Open Conceptual Questions in the Application of Uncertainty Analysis in PRA Logic Model Quantification

Sergio Guarro[a]
[a] ASCA Inc., Redondo Beach, USA

Abstract: The final stage of quantification of Probabilistic Risk Assessment (PRA) or reliability logic models is usually carried out via processes that first obtain estimates of key component-level reliability and risk parameter values, such as the failure rate in the time domain, or the probability of failure (PoF) for a specific function or mission duration, and then propagate such values from the component to the subsystem and system levels according to the component logic arrangements reflected in system reliability and failure logic models. When applying uncertainty analysis techniques to the estimation of reliability or PoF parameters of components belonging to a given system, a conceptual problem arises as to whether the same bottom up process may be applied to the definition of the prior distributions of the parameters of interest, or whether better state-of-knowledge consistency and coherence may be achieved by a top down process that proceeds from the initial construction of a system level prior distribution for the parameter of concern. This paper examines and discusses this and related issues that arise in the application of Bayesian analyses to a system PRA or reliability assessment.

Keywords: Probabilistic risk assessment, uncertainty analysis, Bayesian prior probability.

---

1. INTRODUCTION

The final stage of quantification of Probabilistic Risk Assessment (PRA) or reliability logic models is typically carried out via processes that seek to first obtain estimates of key component reliability and risk parameter values, such as the component failure rate in the time domain, or the probability of failure (PoF) for a specific function or mission duration, and then propagate such values from the component to the subsystem and system levels according to the arrangement of the component logic functions as represented in the combination of logic models used to represent the system reliability or risk, which in the PRA domain most commonly consists of a combination of binary event-trees and fault-trees.

A common estimation process for component failure rates and PoFs applies Bayesian techniques, by which a *prior distribution* of the parameter value is first constructed based on what may be referred to as "soft knowledge," e.g., generic handbook indications and/or engineering judgment concerning where the range of the parameter may lie; then the best data related to the parameter is formally brought into the process via the definition of a Bayesian *likelihood function*; finally, in the last step of the estimation process the prior distribution and the likelihood function are combined, by application of Bayes' theorem, to yield a *posterior distribution*" of the parameter, from which any desired statistics, such as mean or median value, standard deviation, percentiles and *credible intervals* – the Bayesian version of confidence intervals –, can also be extracted.

While the application of this process for a single item or component is relatively straightforward, some challenging questions arise when the process is applied to a number of components that are part of a complex system and contribute to the overall reliability or failure characteristics of the system according to both functional logic arrangement and individual reliability features. Within a probabilistic framework these questions can be viewed as primarily concerning the relation between probability assessments relative to the individual components of a system or subsystem on one hand, and the system or subsystem itself on the other. However the same questions may also be addressed from the broader perspective of what forms of uncertainty representation are appropriate under the conditions of concern. This paper examines these questions and perspectives specifically in relation to the objective of establishing and maintaining the self-consistency of a PRA or reliability model quantification framework, while using knowledge and data that may be applied at different levels of system indenture, i.e., depending on user preference, at different levels in the system functional hierarchy.

2. PRA RELEVANCE OF ISSUES

The issue of consistency between Bayesian assessments for reliability or probability of failure (PoF) conducted at different levels of system indenture was initially discussed in papers published in the 90's and initially thought of as an anomaly of Bayesian probability estimation [1, 2]. The observation generating these discussion was that the application of the same evidence, if alternatively carried out at basic component or whole system level, seemed to produce substantially different results. This subject was also discussed in more general terms in [3] as an issue of "perfect (or imperfect) aggregation" of Bayesian estimates.

Although the initial discussion of the aggregation or Bayesian anomaly issues dates several years back in time, it has gained new relevance in the context of PRAs executed in more recent times for launch and space vehicle applications. The initial PRA applications of the 80's and 90's were primarily for nuclear power plants, for which the bulk of the reliability data resided at the basic component levels, so that it was natural in this context to apply Bayesian estimations from the bottom up, i.e., by first constructing component-level PoF or reliability prior distributions using generic data, and then applying plant-specific evidence at the component and/or higher level, while progressing in the quantification process up the logic structure of a reliability or failure model, such as a fault tree. With many space systems, on the other hand, and more so in the particular case of launch vehicles, it is generally easier to construct system or subsystem-level priors than basic component priors, because generic knowledge of the range of reliability of any such a system is much better and more broadly established than corresponding knowledge for each of the basic components of which the system itself is composed. For example, it is relatively easy and defensible to identify the interval {0.01, 0.10} as a reasonable range for the per mission PoF of a typical medium to heavy class U.S. Government certified launch vehicle. It is not as straightforward to identify any "prior" range of PoF or reliability for a check valve, electronic board, or other low level component of such a system.

