Understanding Relative Risk: An Analysis of Uncertainty and Time at Risk

A. El-Shanawany[a,b]
[a] Imperial College London, London, United Kingdom
[b] Corporate Risk Associates, London, United Kingdom
ashanawany@c-risk-a.co.uk

Abstract: Risk at nuclear facilities in the UK is managed through a combination of the ALARP principle (As Low As Reasonably Practicable), and numerical targets. The baseline risk of a plant is calculated through the use of Probabilistic Safety Assessment (PSA) models, which are also used to estimate the risk in various plant states, including maintenance states. Taking safety equipment out of service for maintenance yields a temporary increase in risk. Software tools such as RiskWatcher can be used to monitor the real time level of risk at plant. In combination with software tools to estimate the instantaneous risk, time at risk arguments are frequently employed to justify safety during plant modifications or maintenance activities. In this paper we consider the effect of using conservative estimates for the probability of failure on demand of safety critical components compared to using a full uncertainty distribution. It is found that conservatism in the base case model translates to a hidden optimism when used in time at risk arguments. While it is known and accepted that quantified risks are necessarily approximate, useful insights can be gained through risk modelling by considering relative risks. Anything that distorts relative risks impacts on the usefulness of the risk modelling. The important point of the effect discussed here is that it has the potential to distort relative risks. The mapping between the base case conservatism and the time at risk optimism is characterised, and the effect is illustrated using simple hypothetical examples. These simple examples show that the shape of the full uncertainty distributions of model parameters have important and direct consequences for time at risk arguments, and must be considered in order to avoid distorting the risk profile.

Keywords: PRA, Uncertainty, Bayesian, Risk, Modelling.

1. INTRODUCTION

Conservative estimates are a mainstay feature of probabilistic risk models used for safety analysis. Conservative arguments are frequently invoked, often in cases when it would be hard to confidently provide an accurate estimate, and are strongly linked with the assessment of uncertainty. The use of conservative arguments implicitly restricts the value of quantitative risk analysis (QRA) to statements such as “the frequency of core melt is lower than x per year”. This fails to do justice to the potential uses of QRA. The value of QRA can extend beyond the identification of “negative insights” and high risk areas, to informing plant operators about “positive insights” such as where the safety margin is very high and could potentially be relaxed [1]. The issue of uncertainty in risk analysis has been discussed extensively in the literature, and the importance of an adequate representation of uncertainty has also been presented [2, 3]. This paper stresses the point that conservative arguments are not an adequate approach to uncertainty by demonstrating that conservative estimates distort the risk profile of the plant, sometimes in non-obvious ways. This lends extra weight to the viewpoint that conservative estimates should always be replaced with best estimates coupled with uncertainty estimates. The conservative distortion is illustrated using the concepts of time at risk and maintenance, in which case conservatisms can hide the true risk.

Time at risk is a fundamental concept when considering risk. In most quantitative risk models, time at risk is used to represent the effect of maintenance, the degradation of plant components, and their susceptibility to various hazards. This paper will explore the concept of time at risk and uncertainty in parameter estimates using the example of maintenance states and considering how plant unavailability due to maintenance affects the prediction of risk.

Maintenance is known to have both potentially positive and negative effects on the risk at a plant. Maintenance is used to identify and fix defects that occur due to anticipated wear and tear on plant equipment due to normal operation, and is vital to ensure the continued operation of the plant. However, incorrectly performed maintenance can also leave a plant in a worse state than before; the classic example is that of misaligned valves. In addition, during a maintenance outage safety equipment is unavailable to perform its duty, hence the plant incurs an increase in risk during the maintenance period. Given that maintenance can have a significant impact on the risk at a plant, it is important to be able to estimate all aspects of maintenance risk as precisely as possible in order to design maintenance schedules that are as close to risk optimal as possible. The significant effects of maintenance on risk at industrial facilities have been extensively discussed in the literature [4, 5, 6]. Although there are numerous dimensions to the interaction of maintenance with plant risk, this paper will start from the reasonable assumption that maintenance is essential and then consider only the impact due to unavailability of plant systems. It is demonstrated that the method used to estimate the failure parameters in the risk model has a significant effect on the estimation of the risk significance of maintenance outages.

