Journal of Reliability Science and Engineering

Avoiding Two Common Pitfalls in Uncertainty Propagation

To cite this article before publication: Matthias G.R. Faes *et al* 2025 *J. Reliab. Sci. Eng.* in press https://doi.org/10.1088/3050-2454/adde61

Avoiding Two Common Pitfalls in Uncertainty Propagation

Journal:	*Journal of Reliability Science and Engineering*
Manuscript ID	JRSE-100044.R1
Manuscript Type:	Paper
Keywords:	Uncertainty Quantification, reliability analysis, Monte Carlo simulation, normal distribution, interval analysis

Avoiding Two Common Pitfalls in Uncertainty Propagation

Matthias G.R. Faes[1,2], Alba Sofi[3] and Marcos A. Valdebenito[1]

[1]Chair for Reliability Engineering, TU Dortmund University, Dortmund 44227, Germany.
[2]International Joint Research Center for Engineering Reliability and Stochastic Mechanics, Tongji University, Shanghai, 200092, China.
[3]Department of Architecture and Design, University Mediterranea of Reggio Calabria, 89124 Reggio Calabria, Italy.

E-mail: matthias.faes@tu-dortmund.de, alba.sofi@unirc.it, marcos.valdebenito@tu-dortmund.de

12 May 2025

Abstract. This work addresses two common pitfalls present in the wider field of Uncertainty Quantification: the usage of a normality assumption for non-negative physical quantities, and the application of sampling or so-called discretisation schemes for the propagation of interval uncertainty. This first part of our work focuses on the normality assumption and its effect on the calculation of moments and probability of exceedance and it is developed around a simple yet illustrative example. Pitfalls associated with the assumption of normality are discussed and highlighted, showing that such an assumption can have a significant detrimental effect when performing uncertainty quantification. Assuming normality for non-negative physical quantities inherently leads to undesirable properties, such as non-existent moments of the response of interest or probabilities of exceedance with tails which become unreasonably heavy. With the second part of our work, we want to elaborate on both analytical and numerical evidence regarding interval uncertainty propagation. Both suggest that using sampling schemes to cope with intervals is extremely inefficient and inaccurate. We illustrate that performing sampling to propagate intervals yields a dramatic underestimation of the worst-case behavior of the problem under consideration at an unreasonable computational cost.

Keywords: Uncertainty Quantification, reliability analysis, Monte Carlo simulation, normal distribution, interval analysis

1. Introduction

Many advanced engineering modeling approaches deal with the solution of systems of partial differential equations (PDEs) that are formulated over continuous domains. Typically, these problems are formulated as:

$$\mathcal{N}_x(\boldsymbol{u};\boldsymbol{\theta}) = \boldsymbol{f}, \boldsymbol{x} \in D, \tag{1}$$

with boundary condition:

$$\mathcal{B}_x(\boldsymbol{u};\boldsymbol{\theta}) = \boldsymbol{b}, \boldsymbol{x} \in \Gamma, \tag{2}$$

where $\mathcal{N}_x$ is a differential operator, $D \subset \mathbb{R}^d, d \in [1,4]$ is the physical domain, $\boldsymbol{u} = u(\boldsymbol{x})$ is the solution of the PDE, and $\boldsymbol{\theta} = \theta(\boldsymbol{x}) \in \mathbb{R}^{n_\theta} \times D$ is the vector-valued field representing the parameters in the PDE. Additionally, $\boldsymbol{f} = f(\boldsymbol{x})$ is the forcing term on D, and $\mathcal{B}_x$ is a boundary condition operator which is defined on the domain boundary Γ. Often, such analyses are performed under the assumption that all parameters $\boldsymbol{\theta}$ can be quantified exactly at any point $\boldsymbol{x} \in D$. This is, of course, unrealistic in engineering practice, as we are faced with both the randomness of the structure and the uncertainty of our own observations of it. These phenomena manifest themselves to us as analysts as, respectively, aleatory [1] and epistemic uncertainty [2]. Often, in practice, both sources appear at the same time, as we are looking through imperfect lens (i.e., clouded by epistemic uncertainty) to a variable environment [3, 4]. To deal with this situation, research of the last four decades accumulated in a plethora of very powerful, practical and efficient techniques to propagate both aleatory and epistemic uncertainty, as well as combinations of those. However, based on our observations of published and unpublished works across all scientific journals related to uncertainty quantification and reliability analysis, we observe two common pitfalls:

- the assumption of a normal distribution for strictly non-negative model inputs $\boldsymbol{\theta}$;
- the application of sampling(-like) schemes for the propagation of epistemic set-valued uncertainty.