3. UNCERTAINTY REPRESENTATION VIA PROBABILITY DISTRIBUTIONS VERSUS UNCERTAINTY INTERVALS

A topic discussed in the literature that is related to what is being discussed here concerns whether "epistemic" uncertainty, i.e. the uncertainty that exists because of lack of knowledge, can properly be represented via probability distributions, as quite commonly is done in PRA. Ref. [4] argues against this, and uses several examples to illustrate the issue. Among these is the example of a quantity AB which is the product of a component A, known to have a value somewhere between 0.2 and 0.4, and a component B, known to have a value somewhere between 0.3 and 0.5, with no additional information available as to where the "true value" of either component may lie. Ref. [4] observes that all that can be said about this situation is that the value of AB must then be somewhere between 0.06 and 0.2. It also argues that equating the lack of knowledge concerning where the values of A and B may be with the probabilistic assumption of a uniform distribution in the respective intervals is beyond the real knowledge about those quantities. Ref. [4] observes in fact that such an assumptions would lead to the probabilistic computation of a distribution for AB that is strongly peaked at an intermediate value between the two extremes, whereas the actual knowledge about AB cannot really go beyond the original assertion that it is somewhere in the interval {0.06, 0.2}.

Translating the above into a PRA context of assessing and propagating uncertainty on probability of failure (or failure rate) values, we can see the analogy with a situation in which the quantities A and B are uncertain PoF values, about which the only knowledge is that they lie in the stated intervals, for two components that represent redundant functionality for a system or subsystem. Under these conditions, the PoF of that system or subsystem is given by the product AB, and the considerations made by [4] about what may be valid or invalid to determine about the AB value apply to this analogous situation.

An additional observations about the situation discussed in [4] can be made if we consider the practice, long established in the probabilistic arena and known as Laplace's Principle of Insufficient Reason, of assuming a uniform probability distribution in a given range when nothing is known about a random variable besides the fact that it must fall somewhere within that range. As shown below, this implies a contradiction between interval representation of uncertainty in the terms argued by [4] and a Bayesian probabilistic representation of the same uncertainty.

The assertion by [4], based on interval analysis, is that all can be said about A, B and AB is that they are somewhere in the respective intervals {0.2, 0.4}, {0.3, 0.5} and {0.06, 0.2}, as represented by Figure 1. Translated in probabilistic terms by application of the above mentioned Principle of Insufficient Reason, this would lead to a representation of all three parameters as random variables uniformly distributed in the corresponding ranges, as shown in Figure 2. However, once probability distributions are defined to represent the A and B factors of a product AB, the axioms and laws of probability, together with the mathematical formulation linking the three variables, univocally define a specific form of the probability distribution of AB itself. In the case where A and B are completely independent variables uniformly distributed in the intervals specified above, the distribution of their product AB looks like what is shown in Figure 3. In summary, within a probabilistic framework, it appears contradictory to claim complete lack of knowledge of where the value of a dependent variable like AB may lie, if the same has been claimed for the values of the associate independent variables, like A and B, and vice versa.