Before proceeding, it is worthwhile to make more precise the key concepts used in this paper. It is noted that there are numerous formulations of the definition of risk, and that the word is often used differently in various contexts. Frequently, risk is defined as probability “multiplied” by consequence, and it is commonly expressed in terms of a frequency of an undesirable event per unit of time. This is an excellent intuitive description of risk, although it is noted that there are certain deficiencies with the definition, such as a precise definition of the multiplication operation. In this document Kaplan and Garrick’s [7] definition of risk as a triplet consisting of scenario, probability and consequence is used. However, the consequence part of the triplet will be a constant throughout this document, and will simply be considered to be “failure to perform a prescribed safety function”. The scenario will switch only between a nominal “normal state” and a “maintenance state”. Hence, of the triplet defined above, the main concern in this paper is the probability component of risk. When considering time at risk, the concept of cumulative risk is key; at its most general, cumulative risk can be defined as the time integral of instantaneous risk. The instantaneous risk will refer to the risk at some specified time point, while the cumulative risk is the total risk experienced over a period of time. In this document “risk” will be used as a synonym for “instantaneous risk”. For example, the risk due to a specified hazard depends on the length of time over which the hazard could potentially occur. Uncertainty itself is a complex topic, but in this document the phrase uncertainty is restricted only to statistical uncertainty. The expansion to the full range of uncertainties is considered in the further work section.

The cumulative risk and instantaneous risk in a maintenance case are shown in Figure 1 below.

Figure 1: Maintenance Outage and Instantaneous Risk

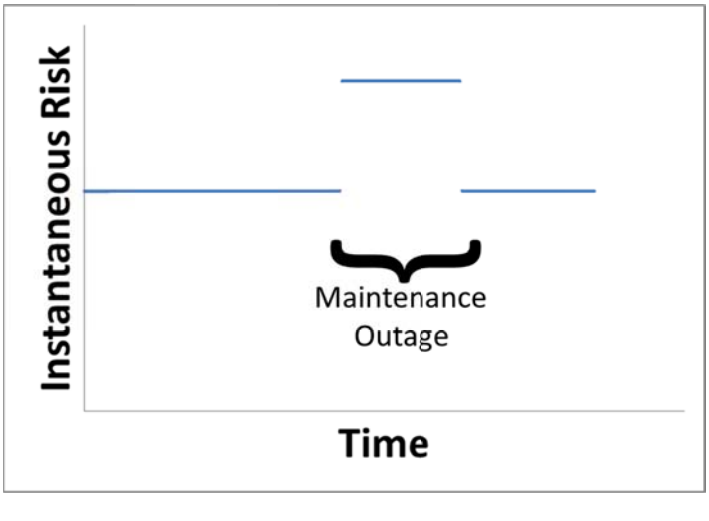

It should be noted that in Figure 1, above, a constant failure rate is assumed and the cumulative risk has been calculated using a Poisson distribution and is not actually proportional to the time; in particular the cumulative probability does not rise linearly to one. For realistic values of the Poisson parameter and short time periods the relationship is very close to being proportional to the time near the origin. In general the failure rate of a component is not constant, for example a bathtub shaped curve could be modelled using a Weibull distribution, but over the short time period that maintenance occurs, a constant failure rate is a reasonable assumption.

To consider maintenance cases in a time varying model, the variation with time of the probability of an event must be considered. A saw tooth curve is produced by the basic assumption that a component's failure rate is constant with time, and that maintenance perfectly restores a component to working order.

Figure 2: Classic Saw Tooth Maintenance Curve

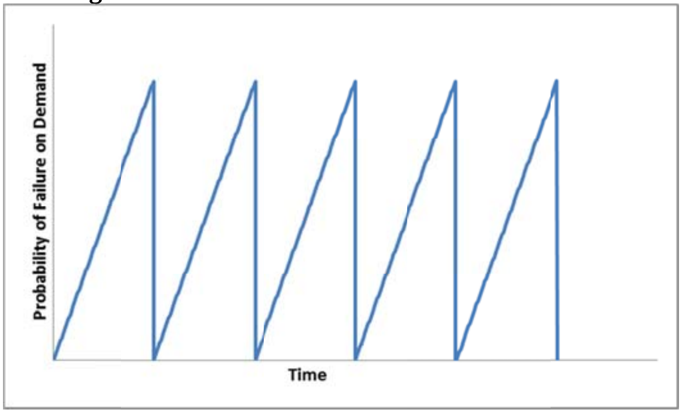

The contribution of maintenance outages can be usefully considered as a proportion of "baseline risk". It is shown that the uncertain nature of the probability component can distort the relative importance of time effects depending on the way in which the uncertainty is handled. Historically, uncertainty in probability estimates has been handled by using conservative estimates for the probability. Despite a shift in thinking towards best estimate values, in practice conservative judgement is frequently invoked in difficult situations. The purpose of making the argument presented in this paper is to reinforce the viewpoint that even where a best estimate is difficult, it should not be replaced by a conservative estimate; in the opinion of the author a large uncertainty distribution (if necessary) is preferable to recourse to a conservative estimate. It is shown that the *combination* of conservative estimates, in the sense of multiple lines of protection, yields an *optimistic* viewpoint of the relative risk during time periods in which a line of protection is removed. This is first demonstrated by comparing conservative estimates and best estimates, and then the case of using best estimates plus uncertainty for the failure parameters is considered. The best estimate plus uncertainty method yields similar results as the best estimate method, but retains the possibility to interpret the results in a conservative way.

