# AN ENGINEERING PRIMER ON EXTREME VALUE STATISTICS

**D.R. Novog[1] F. Hoppe[1], O. Nainer[2], B. Phan[3]**
[1] McMaster University, Ontario, Canada
[2] Bruce Power, Ontario, Canada
[3] Ontario Power Generation, Ontario, Canada

## Abstract

This primer is intended for individuals interested in gaining an understanding of Extreme Value Statistics (EVS). This work provides an explanation of EVS at a level that can be accessible to most people with an engineering or science background. While this work represents a simplification of the discussions from Reference 1, it is hoped that the authors will forgive any liberties taken in this paper. Some of the simplifications presented here may not be rigorous in all aspects, but the sacrifice in rigour is intended to aid the fundamental understanding of the EVS formulation and basic application.

## 1. Introduction

In existing and new nuclear power plants a variety of special safety systems are employed which will trigger fast reactor shutdown in the event of an accident or undesirable plant condition. These special safety systems utilize multiple and redundant measurements of certain process and neutronic variables, known as trip parameters, which are continuously monitored against pre-determined limits. If a measured trip parameter deviates in an unsafe direction in excess of these pre-determined limits, known as trip setpoints, the special safety system will initiate a fast reactor shutdown. Nuclear safety analysis is performed to determine the plant response to hypothetical accident scenarios and to assess the effectiveness of the trip parameters and setpoints in achieving the safety goals (i.e., precluding fuel failures or minimizing public dose). Hence, nuclear safety analysis is a critical component in the operation and regulatory licensing of nuclear power plants.

Historically, a set of bounding analysis methodologies and assumptions were used to determine plant response to these events. As a result of these simplifications it is impossible to determine the exact margins to safety limits. Furthermore, due to scientific discovery issues combined with plant safety margin deterioration due to component aging these traditional methodologies predict consequences which may prohibit full power operation. In addition to the above, changes in the regulatory framework for operating reactors is also driving changes in the methodology used to demonstrate plant safety [2]. Furthermore, Risk Informed Decision (RID) making practices and maintenance optimization at each plant rely on accurate quantification of the impact of upgrades/refurbishment on safety margins [3, [4]]. The Canadian Nuclear Safety Commission (CNSC) and the United States Nuclear Regulatory Commission (USNRC) have recognised that best estimate predictions of plant response, along with accurate assessments of uncertainties, is an acceptable alternative to more limiting and bounding analyses for demonstrating safety system response [5, 6].

At the forefront of these methodologies is EVS, which attempts to provide the most accurate quantification of setpoint or safety limit uncertainty. The key features of EVS include:

a) Division of the uncertainty into two components. The first component is the aleatory uncertainty which result from real changes in power plant conditions which may occur during operation. The second component of uncertainty, the epistemic uncertainty, results from our imprecise knowledge of variables, computer code models, and plant response.
b) Propagation of these uncertainties through all aspects of the physical system, in particular through the natural minimization and maximization functions that result from the mathematical description of the systems involved. This means that all safety system measurements, logic channels and response as well as overall fuel temperature responses in the core are tracked (or a meaningful subset of these channels).

The first application of these approaches was pioneered by Sermer and Olive (Reference 1 and 7) for the application of maximum channel power compliance. Further applications to Neutron Overpower Protection setpoint (8) and setpoint selection (9) are available in the literature. The intent of this paper is to provide a primer on the statistical nature and behaviour of the EVS method along with a basic application within the nuclear safety area.

## 2. Basic Definitions and Formulations

### 2.1 Nature of Uncertainty

Prior to examining some examples, it is important to define a **true value** of a given parameter. This parameter, Q, is the real value that we would observe if an error free measurement of that parameter were possible. We should note that Q may be a function of time and/or space, but it is still the true value for that parameter.
Since measurements free of error are not possible, the outcome is a measurement system that closely monitors the true value. The measured value is denoted as M. The relative difference between the measurement and the true value is defined as the measurement error[1]:

$$\varepsilon_M = \frac{M - Q}{Q}$$

The difference between the measured and true value may result from inaccuracies in the measurement system, time lags within the measurement system (i.e. an offset in time), and/or physical differences in the location of the measurement device (i.e. offsets in space). Reference 10 states that:

*An error in measurement is the difference between the recorded value and the true value.*

If we make many observations, a probability distribution of measurement errors can be obtained. Reference 10 also states that:

[1] For statistical analysis, we often rearrange the equation as $M = Q\left(1+\varepsilon_M\right)$.