In this paper, we want to study the assumptions behind these pitfalls, assess their impact on the analysis results, and illustrate the problematic nature of either of these pathways.

Concerning aleatory uncertainty, it is assumed that the parameters $\boldsymbol{\theta} = \theta_i,\ i = 1, ..., n_\theta$ are affected by uncertainty, which is described by means of independent random variables with probability density function (PDF) $f_{\Theta_i}(\theta_i)$. In view of the assumption of independence, the joint PDF is $f_{\boldsymbol{\Theta}}(\boldsymbol{\theta}) = \prod_{i=1}^{n_\theta} f_{\Theta_i}(\theta_i)$, where $\boldsymbol{\theta} = [\theta_1, \ldots, \theta_{n_\theta}]$. The behavior of the system is synthesized in the so-called performance function $g(\boldsymbol{\theta})$, which assumes a value equal to or smaller than zero

whenever a combination of the uncertain input parameters $\boldsymbol{\theta}$ leads to an undesirable response, for example, loss of serviceability or collapse (see e.g., [5, 6]). Thus, the chance p_f that the system undergoes an undesirable behavior is given by the classical probability integral [7]:

$$p_f = \int_{-\infty}^{+\infty} I(g(\boldsymbol{\theta})) f_{\boldsymbol{\Theta}}(\boldsymbol{\theta}) d\boldsymbol{\theta}, \tag{3}$$

with $f_{\boldsymbol{\Theta}}(\boldsymbol{\theta})$ the PDF describing the aleatory uncertainty in the parameters $\boldsymbol{\theta}$, and $I(g(\boldsymbol{\theta}))$ the indicator function, which returns 1 when $\boldsymbol{\theta}$ is part of the so-called failure domain $\mathcal{F} = \{\boldsymbol{\theta} \mid g(\boldsymbol{\theta}) \leq 0\}$, and which is 0 otherwise. Both methods based on design-point related approaches [8] and simulation methods, such as importance sampling [1], line sampling [9] (also in a Bayesian interpretation [10]), directional (importance) sampling [6], subset simulation [11, 12], Bayesian approaches [13], and many others, have been introduced to estimate p_f. This rich spectrum of available methods, in combination with an unprecedented availability of computational power, enables us to assess, not only the reliability of a structure but also its sensitivity to perturbations [14] long before a prototype has been designed. Nonetheless, many authors default to the assumption of normality when modeling $f_{\boldsymbol{\Theta}}(\boldsymbol{\theta})$. Not only might this provide a biased view of reality, but it could also create serious issues when the quantity being modeled has a strict non-negative nature (e.g., a plate thickness or Young's modulus of a material). A more detailed analysis of this phenomenon and related problems will be discussed in Section 2.

For the propagation of epistemic uncertainty, interval models, in particular, have been shown to offer an objective representation of the extent to which our ignorance reaches. Efficient techniques based on optimization [15, 16, 17], perturbation analysis [18, 19, 20], interval arithmetic [21], affine arithmetic [22], improved interval analysis [23, 24], surrogate modeling [25, 26] and Bayesian cubature [27] have been introduced. Note that sampling methods have also been successfully applied based on Cauchy distributions [28, 29] or scenario optimisation [30]. However, despite this rich wealth of methods for efficient intrusive and non-intrusive interval analysis, we have observed that many authors still use one or the other variant of sampling methods to propagate intervals. This mistake is even found in highly cited papers from reputable and respectable journals. Presumably, this stems either from the false belief that there is some similarity between an interval and a uniform distribution, or from the belief that such propagation is harmless. In the second part of this paper, we aim to illustrate that both beliefs are not only incorrect but may also lead to a catastrophic underestimation of the worst-case behavior of the structure. Arguably even worse, we observe that many authors today, when proposing new approaches for the propagation of epistemic uncertainty, be it pure or in hybrid form, resort to sampling-based propagation to prove the efficacy of their techniques. Obviously, this is not valid, as new approaches need to be bench-marked against the