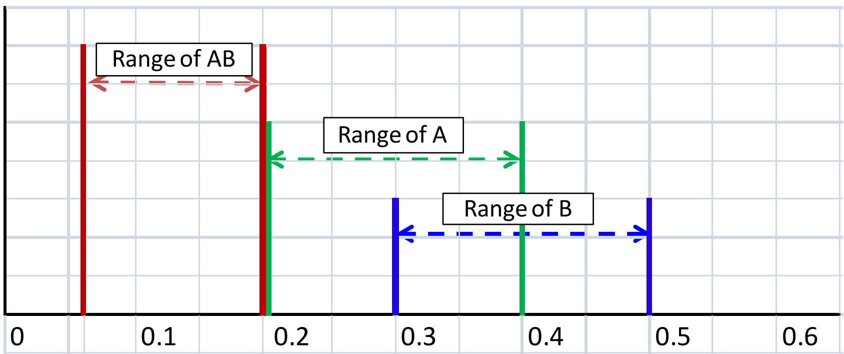

Figure 1: Interval Representation of Variables A, B, and AB

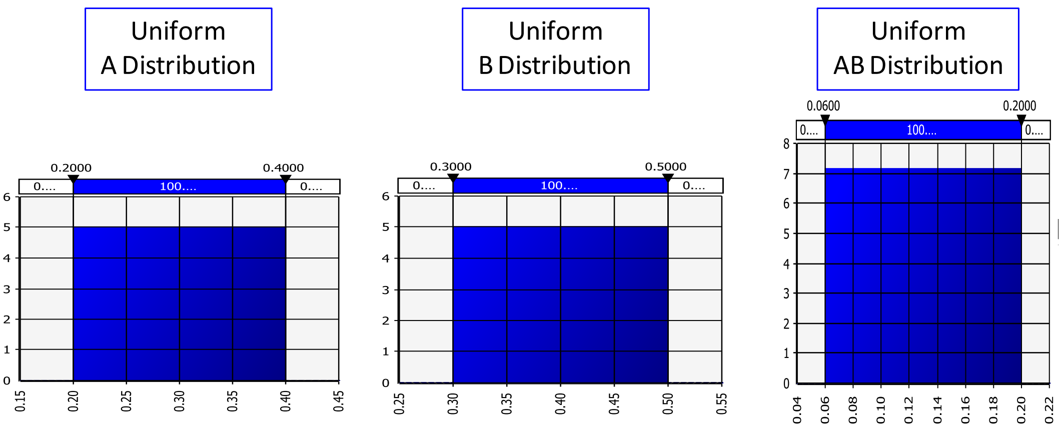

Figure 2: Principle of Insufficient Reason Probabilistic Representation of Variables A, B, AB

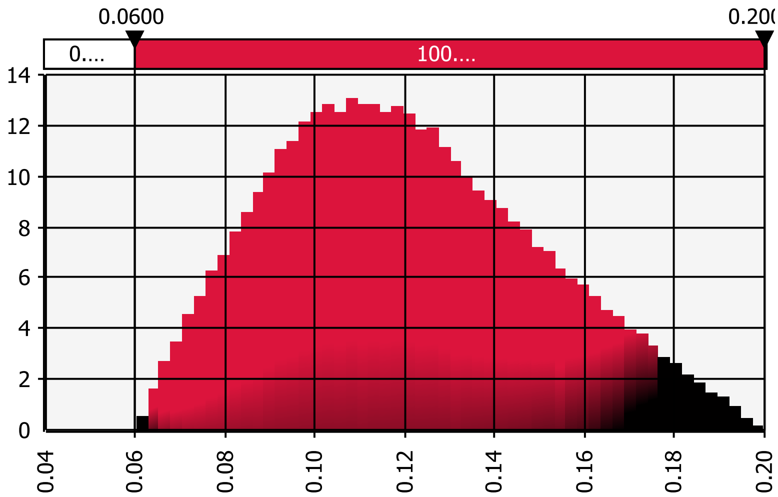

Figure 3: Probability Distribution of AB Derived as Product of Uniform A and B Distributions

For completeness, we observe that a way by which a uniform distribution can be obtained for the product AB when uniform distributions have separately been assumed for A and B does exist. This is if the value taken by one of the two terms, let us say B, is completely determined by the value assumed by the other in a certain specific fashion. For our example, this specific fashion would be that of a deterministic relation between A and B, such that, given that a random value of A occurs at a location corresponding to a fraction F of the A interval, then B is forced to a value that makes the product AB fall in a location corresponding to the same fraction F of the AB interval. This de-facto reduces the situation to one where there is only one independent variable A, and the product AB is linked to it by a relation of proportionality. This extremely correlated condition is certainly not representative of the most common uncertainty assessment conditions that PRA analysts are faced with, thus the basic terms of the issue being discussed are not affected by this observation.

3.1. Relation to the Bayesian Aggregation Issue

Even if it may not be immediately apparent, a relation exist between the issue discussed in [4] and the "Bayesian aggregation anomaly" issue.

Refs. [1] and [2] presented the anomaly subject by considering a situation where two components, which herein we shall refer to as C1 and C2, are arranged functionally in series for the successful operation of a system S. This logic arrangement happens to be in reliability terms the functional opposite of the one discussed above as a reliability analogy of the AB product example discussed in [4], but this is purely coincidental and unimportant for the purposes of the present discussion, as shall be evident from the following. For the stated situation the cited references assume a prior knowledge of the PoF values for the components C1 and C2, and for the system S. Starting from this premise, a Bayesian assessment is applied with the benefit of hypothetical binomial evidence resulting from the system S being tested a number of times, with one observed system failure occurring because component C1, but not C2, fails. Under these conditions, a Bayesian assessment is carried in two different ways, which Refs. [1] and [2] consider to be equivalent:

- In the first mode of assessment the test evidence is applied separately to the two component priors, and then the updated probabilities of failure (PoFs) are combined to provide the system PoF.
- In the second mode of assessment the test evidence is applied directly to the system prior to yield the system PoF.