An appropriately designed maintenance schedule is an important contributor to the safe operation of an nuclear power plant (NPP). The issue of estimating the effect of maintenance is a multi-faceted problem, and numerous models of maintenance have been developed [4, 5], but the argument made in this paper is not dependent on the particular maintenance model used. Hence, for clarity, only the most simple model assuming a constant failure rate with time, as shown in Figure 2 above is considered here. Two cases are considered, using a very basic model in which safety systems have a fixed probability of failure on demand, that is time independent. Using this model we seek to estimate the increase in the risk due to unavailability of a single system. Time at risk is the key concept in this formulation since the risk due to a particular plant state is directly proportional to the time spent in that plant state using the simplifying assumptions. The theoretical consequences of conservative estimates of the probabilities of failure on demand are considered in Section 2. This is illustrated using a simple example in Section 3. Section 4 then further develops the argument to include uncertainty distributions for the failure parameters, rather than point estimates. Section 5 discusses the implications of the results and Section 6 presents areas for further work. Section 7 presents the conclusions of the paper.

2. THEORETICAL JUSTIFICATION

This section provides a theoretical justification for why conservative estimates of failure parameters lead to the underestimation of the contribution to the total risk of maintenance outages. This is done by building up the algebra for a simple system with 'n' diverse lines of protection.

Let X be a system with n protective systems, each of which provides an independent protective barrier to the failure of system X. Note, common cause failures are not considered in this setup, as each line of protection is considered to be a diverse system and hence genuinely independent of each other. Logically, this can be represented by the fault tree shown in Figure 3 below.

Figure 3: Simple Example – A System with n Lines of Protection

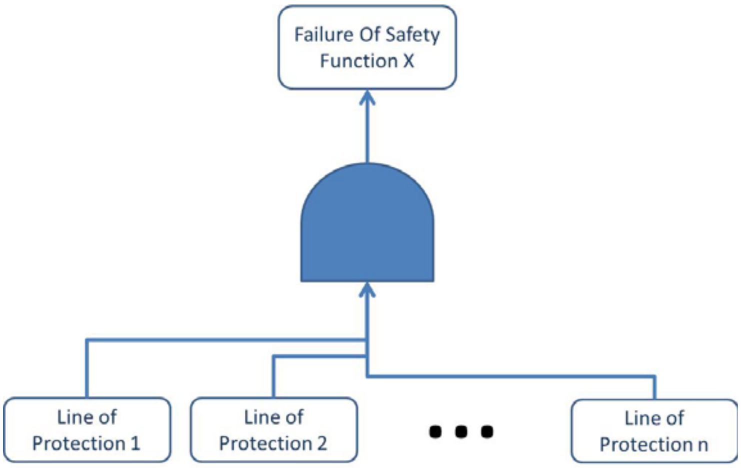

Let $P(B_i)$ be the probability of failure of the i^{th} protective barrier. In general the probability of failure of a barrier can be described by some unknown distribution D_i; i.e. $P(B_i) \sim D(\alpha, \beta)$. The estimation of the distribution D is, in general, a difficult task with numerous associated difficulties for which there is an existing and extensive literature; References 8, 9, 10 and 11 provide a good background on some of the estimation methods used in modern risk analysis and discussion of the associated difficulties. In practice the distribution D is never known, although some of the sources of variability may have been partially estimated. Historically conservative values have been used, usually attempting to estimate the 95^{th} percentile of the distribution D. The conservative estimate will be represented by $\hat{P}(B_i)_{95}$. The "best estimate" of $P(B_i)$ is some measure of central tendency. Usually in risk analysis models the mean is used. However, for purity of the results, and strictness of inequalities the median is used in this document; it is noted that the extension to the use of a mean estimate is straightforward, except for the existence of certain caveats relating to heavily skewed distributions. The best estimate median value is represented by $\hat{P}(B_i)_{50}$. It is noted that for any probability distribution that is not a single point the following strict inequality holds:

$$P(B_i)_{95} > P(B_i)_{50} \quad (1)$$

It is hence reasonable to assume that for any "good" estimate, the following strict inequality will also hold:

$$\hat{P}(B_i)_{95} > \hat{P}(B_i)_{50} \quad (2)$$

It is noted that for most distributions:

$$P(B_i)_{95} > P(B_i)_{\mu} \tag{3}$$

Hence the results presented in this document will almost always also hold if the mean of the distribution D is used in place of the median.