*The uncertainty is a possible value that the error might take on for a given measurement. Since the uncertainty can take on various values, it is inherently a statistical quantity.*

## 2.2 Monte Carlo Methods

Statistical safety analyses have been performed in many areas for longer than 20 years. These methods often employ Monte Carlo computation techniques (although there are others) in order to establish the distribution of possible outcomes given the possible variations in input variables.

For example, suppose a predicted outcome is termed P, and is dependent on two variable parameters, x and y. Furthermore, assume x and y follow some known probability distribution (e.g., normal) with mean values $\mu_x$ and $\mu_y$, and standard deviations $\sigma_x$ and $\sigma_y$, respectively. If we sample these distributions according to their probability distributions, we will get values $x_1$ and $y_1$, and hence a resultant prediction, $P_1$. If we repeat this process n times, we will have n pairs, and a resultant set of P values (also from 1 to n). The terminology for a possible set of input values (say $x_{11}$ and $y_{11}$) is a **realization** of the input parameters. The entire calculation including the realization of x and y as well as the calculation of the output parameter is termed a **Monte Carlo pass**.

If many thousands of passes are performed, the resultant predicted probability distribution for P can be created. This distribution is often used in safety analysis to predict the likelihood of a given output, or the level of confidence, in predicting a certain value. Several alternatives exist to the above, including analytical convolution of the resultant uncertainty. However such methods are generally difficult to apply if the parental uncertainties are non-normal or are not easily modelled using closed formed relationships.

## 2.3 Extreme Functions

Extreme Value Statistics is concerned with the analysis of systems and functions of many variables that use maxima or minima. The fundamental concepts used in statistical analysis (true values, probability functions, levels of confidence, Monte Carlo techniques…) are applied to MIN/MAX problems and there are no additional statistical concepts involved. Essentially, EVS uses statistical concepts that are no different than those taught in undergraduate statistics. For example, if S is a variable that follows some known distribution (e.g., S has a normal distribution with a mean of 0.0 and a standard deviation of 0.5), then:

$$P\{S \le 1\} = 0.98$$

is understood to mean, "that 98% of the measurements of S will have a value less than or equal to 1". More generally:

$$P\{S \le C\} = 1 - \alpha$$

where $\alpha$ is some small number that can be used to specify a probability in our statement (e.g., if $\alpha = 0.02$, then 98% of the measurements of S will be less than or equal to C, some constant value).

Extreme functions (including maximum or minimum operators) operate on a set of variables and return a single value[2]. If S denotes a set of channel powers for a CANDU nuclear power station, it will contain an array of 380 or 480 values (depending on reactor design) corresponding to the channel powers for each channel in the core. Taking the maximum value of the set S, returns a single value, $S_{max}$ (i.e. there is only one maximum value for a given set). EVS is most successful when there are many components of the set S that are close to $S_{max}$, and we are interested in the probability of this maximum (or minimum) exceeding some value. More precisely, EVS is concerned with determining the $S_{max}$, not the particular value of each and every individual component of S (where traditional statistical applications are concerned with application of errors or uncertainties to individual components). The difference although subtle, is important. A robust treatment of the uncertainties in $S_{max}$ is the subject of the following text.

## 3. Basic EVS Formulation

### 3.1 Property of Sets of Measurements

In many engineering applications we are interested in the behavior of a specified parameter (for example, the power in a specific channel). In other applications we may not be interested in the behavior of a single parameter, but rather in the behavior of a **set of parameters** from a number of similar systems. In other cases we are concerned with a single feature of this set such as the maximum or minimum value (e.g., we may be especially interested in the maximum channel power for the set of fuel channels). EVS is related to the solution of the extrema for a given set. For a set we often examine quantities such as:

$$Q_{MAX} = Maximum\{Q_1, Q_2, Q_3, Q_4, \ldots, Q_n\}$$

where Q is some parameter of interest and there are n available values. Although $Q_{MAX}$ will correspond to a specific location in the set, say $Q_7$, in EVS we capitalize on the fact that we do not care where in the array of Q the maximum occurs, only what is the magnitude. If Q corresponds to a set of channel powers at a given CANDU station, for this given set of channel powers the maximum may be 6100kW. The location in the core where this maximum power of 6100kW occurs is not important if we only wish to know the maximum value in the set.