Refs. [1] and [2] observe that the posterior results for PoF(S) turn out to be different in the two versions of the Bayesian assessment and consider this to constitute a probabilistic anomaly in light of the fact that the same test evidence is applied in the two cases, which would lead to the expectation of same ultimate results. In reality, as discussed in [5], the "anomaly" can be explained by the fact that the two assessments start out from assumptions of prior distributions that are not equivalent to each other; thus, although the same test evidence is applied to the two situations, the non equivalence of the assumed priors makes it so that the two overall Bayesian processes considered in are not actually equivalent and cannot produce the same results.

To better illustrate how the anomaly may occur, consider a modified version of the example initially used in Refs. [1] and [2], in which, for the same system configuration, no prior knowledge of the component and system PoFs is assumed to be available, and the test evidence is as summarized in Table I. The main modification of the example with respect to [1] and [2] consists of assuming a lower number of system tests (10 instead of 100), which has the effect of making certain aspects of the issue being discussed even more evident than when using the original case.

Table I: Test Results

	Component C1	**Component C2**	**System S**
Prior PoF	"Unknown"	"Unknown"	"Unknown"
Number of tests	10	10	10
Number of failures	1	0	1

Repeating the steps used in Refs.[1] and [2], in one version of Bayesian assessment the following Process 1 is applied:

A. Reflecting the assumed complete lack of knowledge of the PoF values before the test, uniform priors are assumed for C1 and C2 in the theoretical PoF interval [0, 1];
B. the test evidence of 1 failure in 10 trials for C1, and 0 failures in 10 trials for C2 is applied to obtain posterior PoF distributions;
C. the S system PoF is obtained from the OR-gate formula:

$$PoF(S) = 1 - [1 - PoF(C1)]\,[1 - PoF(C2)] \qquad (1)$$

In a second version of the assessment the evidence is not applied at the component level, but directly at the system level, i.e. according to Process 2:

A. Reflecting the assumed complete lack of knowledge of the PoF value for S before the test, a uniform PoF prior is assumed for the whole system;
B. the evidence of 1 system failure in 10 trials is applied to obtain the posterior PoF for the system S.

Figure 4 shows that indeed the two processes produce different results for the posterior System PoF. However, the reason for the difference lies in the fact that the setting of priors in the two versions of the assessment is not equivalent, which in turn makes the two above processes to be themselves not altogether equivalent. In Bayesian "state-of-knowledge" terms this seems to imply that for a case like this an assessor cannot claim complete lack of knowledge of the PoF whereabouts at the same time at both the system and component levels.

To verify what was just said above about the cause of the difference in results between Process 1 and 2, a Process 3 can be applied to modify Process 2 in a manner that makes the use of the test evidence at the whole system level consistent with the Process 1 utilization of the same evidence at the component level. Process 3 would be applied as follows:

A. Reflecting the assumed lack of knowledge of the PoF values before the test, uniform priors are assumed for C1 and C2 in the theoretical PoF interval [0, 1];
B. a system prior PoF is obtained from the OR-gate formula, eqn.(1);
C. the evidence of 1 system failure in 10 trials is applied to obtain the posterior PoF for the system S.

The results of Process 3 do coincide at the S system level with those of Process 1, i.e., the system S posterior PoF produced coincides with the burgundy-shaded curve of Figure 4, as is actually shown by the direct comparison in Figure 5. It is noted in this regard that the small difference visible in Figure 5 between the Process 1 and Process 3 System POF distribution results is actually due to a distribution-fit approximation introduced to simplify the Bayesian updating process computations.

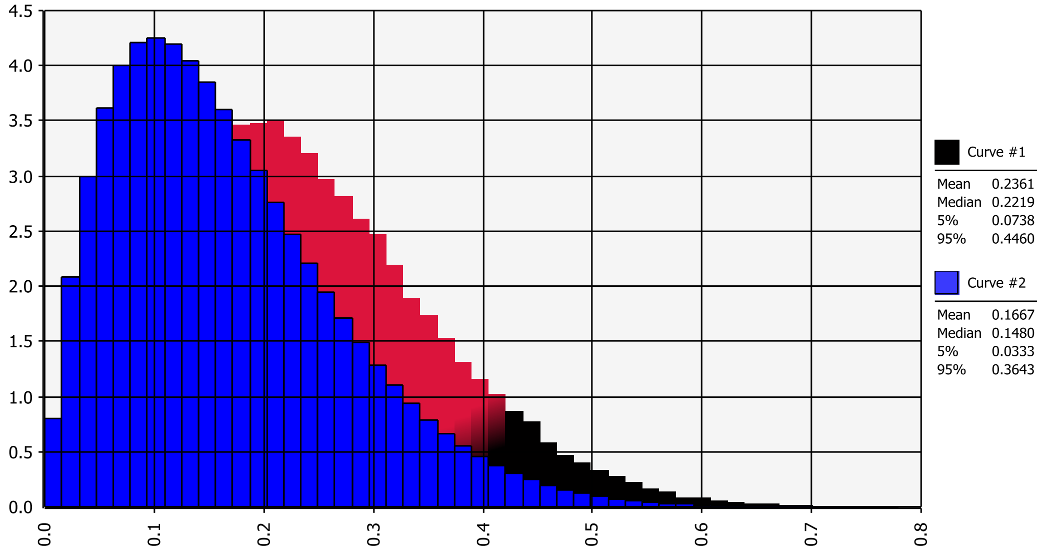

Figure 4: Posterior Probability Distributions for System S PoF per Processes 1 and 2