Let P(X) be shorthand for the probability that system X fails. Then given independence of the lines of protection we see that:

$$P(X) = \prod P(B_i) \tag{4}$$

Using the notation above we can find a best estimate of the probability of failure of X using:

$$\hat{P}(X)_{50} = \prod \hat{P}(B_i)_{50} \tag{5}$$

And a conservative estimate of the probability of failure of X using:

$$\hat{P}(X)_{95} = \prod \hat{P}(B_i)_{95} \tag{6}$$

Now, consider the effect of removing one train of protection, for example for maintenance. Without loss of generality assume that the jth line of protection is removed. Using the subscript 'mj' to denote maintenance of barrier 'j', the system failure estimates now become:

$$\hat{P}(X)_{50,mj} = \prod_{i \neq j} \hat{P}(B_i)_{50} \tag{7}$$

And:

$$\hat{P}(X)_{95,mj} = \prod_{i \neq j} \hat{P}(B_i)_{95} \tag{8}$$

Risk models for complex engineering systems are acknowledged to be approximate tools. Most analysts agree that the absolute value of probabilities calculated using the risk model are very approximate. However, ranking of risks and estimating relative magnitudes is still a useful output, even in absence of good absolute measures. The estimate of the relative risk of different plant components and configurations is a valuable output from risk models. For this reason the risks above can be usefully considered in the context of the baseline risk when all lines of protection are available. The relative risks are:

$$\frac{\hat{P}(X)_{50,mj}}{\hat{P}(X)_{50}} = \frac{1}{\hat{P}(B_j)_{50}} \tag{9}$$

And:

$$\frac{\hat{P}(X)_{95,mj}}{\hat{P}(X)_{95}} = \frac{1}{\hat{P}(B_j)_{95}} \tag{10}$$

Now, noting that the 95th percentile estimate is larger than the 50th percentile estimate we see that:

$$\frac{1}{\hat{P}(B_j)_{95}} < \frac{1}{\hat{P}(B_j)_{50}} \tag{11}$$

This equation tells us that using a 95th percentile estimate of every line of protection gives a lower estimate of the relative increase in risk during maintenance compared to the relative increase in risk that occurs if a median estimate of the probability of failure of each line of protection is used. In general, this means that conservatively estimating failure probabilities results in optimistic estimates

of the relative risk increases during maintenance. Furthermore, the level of optimism is proportional to the level of conservatism, if we define “conservatism” to mean a multiplication factor from the best estimate.
It is noted that the above demonstration did not require the introduction of time at all into the argument. The time argument remains a linear argument that affects only the magnitude of the above effect. The extension to consider a time variant model is trivial but provides the same message with more complicated algebra. The next section considers the effect of the conservatism described above on a simple example model, and shows that the using conservative values gives an under estimate of the proportion of risk that is incurred during maintenance compared to during normal operation with all plant available.

3. SIMPLE EXAMPLE

This section works through an example fault tree representation of a simple system, to demonstrate the effect described in section 2. A system with two lines of protection, instead of n lines of protection, is used for clarity. The failure logic of this simple system is shown as a fault tree in Figure 4 below:

Figure 4: Simple Example – Two Protective Barriers

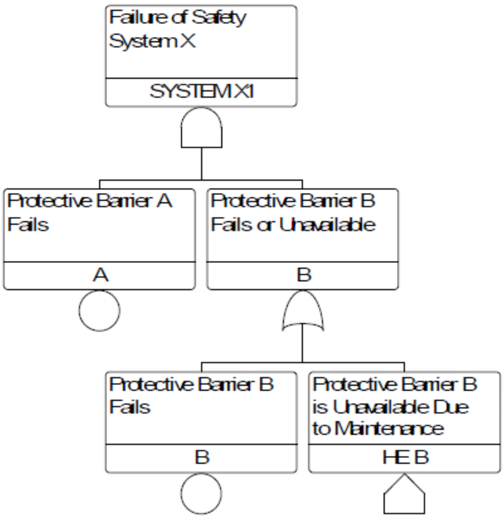

Figure 4 shows the base case for the system, in which the house event is set to false. It is assumed that each line of protection provides has a probability of failure on demand, as per the equations developed in section 2 above. A maintenance case can be considered using the same fault tree by setting the house event to true. This has the effect of taking one line of protection out of service.

A hypothetical conservative and best estimate cases are considered. The probability of failure on demand of each protective barrier in each case is shown in Table 1 below.