The remainder of the primer has many examples related to maximization problems, however the reader is encouraged to understand that the concepts are equally applicable to minimization problems. The details of EVS are rooted in rigorous statistical methods, however the concept is easy to grasp.

***A fundamental principle in EVS is that in the maximization or minimization of a set of randomly varying parameters, only the magnitude of the extreme value is desired, not its location within the set.***

[2] Other function (e.g., mean or median) also return a single value but we will fouc on the extreme functions in this paper.

### 3.2 Basic Prediction in a Single System

For this example we are interested in a variable, S, with an uncertainty that is normally distributed about a mean $\mu_S$ with a standard deviation $\sigma_S$. Let us assume we can measure S many times and create a histogram of the results.

If enough samples are available to construct an accurate histogram that approximates a normal curve, then additional measurements of S will follow this probability distribution function. *If we perform additional measurements for the S, what value u will bound 98% of these additional values*? In other words, we want the value u such that:

$$P\{S \leq u\} = 98\%$$

If S is normally distributed, then the solution is (approximately):

$$S' = \mu_S + 2\sigma_S$$

Another feature for this example system is that for a random realization of the system S, there is a 50% chance of realization above the mean value. That is:

$$P\{S \geq \mu_S\} = 50\%$$

Hence if the problem were such that we wanted to know the probability of a measurement being larger than the mean value, it would be 50%. The discussion of this simple problem may seem obvious, but this sets the stage for more complex and multi-component systems.

### 3.2 Basics of Multiple System Behaviour

For this case we have two identical systems, defined as $S^A$ and $S^B$ with each having a probability distribution function similar to the previous illustration (i.e. both system A and B have the same mean and standard deviation). If we are interested in knowing the probability of a measurement of System A exceeding its true mean, the solution is the same as previously discussed. The same is true if we wanted to know the same for System B. Mathematically this can be written for each system as:

$$P\{S^A \geq \mu\} = 50\%$$
$$P\{S^B \geq \mu\} = 50\%$$

Note for this sample system the mean of Systems A and B are equal, and denoted as $\mu$. The problem to be considered in this example is, *"What is the probability of either System A or System B measuring a value above the mean?"*, or equivalently, *"what is the probability of the maximum of Systems A and B measurements being above the mean?"*

For a given measurement of System A and B, several interesting results are observed when we start to solve the question posed above:

- System A and B independently have a probability of 50% of reading above the mean value.

- Effectively, there are 4 possible outcomes of this combined system for any set of measurements:

| System A | System B |
|---|---|
| Above Mean | Above Mean |
| Below Mean | Above Mean |
| Above Mean | Below Mean |
| Below Mean | Below Mean |

*Since all these outcomes are equally probable, there is a 75% probability that at least one of the systems will read above the mean.* That is,

$$P\left\{\max\left\{S^A, S^B\right\} \geq \mu\right\} = 75\%$$

Comparing this solution to the single system example, the probability of at least one system being greater than the mean has improved from 50% to 75%. Also note that there is equal probability that $S^A$ or $S^B$ will be the maximum value. That is,

$$P\left\{S^B > S^A\right\} = P\left\{S^A > S^B\right\} = 50\%$$

If the number of systems in this example was increased from 2 to 10, then the probability of at least one system measuring higher than the mean value is 99.9%.

$$P\left\{\max\left\{S^A, S^B, \ldots, S^J\right\} \geq \mu\right\} = \left(1 - \frac{1}{2^{10}}\right) \approx 99.9\%$$

This observation demonstrates the characteristics of a combined system.

a) If we perform a measurement for System A, and take that single value, we only have 50% probability that it will exceed the mean.
b) If we have separate measurements of Systems A and B and examine the maximum of the set, the probability he combined maximum is greater than the mean of either system improves to 75%.
c) For a collection of 10 independent systems, we can be very certain that the maximum of all measurements from the 10 systems will exceed the mean (with a probability of 99.9%).