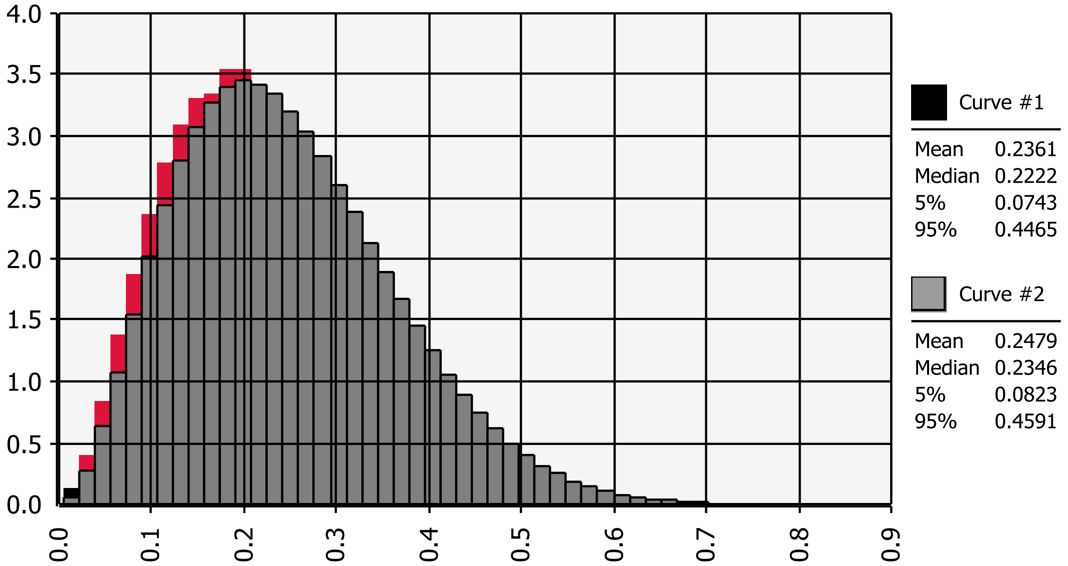

Figure 5: Posterior Probability Distributions for System S PoF per Processes 1 and 3

It is also noted that if conversely one desired, in a situation like the one in the above example, to set up a system level prior PoF and then apply component-level test evidence, to maintain the overall consistency of the estimation process it would then be necessary to decompose the S system prior PoF into component priors, in such a way as to satisfy eqn.(1). That is, the two component prior distributions, when combined by means of eqn.(1), would need to produce the initially defined prior PoF(S) distribution.

Recalling the observation made in [4] with regard to the AB product example discussed earlier in Section 3, the non-equivalence of the two sets of prior distributions assumed in Refs.[1, 2] for the two types of Bayesian assessments can be conceptually related to the difference between obtaining the AB product distribution from assumed uniform distributions of A and B, versus assuming its value to be somewhere in between its known bounds. Ref. [4] is correct in pointing out this difference, regardless of whether one also accepts the associated argument that to assume any probabilistic distribution for a parameter when only its bounds are known represents an overstepping of knowledge. In the Bayesian assessment context of the Refs. [1, 2] discussion, the part of the observation that matters is that the assumption of uniform distributions for the PoFs of C1 and C2, in their possible intervals, is not equivalent to the assumption of a uniform distribution for the PoF of S in its possible interval. This is because the PoFs in questions are tied together by the mathematical relation of eqn.(1), which in turn reflects the logic relation of success and failure between the components C1 and C2 and the system S. In any probabilistic context the eqn.(1) relation is a constraint that must be obeyed at any stage of assessment. Thus in a Bayesian assessment it is a constraint that must be obeyed by both the prior and posterior PoF distributions relative to C1, C2, and S.