Table 1: Probability of Failure on Demand of Each Protective Barrier vs Estimation Technique

	P(A Fails on Demand)	P(B Fails on Demand)
Conservative	1E-02	1E-02
Best Estimate	1E-03	1E-03

Table 2 below shows the probability of failure on demand of safety system X, for each system state and for each failure probability assumption.

Table 2: Probability of Failure on Demand of Safety System X

	Base Case	Maintenance Case
Conservative	1E-04	1E-02
Best Estimate	1E-06	1E-03

The key point is demonstrated in Table 3 below, which shows the increase in the probability of failure on demand of safety system X during maintenance under conservative and best estimate assumptions.

Table 3: Ratio of the probability of failure on demand of system X to failure probability on demand in the base case.

	Ratio of Maintenance: Base case probability of failure on demand of System X
Conservative	100
Best Estimate	1,000

Let the base case failure on demand be $P(X)_{C,BC}$ and $P(X)_{B,BC}$ under conservative and best estimate assumptions respectively, and similarly let the maintenance case failure on demand be $P(X)_{C,M}$ and $P(X)_{B,M}$ under conservative and best estimate assumptions respectively. Further assume that a fixed proportion p_M of the time is spent in the maintenance state. The proportion of time spent in the base case is then 1- p_M. This proportion is a constant across both cases. The proportioned probabilities of failure on demand of the system are $P(TX)_{C,P}$ and $P(X)_{B,P}$ respectively. Then we have:

$$P(TX)_{C,P} = (1 - p_m)P(X)_{C,BC} + p_m P(X)_{C,M} \quad (12)$$

$$P(X)_{B,P} = (1 - p_m)P(X)_{B,BC} + p_m P(X)_{B,M} \quad (13)$$

A sensible question to ask is "what contribution of the weighted probability of failure on demand is does the maintenance state make?" This contribution is p_M $P(X)_{C,M}$ / $P(TX)_{C,P}$ and p_M $P(X)_{B,M}$ / $P(TX)_{B,P}$ for the conservative and best estimate case respectively. Table 4 considers how this contribution changes as the proportion of time spent in the maintenance state changes.

Table 4: The percentage contribution of maintenance

	Percentage Contribution of Maintenance	
Proportion of time in the Maintenance State	**Conservative**	**Best Estimate**
p_M = 12 hours per 365 days	12%	58%
p_M = 4 days per 365 days	53%	92%

The first row of Table 4 represents a case in which the time that spent in the maintenance case is very low, only 12 hours per year. The conservative analysis predicts that only one eighth of the total risk is due to the maintenance state, while the best estimate shows that in fact the majority of the annual risk (58%) is incurred during the twelve hour maintenance period.

Although this observation is very simple, it has important implications for how risk models are interpreted. Most quantitative analysts acknowledge that the absolute numerical prediction of the risk is not the most important contribution of risk models. As this example demonstrates, any conservative bias can result in the distortion of the risk profile, which may affect decisions and the allocation of resources. At present a culture of erring towards conservatism in safety risk models still exists, and this example provides a cogent reason to avoid conservatism.

Section 4 discusses the case when uncertainty in the probability of failure on demand is also taken into account.

4. BEST ESTIMATE PLUS UNCERTAINTY

The preceding sections have established that the use of conservative estimates can significantly alter the risk profile of a plant. However, there is a useful concept implicitly encoded in the use of conservative estimates; namely that of uncertainty in the estimate. Although it is not usually explicitly described, by vague allusions to the 95^{th} percentile confidence interval, conservative values do in fact implicitly take account of uncertainty in estimated values. However, this is not true of a best estimate point value, and represents a significant shortcoming of using best estimates in isolation. Where a conservative estimate overly penalises the risk and in so doing distorts the risk profile, best estimates alone ignore the issue of uncertainty making the results vulnerable to estimates which are not precise. The next logical stage in the process is to include uncertainty explicitly, based on a probability distribution around the best estimate. This has the effect of preserving the uncertainty information that is only encoded implicitly in conservative estimates.

It is useful to note that there are several common operations on distributions such as convolution, product distribution and multiplication of functions; more information on these is available from any standard text. Convolution of distributions is the required operation to combine events using an OR gate, however, this document exclusively considers AND gates so only the product distribution is required. The general form of this for 2 inputs to an AND gate is given below, where Z is the product distribution XY formed from the random variables X and Y.

$$P(Z) = \iint_{x,z} P\left(x, \frac{z}{x}\right) \left|\frac{1}{x}\right| dxdz \tag{14}$$

The general form for n inputs can be found by iteratively re-naming the product distribution. The resultant product distribution can look very different depending on the input distributions. However, a typical form is shown in Figure 5, below, for the product distribution of two lognormal distributions.