The usefulness in the above examples is limited, as we do not often require a solution where we wish to guarantee that a single prediction is above the mean value (although it should be noted that analogous arguments hold for any percentage point, not just the mean). However note that in the 10 component system there is greater than 99% probability of exceeding the mean value, without using the addition of a 2-sigma quantity.

### 3.3 Concept of True Values

Let us assume that System A has a true value, $Q^A$ that is constant in time. Let us make an assumption that we have a very accurate measurement of System A that has no bias, but that the instrument has some random uncertainty. The reading of the instrument has a signal S and, due to the uncertainty in the instrument, this signal varies according to a probability distribution function. We can get some idea of the true value by taking many measurements (say some large number of measurements, n) and calculating the mean value of the resultant distribution. Since the measurements are assumed bias free, the mean of many measurements should approximate the true value very well.

The problem to be considered in this example is, *"what is the probability that the next measurement will exceed the true value for System A?"* In System A for the next measurement, this is:

$$P\left\{S^{A} \geq Q^{A}\right\} = 50\%$$

Similar to the previous section, let us assume there is an identical measurement (i.e., at the same location and at the same time) made with a different instrument (that varies with the same normally distributed random uncertainty). If System B is independent of System A, but has the same characteristics (i.e., the same mean and standard deviation), we have:

$$P\left\{S^{B} \geq Q^{B}\right\} = 50\%$$

and we can drop the superscript on Q. If we have the accurate measurements on both systems, *"what is the probability that at least one system measurements will be higher than the true value?"* The answer is:

$$P\left\{\max\left\{S^{A}, S^{B}\right\} \geq Q\right\} = 75\%$$

If we were to increase the number of systems in the set, this probability would increase. Note a fundamental concept here is that we have not increased the precision of the individual measurements in any way, nor has an additional allowance or multiplier been included, but the probability of measuring a value in excess of the true value has increased from 50 to 75%, simply by factoring the set of measurements in Systems A and B.

## 3.4 Concept of Participants

There was a noticeable difference in the predicted probability of exceeding some threshold when we examine multiple systems and the maximization function. Specifically the probability of a maximum exceeding some value in the examples increased from 50% for the case of a single system in isolation, to 99.9% for the case where we considered the overall maximum of 10 systems. This leads us to the idea of **participants**. In short, in a maximization (or minimization) problem, increasing the number of systems that have a reasonable probability to contribute to the maximum has a benefit[3] when examining EVS behavior. Consider the cases discussed below.

---

[3] The term "benefit" is applied loosely, in some applications EVS behaviour may contribute negatively to the desired result. The term benefit is used here as applied to the examples being considered.

Suppose there is a measured variable $S^A$ that approximates the true value, $Q^A$, of System A. We will also simplify the problem and assume again that the measurement closely approximates $Q^A$ and that the set of measurements denoted by $S^A$ varies normally about $Q^A$. We are interested in gaining an understanding of the probability that a measurement S will exceed the true value $Q^A$ for System A. The probability that some measurement $S^A$ will exceed $Q^A$ is 50%;

$$P\{S^A \geq Q^A\} = 50\%$$

Consider including System B, which also has measurements $S^B$ whose mean closely approximates the true value for the system, $Q^B$. However, let us assume that the true value of System B is significantly below the true value of System A. That is:

$$Q^A >> Q^B$$

A result of this assumption[4] is that the mean value of the measurements in System B will be less than the mean value for System A. We will also assume that both systems are characterized by the same standard deviation. Graphically this can be shown as:

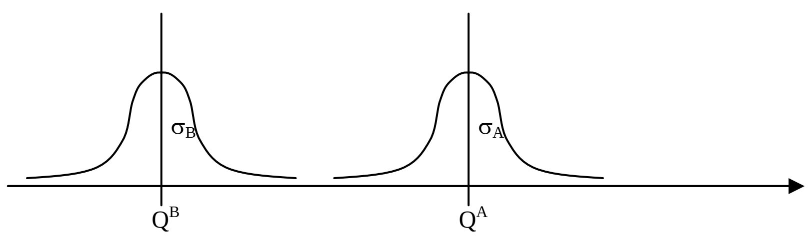


If we examine System B in isolation, we also see:

$$P\{S^B > Q^B\} = 50\%$$

The question is, *"what is the probability of either system measurements exceeding the maximum of all true values?"* The problem can be written mathematically as:

$$P\{\max\{S^A, S^B\} > \max\{Q^A, Q^B\}\} = ?$$

Note that for the combined system, the global maximum of the true values is $Q^A$:

$$Q_{\max} = \max\{Q^A, Q^B\} = Q^A$$

[4] In fact we require that $(Q^A - Q^B) / \sigma_b$ be large, where $\sigma_b$ is the standard deviation of measurements in System B. Interestingly, it does not depend on the uncertainty in channel A predictions.