In the context of the Ref.[4] discussion, a constraint conceptually corresponding to that expressed by eqn.(1) is represented by the relation between A, B, and the product AB. Despite the difference between the mathematical formulations of the two constraints, their effect is conceptually equivalent, in that any assumption of probability distributions for A and B in the Ref.[4] example, or the PoFs of C1 and C2 in the Refs.[1, 2] example, respectively forces the product AB and the PoF of the system S to have themselves distributions of nature and form which are determined by what has been assumed for the composing elements (A and B in the first example, PoF(C1) and PoF(C2) in the second). The converse is also true, so that if probability distributions are assumed for AB or PoF(S), any distributions of the composing terms (A and B in one case, PoF(C1) and PoF(C2) in the other) must be such that, when combined according to the mathematical formulations for AB and PoF(S), they produce the distributions assumed for the latter entities.

The bottom line of what we have discussed above is that, in the presence of a logic-mathematical relation between uncertain quantities or parameters, an assessor must be very careful not to violate the associated constraints when making assumptions concerning ranges of uncertainty and translating these into the assumption of "state of knowledge" probability distributions. This is further discussed in the following sections.

4. PRA REPRESENTATION OF UNCERTAINTY FOR COMPONENTS AND SYSTEMS

This section attempt to more directly connect the above discussion and observations to the question of how to maintain consistency in the representation of uncertainty and prior state-of-knowledge in probabilistic risk assessment (PRA) of multi-component systems. For this purpose, we will initially set aside the question of whether other ways of representing uncertainty should be considered as an alternative to treating uncertain parameters as probabilistically distributed random variables, but will come back to this question after discussing the key issue of consistency that a PRA assessor may face even in this more limited context.

4.1. Top-down versus Bottom-up Prior Knowledge Representation in PRA

The discussion in Section 3.1 has brought to attention the constrained nature of the process of constructing probability distributions in the presence of a functional and/or mathematical relation among the parameters for which such distributions are constructed.

In the earlier classical forms of PRA application the consistency issues discussed in Section 3.1 did not come to the forefront, as these applications, primarily concerning nuclear power plants, carried out probabilistic quantification entirely from the bottom up, i.e., defining first "generic prior" probability distributions for reliability parameters at the component level, and then updating these priors with more plant-specific data available at increasingly higher levels of indenture in the overall plant PRA models. In the simple conceptual example discussed in Section 3.1, this would be equivalent to defining prior PoF distributions for the two components C1 and C2, without having to worry about what definition of system S PoF prior distribution would be consistent with the priors assumed for C1 and C2.

As previously mentioned in Section 2, the situation may be different in other types of PRA applications, such as the space system PRAs which have become more common in recent years, where "prior generic knowledge" may be better assessed at the whole system level. After this is done first, mission-specific evidence is usually applied at lower levels in order to obtain output results that provide insights into subsystem or component behaviour, e.g. for the purpose of identifying contributors to which reliability and risk-reduction measures may be applied to obtain better overall system performance. Under these conditions, and specifically when Bayesian techniques are used in the probability assessments, it becomes necessary to construct individual subsystem and component prior distributions that, taken all together, maintain consistency with the overall uncertainty distribution assumed at the whole system level.

4.2. Constructing Component Priors from a System Prior

An approach to the problem posed at the end of the above section has been discussed in [5], and it can be related back to what has been presented in Section 3.1. Using the same example of a system S constituted of two components C1 and C2, let us then assume the following:

a) the assessor's initial knowledge of possible PoF values concerns in absolute terms primarily the system S;

b) PoF results at the "posterior" level are desired for practical purposes at the component C1 and C2 level;

c) test data for the two components are available to carry out assessments at the component level;

d) before conducting any component level assessment with the available data, the assessor has no other direct knowledge or information concerning the possible values of the two component PoFs.

Under the stated conditions, and more specifically because of what stated above in a), it would be reasonable for the assessor to initially start a Bayesian PoF estimation process from the definition of a system-level prior distribution for PoF(S). However, points b) and c) make also clear that results are desired at the component level and that test data to carry out a statistical estimation are available at that level. Because of this it becomes then necessary for the assessor to recast the system-level knowledge represented by the PoF(S) prior distribution into individual prior distributions for the component PoFs. Per the discussion carried out in Section 3.1, this cannot be done loosely, but in such a way that the interdependence of the three PoFs represented by eqn.(1) of Section 3.1 is satisfied. That is, whatever form of distribution is assumed for the prior PoF(S), the two component PoFs, as distribution functions, must be such that their combination in the form of eqn.(1) yields back exactly the distribution function assumed for PoF(S).

In the above we essentially have a "decomposition" problem for a distributed parameter, i.e., a distribution function (rather than a single number) must be decomposed into sub-distribution functions that obey a given mathematical constraint, more specifically eqn.(1) for the example situation considered here. Ref.[5] discusses in detail how this problem can be set up and handled in mathematical terms for the same example situation. It must be noted that, unless additional information is available, the solution to the decomposition problem is not unique when only one constraint applies, as represented by eqn.(1) or by any other single equation dictated by a different reliability logic linking component PoFs to system PoF. Thus, for example, in the solution discussed in [5] an additional constraint is introduced by assuming that the two component PoFs have the same mean values. In the case of real-life systems with multiple components arranged in series / parallel logic combinations, the decomposition problem may become correspondingly more complex and challenging, yet it remains conceptually the same as that discussed here and in [5] as an example. Solution processes for these more challenging conditions have been carried out and documented for large scale launch vehicle PRA applications concerning planetary NASA missions [6, 7].

