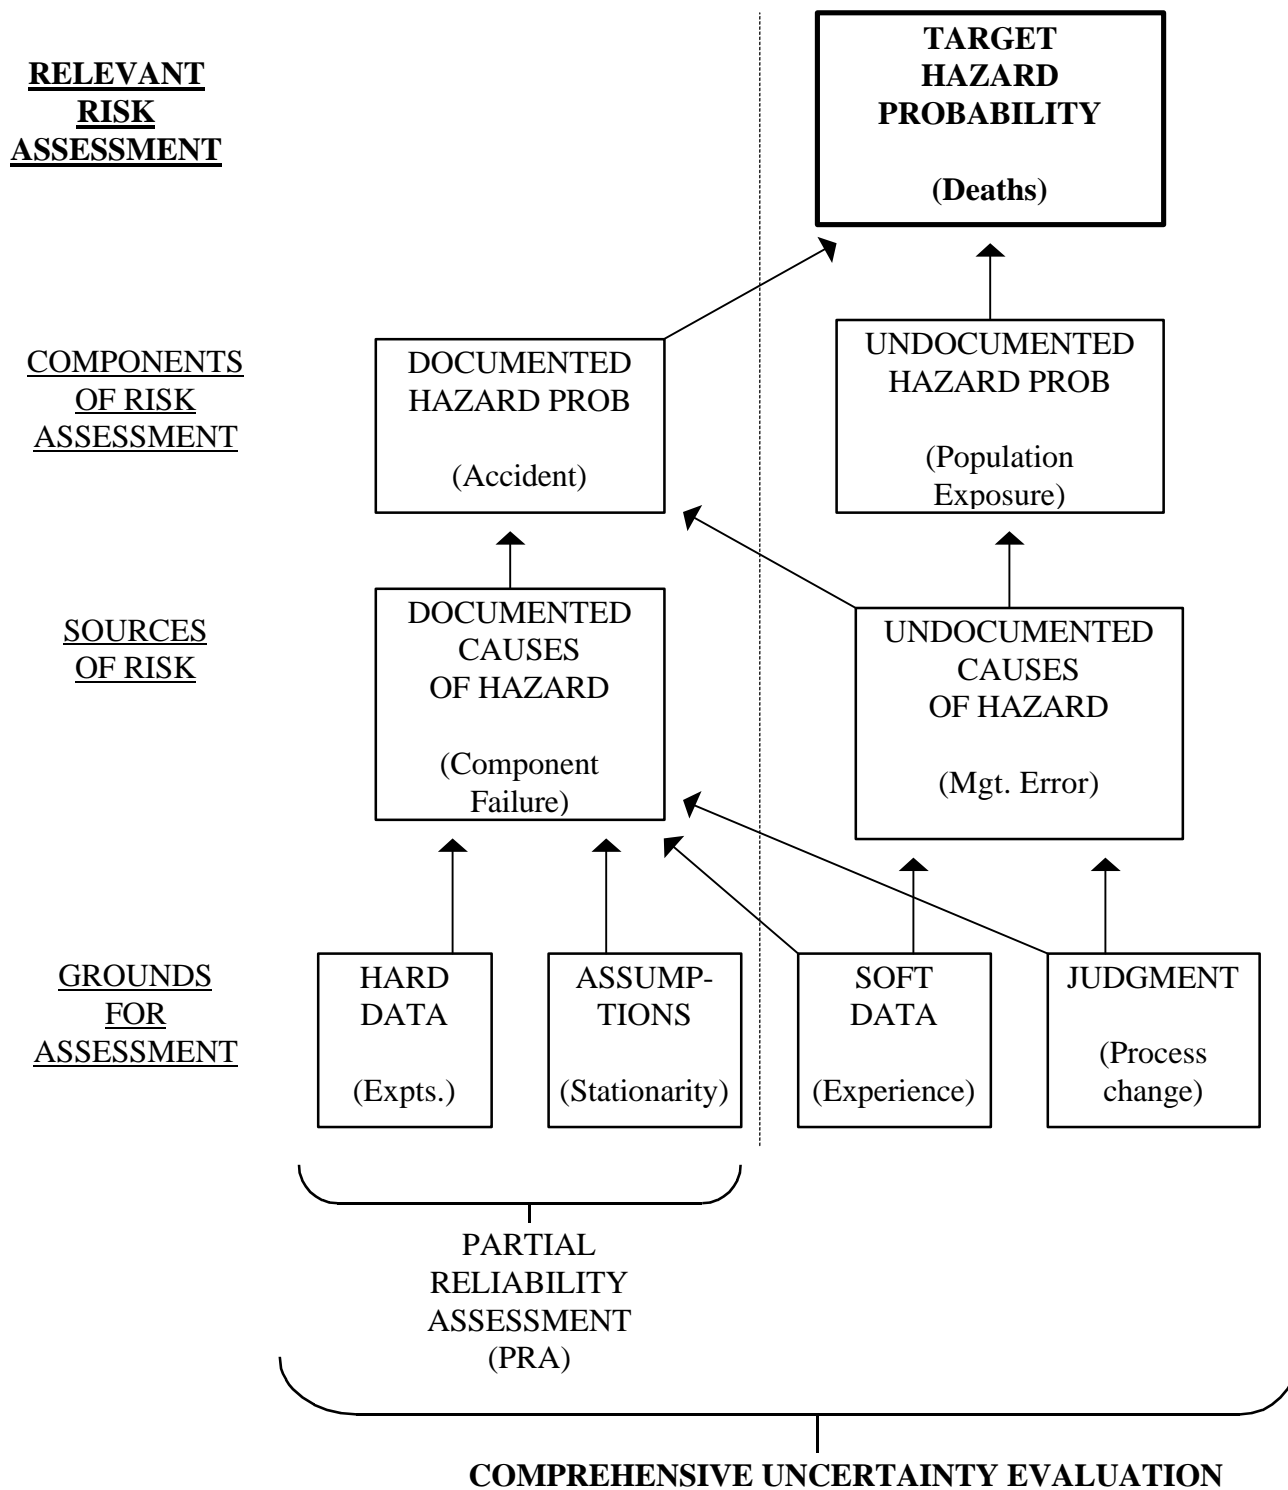
#### About the author

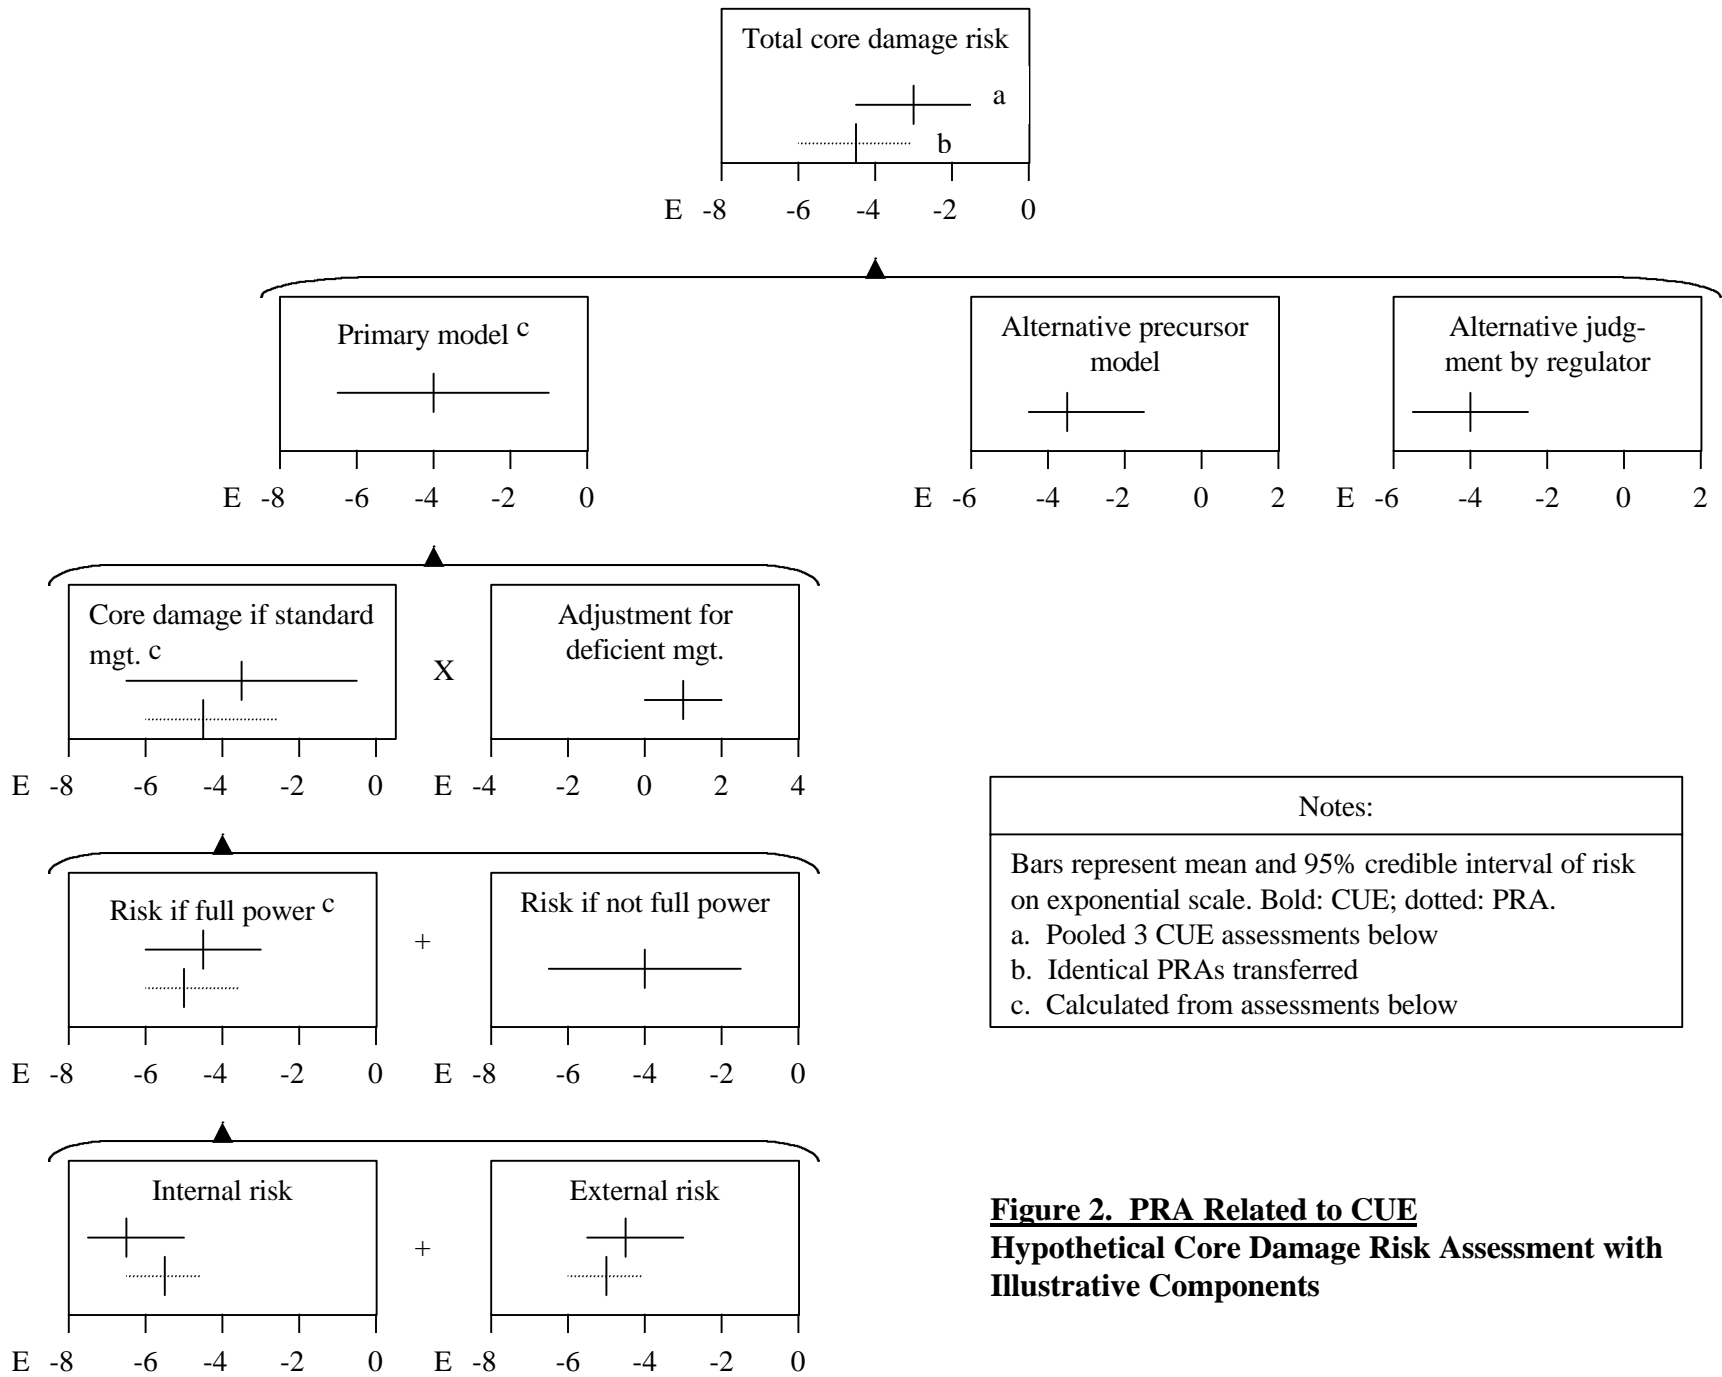
Rex Brown is currently Research Professor in the Institute of Public Policy at George Mason University and Visiting Professor at the London School of Economics. He was previously a consultant on decision processes and risk analysis, and served as advisor to the Head of Nuclear Reactor Regulation at NRC and to the Head of the Office of Civilian Radioactive Waste Management at DOE.

Contact: rbrown@mason1.gmu.edu; fax: 703-620-2839; phone: 703-620-2443

Figure 1

CONTRIBUTORS TO COMPREHENSIVE UNCERTAINTY EVALUATION  
(with examples)





**Figure 2. PRA Related to CUE**  
**Hypothetical Core Damage Risk Assessment with**  
**Illustrative Components**