Since $Q_A >> Q_B$, and σ in each system is significantly smaller than $(Q_A - Q_B)$, $S^A$ will be greater than $S^B$ with near certainty. That is:

$$S_{\max} = \max\{S^A, S^B\} = S^A$$

for almost all measurements. Therefore, the overall solution to this example is:

$$P\{\max\{S^A, S^B\} > \max\{Q^A, Q^B\}\} = 50\%$$

This results indicate that System B does not participate significantly in the maxima function (i.e., with the distribution os System A much greater than that of System B, B does not participate the maximization). For this example there was no change in the combined maximal behavior, as System A was sufficiently different than System B. In some discussions this is referred to as **peaked** behavior and will be discussed in more detail in Section 4. We should note that if the width of the distribution of System B was increased, or if the mean values were closer together, System B would participate in the maximization function, as is shown in the next example.

Assume we have the same systems as previously discussed, except that the true value of System B is much closer to that of System A. Graphically this can be shown as:

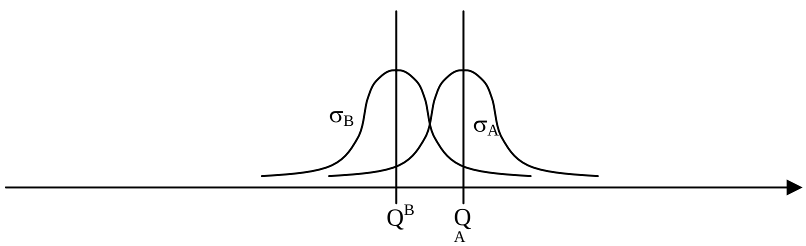


Note that $Q^B < Q^A$, but the means of the distributions are close together. Note that the maximum of the true values for the system remains $Q^A$. That is:

$$Q_{\max} = \max\{Q^A, Q^B\} = Q^A$$

Similar to the previous illustrations, when looking at each system in isolation we have:

$$P\{S^A > Q^A\} = 50\%$$

and

$$P\{S^B > Q^B\} = 50\%$$

However, there is now a significant probability that for some measurements $S^B$ will be greater than $S^A$. In fact there is a non-zero probability that $S^B$ will exceed $Q^A$ (which is the maximum true value over both systems). Hence for the measurements where the measurement $S^A$ under predicts $Q^A$, there is now a possibility that the measurement of $S^B$ will be sufficiently high that it will capture the maximum true value (which is not in System B, but is in System A, and has a value $Q^A$). Therefore:

$$P\left\{\max\left\{S^A, S^B\right\} > \max\left\{Q^A, Q^B\right\}\right\} > 50\%$$

The amount that the probability exceeds 50% will be dependent on the relative proximity of the mean of System B to the mean of System A. It will also be dependent on the width of distributions B and A. In other words, the distribution of the maximum is affected by the distribution widths[5].

This is termed the **participant effect**, which implies that by increasing the number of participant systems that are involved in the maximum operation, we can increase the confidence that at least one of the measurements will be greater than the true maximum. The concept of participants is related to the properties of combined systems and of the maxima/minima operators and can be summarized as:

***The number of participating systems (i.e., the systems that have a possibility to predict the maximum) will influence the probability that the maximum over these systems will exceed some fixed value (such as the true maximum over all systems).***

The participation effect is closely related to the concept of **flatness**. If the true values of all the individual systems are (relatively) close together, then the problem is such that it will likely have many participants and EVS will show large benefits. Problems with the potential for many participating systems are referred to as **flat**. If a small number of systems (or a single system) have true values significantly larger, than the number of participants may also be limited. This type of problem is referred to as **peaked** in nature, and the benefits of applying EVS to these problems may not be significant[6].