4.3. Practical Consideration in Uncertainty Analysis for Space System PRAs

A question that deserves consideration, given the mathematical challenges posed by the decomposition of a system-level PoF distribution into lower level priors for Bayesian estimation purposes, is the why such a process may be given preference over the more traditional one of setting priors up from the lowest level at which any estimation is carried out.

In the execution of a space launch vehicle PRA, the top-down choice of process for developing prior distributions is primarily suggested by the lack of generic component-level statistical data that reflect the high stress conditions (vibration, acoustic, pyro-shock, thermal) of a space launch. The approach suggested by existing reliability data compilations or formulations (such as MIL-HDBK-217, the more recent "217-Plus, and NPRD-95) is to apply multiplier factors to failure rate and PoF data compiled for components operating in conditions without the stress factors mentioned above. The lack of validation for these factors and the large discrepancies between the factors suggested by different sources for the various classes of components that are of interest, however, give reason to doubt the validity and accuracy of the outcomes of this approach. When applied in the past, in fact, approaches following this route typically resulted in launch vehicle reliability predictions that were unrealistically high. I.e., when compared with the launch vehicle system record at the end of a sufficiently high number of actual launches, the PoF initially predicted by following such a kind of approach would be systematically underestimated by as much as an order of magnitude or more.

Considering the above, some launch vehicle PRA practitioners have concluded that it would be more reasonable to start the prior distribution setup process from the assessment of ranges of reliability for entire launch vehicle systems, which, being based on generic but well established knowledge of actual launch mission outcomes, would be more defensible than processes based on data compilations of dubious applicability to space launch conditions. From the launch vehicle system top level one would then proceed, using considerations of relative ratios of failure rates or PoFs between different launch vehicle subsystems and major component, to the definition of prior distributions at progressively lower levels of indenture. The considerations and discussion carried out above in Sections 3.1, 4.1 and 4.2 are directly relevant to this alternative process of defining Bayesian priors.

5. ALTERNATIVE REPRESENTATION OF EPISTEMIC UNCERTAINTY

Much discussion can be found in the literature on the nature and proper treatment of epistemic and aleatory uncertainty. Refs. [4, 8-10] provide a sample of these discussions, and [10] offers an exhaustive bibliography on the subject. Epistemic uncertainty is commonly referred to as arising from lack of knowledge, while aleatory uncertainty is considered as originating from the intrinsic variability of some physical or other type of process, such as in the toss of a die. At the philosophical level it may be argued that aleatory uncertainty is also the result of lack of knowledge, as after all if one could exactly know and represent the impulse imparted to the die and all the boundary conditions for the bounces that follow, one may in theory be able to predict the outcome of each toss. However, the distinction is well established and accepted in the PRA community. In the practical terms of interest in the PRA context, epistemic uncertainty, unlike aleatory, typically refers to situations where the unknown parameter or variable is not really changing from case to case in the theoretical set of cases of interest for a specific assessment. Rather, it may have a very specific constant value, but that value is unknown within a certain range of uncertainty.

Based on the arguments that we have tried to summarize in this paper, some authors [4, 9] have argued that the probabilistic representation of epistemic, i.e. "lack of knowledge", uncertainty is inappropriate and leads to under-estimation of the true range of uncertainty when probabilistically propagated up the logic structures of hierarchically arranged logic-probabilistic models, such as those routinely used in PRA. To counter this, these authors have proposed alternative methods of epistemic uncertainty representations and propagation within a system model, such as "interval analysis" [4] and Dempster-Schafer evidence theory [9, 10]. Both Refs. [4] and [9] present examples of propagation and presentation of both aleatory and epistemic uncertainty within a given assessment, using separate frameworks and techniques for each.

Aside from other types of considerations that have been presented in the literature debate on the probabilistic versus non probabilistic representation of uncertainty, the one that appears to perhaps have greater relevance in a practical sense is the distinction between:

a) uncertain quantities with true full wide variability within a given range;
b) uncertain quantities which are believed to have an unknown fixed value (or unknown narrow range of variability) within a wide "ignorance range."