most efficient and accurate tools currently available. Sampling-based propagation of intervals possesses neither of these two properties. This creates a situation where 'novel' approaches are being bench-marked against a wrong golden-standard. These phenomena will be discussed in detail in Section 3. As an additional remark, it should be noted that considering a uniform (or bounded) distribution may be challenging even in the context of classical probabilistic analysis, as it introduces strong non-linearities when mapping to the standard normal space [31].

2. Normality assumption

2.1. Context

The normal distribution is one of the most commonly used in structural reliability analysis to describe uncertainty associated with model inputs of a numerical model [6]. Its widespread application is driven by several factors that make it both practical and mathematically convenient. Some of these factors are the following.

- **Data scarcity and maximum entropy**. In situations where data is scarce, the only available information may be limited to the mean and standard deviation of the input variables. When this is the case, the normal distribution becomes a natural choice due to the principle of maximum entropy [32, 33]. The normal distribution maximizes entropy for a given mean and variance, making it the least biased assumption. It does not introduce any additional information beyond what is provided by these two statistical moments, making it a reasonable model in the absence of further data.
- **Standard normal space**. A significant advantage of modeling uncertainty via normal random variables is its ease of transformation into standard normal space. In structural reliability analysis, methods such as the First-Order Reliability Method (FORM) and the Second-Order Reliability Method (SORM) frequently rely on such transformations[34]. The ability to map uncertain inputs to a standard normal distribution with mean zero and variance one simplifies failure probability calculations, especially for multidimensional problems.
- **Central limit theorem**. The central limit theorem further supports the use of the normal distribution. Indeed, according to the central limit theorem, the sum of a large number of independent, identically distributed random variables tends to follow a normal distribution, regardless of their initial distributions[6]. In structural systems, where uncertainties often result from the combination of several factors, the aggregated effect of these uncertainties may naturally approximate a normal distribution.
- **Mathematical convenience**. The normal distribution offers significant mathematical convenience. Its properties, such as symmetry and smoothness,

allow for analytical solutions and efficient numerical integration. These features are especially valuable in structural reliability analysis, where computational efficiency is often critical.

However, while the normal distribution has appealing properties, its adoption may lead to issues in particular cases. Indeed, as discussed in [35, 36], using a normal distribution to describe the uncertainty associated with a strictly positive quantity assigns nonzero probability to negative values. This in turn can lead to loss of coercivity, which is a mathematical condition ensuring that the differential operator remains well-behaved (i.e., bounded below by a positive constant). When coercivity is lost, the associated boundary value problem may become ill-posed. To address this issue, the use of uniform or lognormal distributions has been investigated in the literature, as discussed in e.g. [37, 38]. In addition, using the maximum entropy principle based on moments only may not be appropriate. In fact, when considering the first two moments of an uncertain variable plus the condition of strict positivity, the maximum entropy distribution is no longer normal. Therefore and based on the previous discussion, application of the normal distribution for modeling strictly positive quantities is wrong from both a mathematical viewpoint and the maximum entropy principle.

While there are several cases where the normal distribution may not be an appropriate choice for modeling uncertainty, this paper will not attempt to address all such instances, as that would be a daunting task. Instead, the focus will be limited to a specific case study where the uncertain input parameter of a structural model is strictly positive due to physical constraints. In such a case, using a normal distribution is particularly problematic, as it allows for negative values, which are physically impossible. This analysis will demonstrate how the assumption of a normal distribution in such scenarios can lead to inaccurate representations of uncertainty and, consequently, unreliable model predictions.

2.2. Case Study: Linear Spring

The case study considered here is a simple linear spring subjected to a unit force $f = 1$, as depicted schematically in Figure 1.