There are interesting and novel ways to achieve flatness. In fact one of the most difficult aspects of implementing EVS on more complicated systems is that the problem should be formulated such that it is as flat as possible. The channel power problem compliance problem is relatively straightforward EVS problem because it is naturally flat. Other systems can be made more flat by normalizing the prediction with respect to certain other parameters[7].

## 4. Channel Power Compliance

### 4.1 Basic Formulation

A complete derivation of the channel power compliance problem and the EVS application is provided in Reference 1. The following explanation is meant to highlight the features of EVS application within the context of the description provided in the preceding sections.

---

[5] However, as noted previously, the probability of the measured maximum over both systems exceeding the true maximum is unaffected by the width of distribution A.
[6] Note that flatness and peaked behaviour are concepts introduced here and are related to how near the true quantities being studied are to each other. They do not refer to the shape of the histograms involved.
[7] For example, the minimum channel power limit for heat transport low flow trip channels is normalised using a combination of the maximum and minimum channel powers.

If we could measure the channel powers with 100% accuracy, there would be 380 true channel powers. The channel powers must be below the license channel power limit. Mathematically this implies that:

$$\max_{i=1 to 380}\{Q_i\} \leq L$$

where Q denotes the true quantities and L is the limit. If all the true values were known, compliance would be straightforward, as we only require that the maximum of the true values is below the license limit. Unfortunately, the true values of the channel power are not known. At best there are a limited number of FINCH where the channel power is inferred from channel measurements. Therefore we must rely on simulations to predict the channel powers for each channel in the core.

As a first step, let us define the true value of the channel powers in the core. Typically, for each channel the true channel power is denoted as $Q_i$ where i corresponds to the channel number in the core (E.g., 1 to 380 for Pickering B). For the entire set of true channel powers there is maximum $Q_{max}$ defined as:

$$Q_{\max} = \max_{n=380}\{Q_1, Q_2, Q_3, Q_4, \ldots, Q_n\} = \max_{j=1 to 380}\{Q_j\}$$

For compliance we must satisfy:

$$Q_{max} < L$$

In Reference 1, a detailed discussion of hypothesis testing and null hypotheses is provided, and the approach presented there is rigorous in demonstrating channel power compliance. This primer will attempt to simplify the description by concentrating only on probabilistic treatment of the channel power predictions and the necessary compliance uncertainty to ensure the true value is exceeded with high confidence. The application of null hypothesis and testing is related to rigorous problem formulation, and this description of EVS formulations will not address the hypothesis testing issue.

We do not know the true channel powers, hence we do not know the true maximum channel power, $Q_{max}$. What we have available are SORO simulations of the channel powers. If we take the maximum value of the SORO channel powers from every channel we get:

$$S_{\max} = \max_{n=380}\{S_1, S_2, S_3, S_4, \ldots, S_n\} = \max_{i=1 to 380}\{S_i\}$$

So we know $S_{max}$ for a given reactor state. As stated above *"what multiple (or what value of compliance uncertainty allowance) of $S_{max}$ is required in order for us to be sufficiently confident that the result will be greater than the maximum true channel power"*. In other words:

$$P\left\{\max_{i=1 to 380}\{S_i\} \times (1+\eta_\alpha) \geq \max_{j=1 to 380}\{Q_j\}\right\} \geq 1-\alpha$$

Where $\alpha$ is some small percentage point.

## 4.2 Flatness in Channel Power Compliance

It should be clear by now that the reasons for EVS effectiveness[8] are:

- For any simulation of the channel powers, the channel where the simulated maximum channel power occurs (e.g., channel L12) does not necessarily have to be the location where the true value of the maximum occurs. For example, the channel with the true maximum may be channel M14. In fact, if the problem is nearly flat than for many simulations the predicted location may alter its location for each case.

- For certain cases where we do not predict the true value of the maximum channel power at its real location, there is a chance that for that simulation we may predict a value greater than the true value, although not at its correct location.