Figure 5: Product Distribution of Two Lognormal Distributions

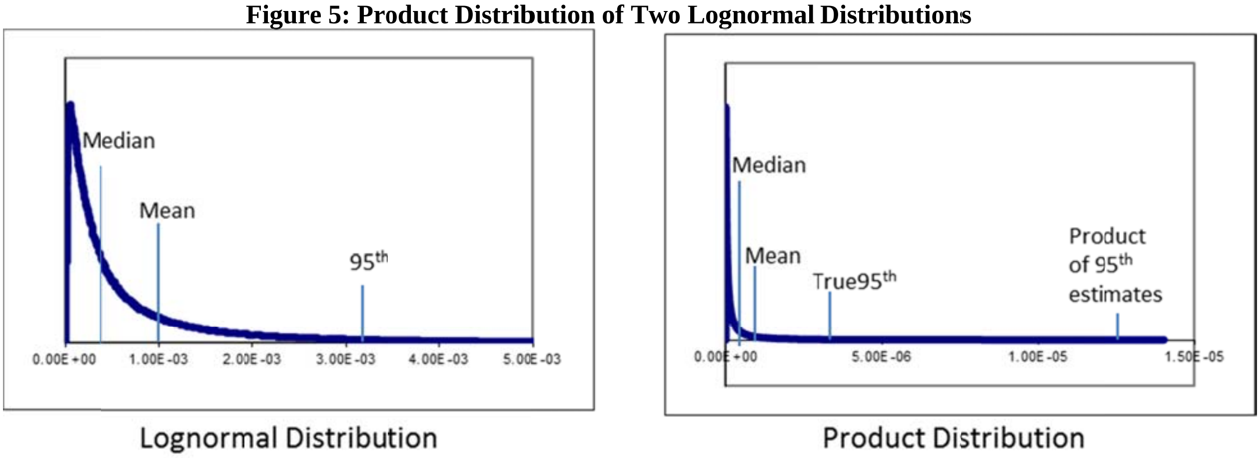

The mean, median and 95^{th} percentile are marked on the distribution in Figure 5. The product of two 95^{th} estimates is also marked on the product distribution to illustrate that it does not coincide with the true 95^{th} percentile of the product distribution. The true 95^{th} percentile of the product distribution (in the given example) is 3.7E-06, while the 95^{th} percentile calculated by just multiplying two 95^{th} percentiles together is 1.4E-05, which is actually the 99^{th} percentile of the true product distribution. This is a further distortive effect of conservative estimates that occurs after the use of just a single AND gate, and provides additional justification for not using them. If the mean is used as the best estimate then this estimate at least is not distorted by the process of using an AND gate; i.e. the value of the product of mean best estimates equals the mean of the product distribution of the two "true" distributions. However, using the mean alone fails to recognize the uncertainty information. The purpose of including a full uncertainty distribution is to add back in the information which is "lost" when moving from a conservative estimate to a best estimate paradigm. This "loss" of information is easily seen by considering two distributions with the same mean, but different levels of uncertainty. One distribution may be very narrow, while the other may be broad, but using the mean estimate for both would appear to equate the two. Hence, in most scenarios, simply using a best estimate will result in a naïve understanding of the risk. To complete the picture it is essential to have an explicit estimate of the uncertainty as well as using a best estimate of the central tendency.

5. DISCUSSION

The results presented here provide additional confirmation that the use of conservative judgements in risk analysis does not provide results which properly reflect the risk profile. In particular it has been shown that in the case of maintenance arguments, it leads to undervaluing the contribution to risk of removing lines of protection which have been conservatively assessed. This extent of this distortion increases the greater the conservatism. This can have a real impact on decision making; for example a particular barrier may afford (in reality) excellent protection but where analysis for it is very conservative, this would lead to a significant under-valuing of its protective capability. This conservatism may not be apparent in the cutset results and risk importance results. However, during maintenance it means that the risk model would fail to inform the analyst about the risk spike which would occur when that excellent barrier was unavailable due to maintenance.

There are parallels here with assessment of software. Software testing is an exceptionally hard problem which continues to challenge the software community [12]. A major contributor to the difficulty is the high dimension of the parameter space which needs to be checked, meaning that only a small volume can practically be checked. For this reason the current approach in assessing software in risk analyses is to use an ultra-conservative approach. To the author this is an outdated viewpoint, which needs to be addressed. While it is acknowledged that predicting software reliability is hard, it should still be subject to the best estimate philosophy; uncertainty estimates then provide a way to qualify that best estimate and to, rightfully, acknowledge that the software reliability is currently approximate.