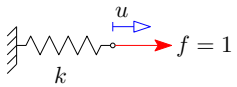

Figure 1. Linear spring with stiffness k subject to unit load f.

The spring is characterized by its stiffness, denoted as k. The system follows Hooke's law, where the displacement u under the unit force is inversely proportional

to the stiffness k, leading to the expression for displacement:

$$u = \frac{1}{k}, \tag{4}$$

where physical units have been omitted for the sake of simplicity. Note that the numerical model as cast in Eq. (4) is extremely simple. Such a model is selected on purpose to allow for analytical derivations. However, even such a simple model allows to extract very relevant conclusions which can be straightforwardly extended towards more complex models.

It is assumed that the spring stiffness k is affected by uncertainty and hence, it is modeled as a random variable. Note that due to physical considerations, it is known that $k > 0$. In particular, three different random variable models are chosen to characterize uncertainty, as described below.

- The first model considers a random variable K_N following a **normal distribution**. Specifically, $K_N \sim N(\mu, \sigma)$, where μ is the mean stiffness, and $\sigma > 0$ is the standard deviation representing the uncertainty. However, this model presents a limitation: the normal distribution allows K_N to take negative values, which violates the physical condition that the stiffness k must be positive. The probability density function (PDF) associated with K_N is given by:

$$f_{K_N}(k) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(k-\mu)^2}{2\sigma^2}\right), \tag{5}$$

  where $k \in (-\infty, \infty)$, even though negative values are non-physical.

- The second model is based on a **truncated normal distribution**. Here, the random variable K_{TN} follows a truncated normal distribution with shape parameters μ and σ and whose support is restricted to $[a, \infty)$, where a is a real constant such that $a \geq 0$. This restriction ensures that the stiffness k remains non-negative, in contrast to the normal distribution which allows negative values. The truncated normal distribution modifies the standard normal distribution by renormalizing the probability density function (PDF) over the interval $[a, \infty)$, ensuring that the total probability over this range equals 1. The PDF associated with K_{TN} is given by:

$$f_{K_{TN}}(k) = \frac{1}{\sigma\sqrt{2\pi}} \frac{\exp\left(-\frac{(k-\mu)^2}{2\sigma^2}\right)}{1 - \Phi\left(\frac{a-\mu}{\sigma}\right)}, \; k \geq a \tag{6}$$

  where $\Phi(\cdot)$ is the cumulative distribution function (CDF) of the standard normal distribution. Note that the shape parameters μ and σ are in general different from the mean and standard deviation, respectively, due to the parameter a restricting the support. However, if the renormalization constant $1 - \Phi\left(\frac{a-\mu}{\sigma}\right)$ is close to 1, then μ and σ become very close to the mean and standard deviation of the distribution.

- The third model assumes that k follows a **lognormal distribution**. This means that the natural logarithm of k, denoted by $\ln(k)$, follows a normal distribution with mean μ_G and standard deviation σ_G. The lognormal distribution ensures that the stiffness k remains strictly positive, making it a natural choice for modeling physical quantities that cannot be negative. The PDF associated with the lognormal distribution K_L is given by:

$$f_{K_L}(k) = \frac{1}{k\sigma_G\sqrt{2\pi}} \exp\left(-\frac{(\ln k - \mu_G)^2}{2\sigma_G^2}\right), \quad k > 0 \tag{7}$$

  The parameters μ_G and σ_G (mean and standard deviation of the underlying normal distribution) can be computed in terms of the mean μ and standard deviation σ of the lognormal distribution as follows:

$$\mu_G = \ln\left(\frac{\mu^2}{\sqrt{\mu^2+\sigma^2}}\right) \tag{8}$$

$$\sigma_G = \sqrt{\ln\left(1+\frac{\sigma^2}{\mu^2}\right)} \tag{9}$$

  Please note that in all definitions in Eqs. (7)-(9), it is implicitly considered that $\mu \geq 0$ and $\sigma \geq 0$.

Sections 2.3 and 2.4 focus on analyzing the consequences of adopting any of the three probabilistic models described above for calculating the second-order statistics and probability of exceedance of the displacement of the linear spring problem described in Eq. (4).