In the course of application of EVS to the channel power compliance problem, large benefits were realized in terms of operating margin and reactor power. Having completed all the discussions on EVS, we can now understand why it was so successful in this application. As fuelling engineers select channels for refueling they are trying to balance many parameters. One of the key parameters that are balanced on a day to day basis is channel power. Channels are selectively fueled to not only minimize zone and reactor power tilts, but channels are avoided if that channel or a neighbour may become limiting with respect to maximum channel power or CPPF.[9] This inherently means that one of the fuel engineers goals is to avoid peaking the core, i.e., to keep the channel powers in the CPPF region flat. This flatness is a major advantage when considering EVS and its important to note than the flatness results in [[11]]:

$$\eta_\alpha << 2\sigma$$

Where σ is the standard deviation in the measurement of channel power. The result indicates the improvement relative to the historical compliance uncertainty.

## 5. Conclusion

This primer has attempted to provide the basic statistical framework and physical understanding of the Extreme Value Statistics and its application. The following general observations and conclusions are made:

i. Quite often the literature associates the EVS terminology with the fitting of EVS or Gumbel type distributions and the analysis of rare events [12], however, in this application no distribution fitting is performed. Rather the distribution of the resultant histograms are a driven by the physical system, its modelling, and the input uncertainties provided.
ii. EVS benefits for safety and trip setpoint analysis benefit from the “cancelation of errors”, that is to say if errors are applied to channels randomly, then it is probable that even if the specific channel power is in error in a downwards direction, there is a high probability that a different channel will err on the high side, thus ensuring that the simulated maximum will tend to exceed the true maximum. Hence the location in the core where the simulated maximum occurs may differ from that of the true maximum.

[8] or efficiency as discussed in Reference 1
[9] CPPF, Channel Power Peaking Factor, is the ratio of a maximum channel power in the core to its NOP reference power in the CPPF region of the core.

iii. The "benefits" of EVS relative to traditional approaches for accounting for uncertainty result from the flatness of the set of variables being studies. In some cases, such as the channel powers in a CANDU reactor, the distribution may be naturally flat due to operational restrictions and performance. For other applications, the variables can be made flat by some form of normalization, such as through channel specific normalization.

## Acknowledgements

The author would like to sincerely thank the individuals who contributed to the work presented above. Specifically, Paul Sermer and Charles Olive, the authors of Reference 1 should be thanked for their development and implementation of EVS. Edward Panyan, Keith Weaver and Gabe Balog also provided considerable technical input. We would also like to thank the University Network of Excellence in Nuclear Engineering (UNENE) and NSERC for their support.

## References

[1] P. Sermer and C. Olive, "Probabilistic Approach to Compliance with Channel Power License Limits Based on Optimal Maximum Uncertainty", 1995 American Nuclear Society Annual Conference, Philadelphia, 1995.

[2] Canadian Nuclear Safety Commission, Proposed Regulatory Standard - S310, 2006.

[3] G. Geisler, S. Hellweg, K. Hungerbuhler, Int. J. of Life Cycle Assessment, 10(3), pp. 184-192, 2005.

[4] Pandey, M.D., "Extreme Quantile Estimation using Order Statistics with Minimum Cross-Entropy Principle", J. Probabilistic Engineering Mechanics, 16(1), 31-42, 2001.

[5] Canadian Nuclear Safety Commission, Regulatory Guide G-144, "Guidelines for Establishment of Shutdown System Trip Parameter Effectiveness", May, 2006.

[6] Technical Program Group, "Quantifying Reactor Safety Margins – Application of CSAU Methodology to a LBLOCA", EG&G Idaho, Inc., NUREG/CR-5249, December, 1989.

[7] C. Olive and P. Sermer, "Statistical Approach To Compliance With Maximum Channel Power Licence Limits In Ontario Hydro Reactors", Canadian Nuclear Society Conference, 1991.

[8] D. Novog, K. Atkinson, M. Levine, O. Nainer, B. Phan, Proc. Of the ASME – ICONE 16, 2008.

[9] Novog and Sermer, J. of Science and Technology of Nuclear Installations, 2008

[10] S.J. Kline, "The Purpose of Uncertainty Analysis", Trans. of the ASME, Vol. 107, June 1985, pp. 153-160.

[11] A. Firla and P. Sermer, "Probabilistic Analysis Of Maximum Channel Power Error - A Parametric Study", Procs. of the Canadian Nuclear Society, 1996.

[12] P. Embrechts, C. Klüppelberg, and T. Mikosch, Modeling Extremal Events for Insurance and Finance, Springer Verlag, 1997.