The above could be used as a working level criterion for the classification of aleatory versus epistemic uncertainty. Regardless of whether this is 100% valid at the conceptual and philosophical level, if the above distinction is al least viewed as a valid means of categorizing the large majority of uncertain parameters that appear in a typical risk assessment, then it can be argued that the application of probabilistic mathematical models and rules to represent uncertainty relative to quantities of type a) seems to be both justified and defensible. Equally justified, however, may be some of the doubts and questions cast against the use of the same approach to represent, and carry through a complex logic-mathematical model, the uncertainty relative to quantities of type b).

To see why the difference between the two situations depicted above may possibly affect the legitimacy of a purely probabilistic view for both cases, we may in fact visualize a theoretical experiment, in which it were possible to repeatedly sample variables of the two types. In such an experiment the sampled values of the type a) variables would indeed be varying across the respective wide variability ranges according to same specific form of distribution. Each of the type b) variables, however, would have essentially the same value appear over and over in each of the sampling tests, although this essentially fixed value would be somewhere within the initially identified wide range of possible values. The implication is that any dependent variable produced by the combination of type a) independent variables would indeed be distributed according to the probabilistic-combinatorial laws governing the independent variables, but any dependent variable produced by the combination of type b) independent variables would instead simply assume a fixed value, somewhere within a range that could be pre-identified by means of interval-analysis from the initially identified ranges of the independent variables. Essentially based on this observation, Ref.[4] proposes a means of combining probabilistic (for aleatory) and interval (for epistemic) representation of uncertainty. Unfortunately such a method appears to be practically viable only for relatively simple logic-probabilistic system models and not scalable to the very complex PRA models that are common in today's practical contexts.

6. CLOSING OBSERVATIONS AND COMMENTS

This paper has discussed a set of issues relative to the representation and handling of uncertainty within the practical objectives of a probabilistic risk assessment. Some of these issues have been previously been discussed in the literature, but further discussion in the PRA technical community of their conceptual and practical relation to the overall problem of effective representation of uncertainty in PRA may be beneficial.

At the present time the use of Bayesian-style state-of-knowledge probability distributions is the preferred means of uncertainty representation and handling in large scale PRAs. Aside from any conceptual preferences in regard, this type of representation offers without doubt practical advantages, as the computational framework to support it and executed it is well established within the software tools at the disposal of the PRA technical community. The discussion in this paper shows, however, that issues of consistency may arise in the setting of Bayesian prior uncertainty distributions, depending on whether the assessor seeks to a bottom-up or top-down representation of his/her prior state-of-knowledge concerning the elements of the system being analyzed. It also shows that the concerns that have been raised by some experts with regard to the use of Bayesian probability for the representation of epistemic uncertainty may still deserve further attention and discussion, as neither probabilistic formulations nor the alternative uncertainty representations proposed in the academic literature of the more recent years appear yet to provide a widely applicable and practically implementable solution to the related issues.

7. REFERENCES

[1] L. Philipson, *"Anomalies in Bayesian launch range safety analysis,"* Reliability Engineering and System Safety, vol. 49, pp 355-357, (1995).

[2] L. Philipson, *"The Failure of Bayes System Reliability Inference Based on Data with Multi-level Applicability,"* IEEE Transactions on Reliability, vol. 45(1), pp. 66-69, (1996).

[3] V. Bier, *"On the concept of perfect aggregation in Bayesian estimation,"* Reliability Engineering and System Safety, vol. 46, pp. 271-281, (1994).

[4] S. Ferson and L. Ginzburg, *"Different methods are needed to propagate ignorance and variability,"* Reliability Engineering and System Safety, vol. 54, pp. 133-144, (1996).

[5] S. Guarro and M. Yau, *"On the Nature and Practical Handling of the Bayesian Aggregation Anomaly,"* Reliability Engineering and System Safety, vol. 94 (6), pp. 1050-1056, (2009).

[6] New Horizons Mission, *"Atlas V 551 Final SAR Databook, Revision A,"* (May 2005).

[7] Mars Science Laboratory Mission, *"Safety Analysis Report Databook, Revision B,"* (February 2010).

[8] A. DerKiureghian and O. Ditlevsen, *" Aleatory or epistemic? Does it matter?,"* Special Workshop on Risk Acceptance and Risk Communication, Stanford University, (March 26-27, 2007).

[9] L. Swiler, T. Paez and R. Mayes, *" Epistemic Uncertainty Quantification Tutorial,"* Proceedings of the IMAC-XXVII, Society for Experimental Mechanics Inc., Orlando, Florida USA, (February 9-12, 2009).

[10] E. Zio and N. Pedroni, *"Literature review of methods for representing uncertainty,"* Cahiers de la Sécurité Industrielle, vol. 2013-03, Fondation pour une Culture de Sécurité Industrielle, Toulouse, France (April 2013) – (Available at http://www.FonCSI.org/fr/).