2.3. Mean and Variance of the Response

In the analysis of systems with random inputs, *second-order statistics* (namely, mean and standard deviation) are fundamental in quantifying the uncertainty of the response output [39]. With reference to the spring with uncertain stiffness, the mean μ_U of the random displacement U gives an idea of the system's typical behavior, while the standard deviation σ_U quantifies the spread or uncertainty of the displacement about the mean. However, while second-order statistics are important, they are often not sufficient for capturing the complete behavior of a random system. The *third-* and *fourth-order statistics* (that is, skewness and kurtosis, respectively) provide deeper insights into the asymmetry and tail behavior of the displacement distribution [40]. These higher-order moments are particularly important when the system exhibits non-normal behavior or when the tails of the distribution significantly influence the system's risk and performance. However, for the sake of simplicity, the focus is on second-order statistics in the following.

The PDF of the random variable K modeling the uncertain stiffness of the spring is considered to compute the mean and standard deviation of the displacement. Note that K can assume any of the three PDFs defined previously, that is: normal, truncated normal, or lognormal. Thus, the mean of the displacement μ_U is defined as [5]:

$$\mu_U = \mathbb{E}[u] = \int_{\Omega_k} \frac{1}{k} f_K(k)\, dk \tag{10}$$

where Ω_k denotes the domain associated with K and $\mathbb{E}[\cdot]$ denotes expectation. The standard deviation σ_U is equal to the square root of the variance, which is defined as[5]:

$$\sigma_U^2 = \mathbb{V}[u] = \int_{\Omega_k} \left(\frac{1}{k} - \mu_U\right)^2 f_K(k)\, dk \tag{11}$$

where $\mathbb{V}[\cdot]$ denotes variance. In the following, the values assumed by the mean μ_U and standard deviation σ_U of the displacement are discussed, considering the three distribution models for the random stiffness, namely normal, truncated normal and lognormal distributions.

2.3.1. Mean and Variance of the Response – Case of Stiffness Following Normal Distribution To start calculating specific values for the mean and standard deviation of the displacement, the case of a **normal distribution** for the stiffness is considered first. The expected value of U is given by the following integral:

$$\mu_U = \int_{-\infty}^{\infty} \frac{1}{k} f_{K_N}(k)\, dk. \tag{12}$$

The last integral is split into two parts to facilitate its calculation.

$$\mu_U = \int_{-\infty}^{\infty} \frac{1}{k} f_{K_N}(k)\, dk = \underbrace{\int_{-\infty}^{0} \frac{1}{k} f_{K_N}(k)\, dk}_{I_1} + \underbrace{\int_{0}^{\infty} \frac{1}{k} f_{K_N}(k)\, dk}_{I_2} \tag{13}$$

To analyze the behavior of I_2, it is split into two parts:

$$I_2 = \int_{0}^{\infty} \frac{1}{k} f_{K_N}(k)\, dk = \underbrace{\int_{0}^{\mu} \frac{1}{k} f_{K_N}(k)\, dk}_{I_{2a}} + \underbrace{\int_{\mu}^{\infty} \frac{1}{k} f_{K_N}(k)\, dk}_{I_{2b}} \tag{14}$$

where I_{2a} covers the range $[0, \mu]$, and I_{2b} covers the range $[\mu, \infty]$. Now, focusing on I_{2a}:

$$I_{2a} = \int_{0}^{\mu} \frac{1}{k} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(k-\mu)^2}{2\sigma^2}\right) dk. \tag{15}$$

This last integral can be bounded from below by replacing $(k-\mu)$ by μ as argument of the exponential function, leading to:

$$I_{2a} \geq \int_0^{\mu} \frac{1}{k} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) dk = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) \int_0^{\mu} \frac{1}{k}\, dk \tag{16}$$