An explanation for this attraction to conservative estimates is found in what appears to be the basic psychological wiring of humans showing an aversion to uncertainty, which has been dubbed the "uncertainty effect" [13]. Indeed, the uncertainty effect goes further than merely devaluing a package compared to the mean due to uncertainty; a package including uncertainty is often valued, subjectively, as worth less than the worst possible outcome. For example an uncertain lottery in which payouts are gift certificates with a face value between $50 and $100 is valued, by a significant proportion of people, as being worth less than a certain payout of a gift certificate with a face value of $50 [13]. This is a surprising result indeed, but the only point drawn from this result here is that this type of observation is indicative of the level of human aversion to uncertainty, even if the precise characterisation of that aversion requires further investigation in the psychological literature. However, this type of psychological bias must not be allowed to creep into the way in which risk analysis is performed. Even if psychological biases are unavoidable in the eventual decision making, psychological effects should be deferred as far down the process as possible; i.e. it should not feature until the "end" of the quantitative risk analysis estimation problem, after the risk profile has been estimated as faithfully as possible, including estimates of uncertainties as far as possible. This helps to avoid a compounding of these effects throughout the process.

Human perception of risk and reward is known to be a complex topic and subject to numerous psychological effects [14, 15]. Humans are bad at internalising small probabilities, and are known to distort the value of small probabilities, with a strong tendency over-value them. The value gradients are observed to be steep near certainty and near impossibility [14]. Since the absolute magnitude of small probabilities is not readily processed by humans, this presents a strong argument for the use of relative probabilities in assessing different scenarios, as far as possible. Relative probabilities are, in general, closer to the range of 50:50, which is a probability region in which humans tend to respond more rationally [14].

The aim of this type of work is to attempt to erase the prevailing mindset that it is better to be conservative than optimistic in risk estimates. This type of thought process certainly makes sense when it comes to design but is absolutely flawed when it comes to quantitative risk assessment, since it is like trying to push down a lump in a carpet. If you are conservative in one area (push down the lump) then you inadvertently neglect another area by distorting the risk profile (the lump pops up somewhere else). Not only does this type of behavior reduce how informative the analysis we have performed on areas of the system we understand well, but it also has the less well defined and pernicious effect of permitting the belief that, since we have been conservative in all our assessments, the overall values we are calculating are themselves conservative. The ill-stated implication is that, by being conservative for known sequences, we have implicitly allowed for model completeness uncertainty. Unfortunately, this is demonstrably untrue by a comparison between predicted values from PSA studies and the observed figures of reactor core melts and total reactor operating years accumulated worldwide, which is (at least) 3 severe accidents (Fukushima, Chernobyl, and Three Mile Island) in approximately 15,000 operating reactor years. There are strong mitigating arguments against this type of simplistic frequency observation, including the location dependence of hazards and the evolution of reactor design compared to reactors that suffered severe accidents. Nonetheless it provides a strong indication that current risk models may be missing significant risk contributors. The use of conservative assessments in the development of risk models could be acting to mask this conclusion by appearing to imply that risk models in their totality are also conservative; this in turn provides a loose rationale for the neglect of model completeness issues.

6. FURTHER WORK

This paper has demonstrated the distortion of the risk profile due to maintenance outages for a simplified model of a system in which there are ‘n’ lines of diverse protection, resulting in a particularly simple class of fault tree, but this analysis could be extended to more complex models. The analysis was greatly simplified by assuming that only AND gates were necessary. Events under OR gates could be replaced by a single new basic event with a different failure parameter; it is noted that by doing this the uncertainty distribution associated with the new basic event would be more complex, but this does not affect the overall argument above since no assumptions were made about the form of the uncertainty distribution $D(P_i)$. For this reason the results presented here are applicable to a general fault tree model, although further work could be done to definitively prove this claim using more complex fault trees. In addition to the level of complexity of the model, other aspects of risk models typically found in PSA models could also be included in the analysis. For example the use of time varying models, event trees and the use of boundary conditions to define scenarios of interest. This would lend even greater weight to the need for the use of best estimates, especially when the results of risk modelling are being used to inform decision making. Beyond strengthening the motivation for quantitatively assessing uncertainty, there are numerous maintenance analyses that could be usefully re-evaluated including uncertainty. For example the design of ‘optimal’ maintenance schedules could be strongly affected by the inclusion of uncertainty in failure parameter estimates.

The type of uncertainty considered in this document has only been statistical uncertainty. There are numerous other forms of uncertainty in PSA models, for example scenario uncertainty, success criteria uncertainty, accident progression and operator reliability uncertainty. Incorporating, explicitly, uncertainty from these sources would greatly benefit the predictions and insights that can be gained from PSA models. It is acknowledged that this represents a significant body of work and many uncertainties will require a bespoke method to incorporate. An example of the assessment of a “hidden” conservatism resulting from supporting neutronic analysis is presented in Reference 16.