The integral $\int_0^{\mu} \frac{1}{k}\, dk$ is known to diverge as $k \to 0$. Therefore, I_{2a} diverges to $+\infty$ due to the behavior of $\frac{1}{k}$ as $k \to 0$. Given this result, it can be readily shown that the integral I_2 tends to $+\infty$ as well. Using similar arguments, it can be shown that I_1 tends to $-\infty$. This implies that the expected value μ_U is undefined because of the behavior of the integrals I_1 and I_2. As the expected value is undefined, it becomes evident that the standard deviation σ_U is also undefined, because its calculation depends on the mean value μ_U; for a formal proof, please refer to [41], as it covers precisely the calculation of the second-order moment. To gain further insight into the non-existence of the mean and standard deviation of the displacement, the following numerical experiment is carried out. It is assumed that the mean and standard deviation of the normally-distributed stiffness are $\mu = 2$ and $\sigma = 1$, respectively. Then, a set of N samples of the stiffness is generated, which is used to calculate the corresponding samples of the displacement with Eq. (4). Thereafter, we attempt to calculate the mean value of the displacement using the classical formula $\frac{1}{N}\sum_{i=1}^{N} u^{(i)}$, where $u^{(i)}$ is the i-th sample of the displacement. The results obtained are shown in Fig. (2) as a function of the sample set size N for three independent runs, which are shown with yellow, orange and blue colors. For these three independent runs, it is possible to observe that the estimate of the mean does not converge to any particular value. Thus, this numerical experiment confirms the analytical results.

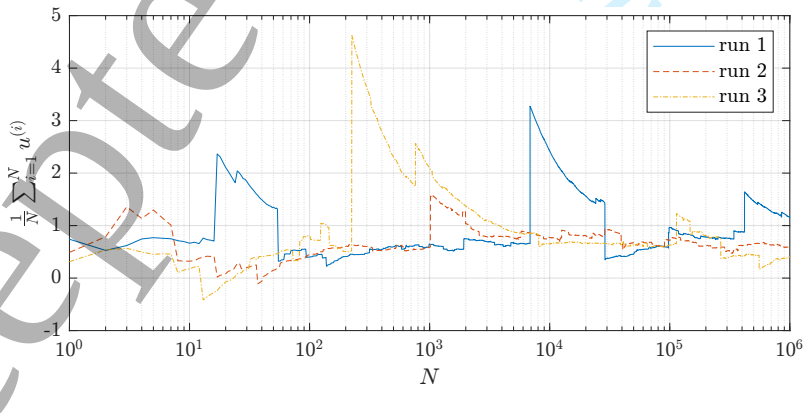

Figure 2. Estimator for the mean value of the displacement of the spring when the stiffness follows a normal distribution.

The fact that both the mean and standard deviation of the displacement do not exist when the stiffness of the spring follows a normal distribution may seem surprising at first. However, such behavior has already been acknowledged before,

as discussed in, e.g. [42, 43]. In fact, as long as the PDF associated with the stiffness value $k = 0$ is larger than zero, the mean and standard deviation of the displacement become undefined or tend to infinity, depending on the specific situation. Furthermore, it should be noted that the non-existence of the mean and standard deviation of the displacement does not depend on the particular values assumed by the statistical properties of the spring stiffness. In other words, assuming a normal distribution for the stiffness of the spring immediately leads to a displacement without mean or standard deviation, no matter how small the tail of the normal distribution is for negative values of the stiffness.

2.3.2. Mean and Variance of the Response – Case of Stiffness Following a Truncated Normal Distribution The second case under consideration corresponds to the stiffness following a truncated normal distribution over the interval $[a, +\infty)$. For such a case, the expected value of the displacement is:

$$\mu_U = \int_a^{\infty} \frac{1}{k} \frac{1}{\sigma\sqrt{2\pi}} \frac{\exp\left(-\frac{(k-\mu)^2}{2\sigma^2}\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)} dk \tag{17}$$

while the standard deviation of the displacement is:

$$\sigma_U = \sqrt{\int_a^{\infty} \left(\frac{1}{k} - \mu_U\right)^2 \frac{1}{\sigma\sqrt{2\pi}} \frac{\exp\left(-\frac{(k-\mu)^2}{2\sigma^2}\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)} dk}. \tag{18}$$

For the case where $a = 0$, the expected value integral in Eq. (17) is similar to the one in Eq. (14), except for the factor $1 - \Phi\left(\frac{a-\mu}{\sigma}\right)$. Thus, it is readily seen that for the case where the stiffness is modeled as a truncated normal distribution over the support $[a = 0, \infty)$, the mean value of the displacement tends to infinity. To verify such a result, the following numerical experiment is carried out. N samples of the stiffness following a truncated distribution with parameters $\mu = 2$, $\sigma = 1$ and $a = 0$ are generated. Then, the mean value of the displacement is estimated, as shown in Figure 3. The results obtained for three independent runs (which are denoted with yellow, orange and blue colors) show that the estimated mean values increase as the number of drawn samples N increases.