7. CONCLUSION

The distortive effect of conservative estimates has been examined. This paper has demonstrated that the risk due to maintenance outages is underestimated if conservative values are used for failure parameters instead of best estimate values. It was then acknowledged that the conservative estimates, relied upon historically in the risk community, actually have intrinsic estimates of uncertainty bound up in them, and this partially justifies their use. It was shown that, in order not to distort the risk profile of a system, but while also retaining the uncertainty information implicit in conservative estimates, that best estimates alone are not sufficient and that best estimates plus uncertainty distributions are required. While there are clearly challenges in quantitatively finding a best estimate, and an estimate of the uncertainty, the author maintains that this is not a fundamentally different or more difficult task than producing a conservative estimate. A major difference is in fact exposure; whereas it is almost always possible to find a conservative number that few people would challenge, a best estimate is intrinsically more vulnerable to criticism. This is not necessarily a trite consideration; in some legal settings this could be of significance. However, from a pure risk quantification perspective it should always be desirable to develop the most accurate risk profile possible.

Acknowledgements
The author is very grateful to the EPSRC for funding and to Corporate Risk Associates for additional funding and technical support. Thanks in particular to Garth Rowlands and Rebecca Brewer for reviewing this paper and providing useful feedback. The author would also like to thank Dr. Simon Walker of Imperial College London and Jasbir Sidhu of Corporate Risk Associates for their continued support of this work.

REFERENCES

[1] G. E. Apostolakis, *"How Useful Is Quantitative Risk Assessment?"*, Risk Analysis, Vol.24, No.3, 2004.
[2] E. Zio and T Aven, *"Industrial disasters: Extreme events, extremely rare. Some reflections on the treatment of uncertainties in the assessment of the associated risks"*, Process Safety and Environmental Protection 91, 2013, pp 31–45.
[3] J. M. Reinert and George E. Apostolakis, *"Including model uncertainty in risk-informed decision making"*, Annals of Nuclear Energy 33, 2006, pp 354–369.
[4] N.S. Arunraj and J. Maiti, *"Risk-based maintenance—Techniques and applications"*, Journal of Hazardous Materials 142, 653–661, (2007).
[5] S. Turner, *"The Representation of Unavailability in Fault Trees and the Optimisation of Maintenance Actions"*, PhD Thesis, University of Birmingham, (1997).
[6] M. Bertolini, M. Bevilacqua, F.E. Ciarapica, G. Giacchetta, *"Development of Risk-Based Inspection and Maintenance procedures for an oil refinery"*, Journal of Loss Prevention in the Process Industries 22, pp. 244–253, (2009).
[7] S. Kaplan and B. J. Garrick, *"On The Quantitative Definition of Risk"*, Risk Analysis, Volume 1, No.1, pp.11-27, (1981).
[8] C. Bunea and T. Charitos & R.M. Cooke, *"Two-stage Bayesian models – application to ZEDB project"*, Reliability Engineering and System Safety, pp 321-329, (2005).
[9] E.L. Droguett, F. Groen & A. Mosleh, *"The combined use of data and expert estimates in population variability analysis"*, Reliability Engineering & System Safety, pp311-321, (2003).
[10] A. El-Shanawany, *"A Comparison of Bayesian Analysis Methods for Reliability Parameter Estimation in PSA"*, IET International System Safety Conference, 2010.
[11] K. Pörn, *"On Empirical Bayesian Inference Applied to Poisson Probability Models"*, Linköping Studies in Science and Technology, Dissertation No. 234, (1990).
[12] V. R. Basili and R. W. Selby, *"Comparing the Effectiveness of Software Testing Strategies"*, IEEE Transactions on Software Engineering, vol. se-13, no. 12, December 1987.

[13] Gneezy, Uri, John A. List, and George Wu. "The uncertainty effect: When a risky prospect is valued less than its worst possible outcome." The Quarterly Journal of Economics 121, no. 4 (2006): 1283-1309.

[14] Y. Rottenstreich, and K. H. Christopher, *"Money, kisses, and electric shocks: On the affective psychology of risk"*, Psychological Science 12, no. 3 (2001): 185-190.

[15] D. Kahneman and D. Lovallo, *"Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking"*, Management Science, Vol. 39, No. 1 (1993), pp. 17-31.

[16] A. El-Shanawany et al, *"Propagating Uncertainty in Phenomenological Analysis into Probabilistic Safety Analysis",* Submitted to PSAM12, 2014.