Arguments similar to the ones discussed before allow deducing that the standard deviation of the displacement becomes undefined for the case where the stiffness follows a truncated distribution over the interval $[a = 0, \infty)$. Demonstrating such a result is straightforward, however, the detailed steps are omitted here for the sake of brevity.

Whenever the uncertainty of the stiffness is modeled using a truncated normal distribution with $a > 0$, then both the mean μ_U and standard deviation σ_U of the displacement become real numbers. While there are no closed-form formulas for the

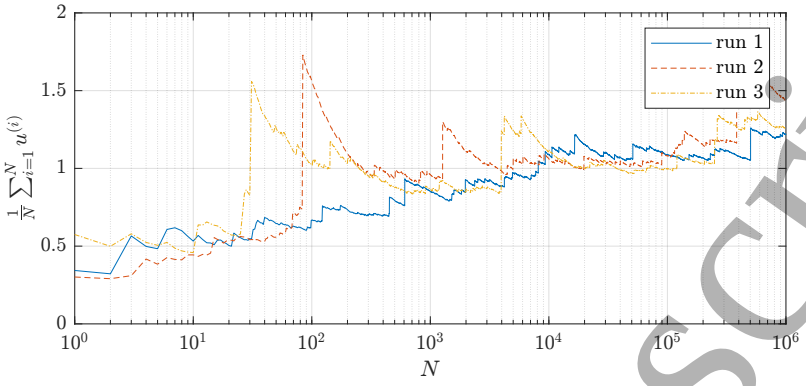

Figure 3. Estimator for the mean value of the displacement of the spring when the stiffness follows a truncated normal distribution with $a = 0$.

mean and standard deviation of the displacement in Eqs. (17) and (18), respectively, efficient one-dimensional quadrature schemes can be implemented to calculate these second-order statistics. Nevertheless, when characterizing the uncertainty of the stiffness with a truncated normal distribution, care must be taken in choosing the truncation parameter a, as it may possess a significant effect on the calculated second-order statistics. To demonstrate the latter point, consider the following example. The stiffness in Eq. (4) is modeled following a truncated normal distribution as shown in Eq. (6) with parameters $\mu = 4$, $\sigma = 1$ and $a \in [10^{-6}, 10^{-1}]$. The results obtained for the mean μ_U and standard deviation σ_U of the displacement are shown in Figure 4. It is noted that while the mean is more or less stable around the value of 0.27, the standard deviation varies between two orders of magnitude depending on the specific value chosen for the truncation parameter a. This highlights that the standard deviation of the displacement is extremely sensitive to the particular selection of the truncation parameter. In consequence, considering a truncated normal distribution for modeling the uncertainty of the stiffness may not be convenient unless the effect of the truncation parameter is studied in depth, or solid engineering arguments exist to impose such a bound.

To gain further insight into the challenges associated with the application of a truncated normal distribution, a particular setting of the last example is investigated. The parameters of the truncated normal distribution that characterize the stiffness are selected as $\mu = 4$, $\sigma = 1$ and $a = 10^{-6}$. Then, both the mean μ_U and standard deviation σ_U of the displacement are estimated using Monte Carlo simulation (MCS) with a set of samples of size N. The evolution of the estimates of μ_U and σ_U as a function of N are shown in Figure 5 with solid blue line. In addition, the reference values calculated with quadrature (Q) are shown with a dashed red line. An analysis of the results obtained with Monte Carlo indicates that while the mean μ_U can be reasonably well estimated with about $N = 10^3$ samples, the estimator for the

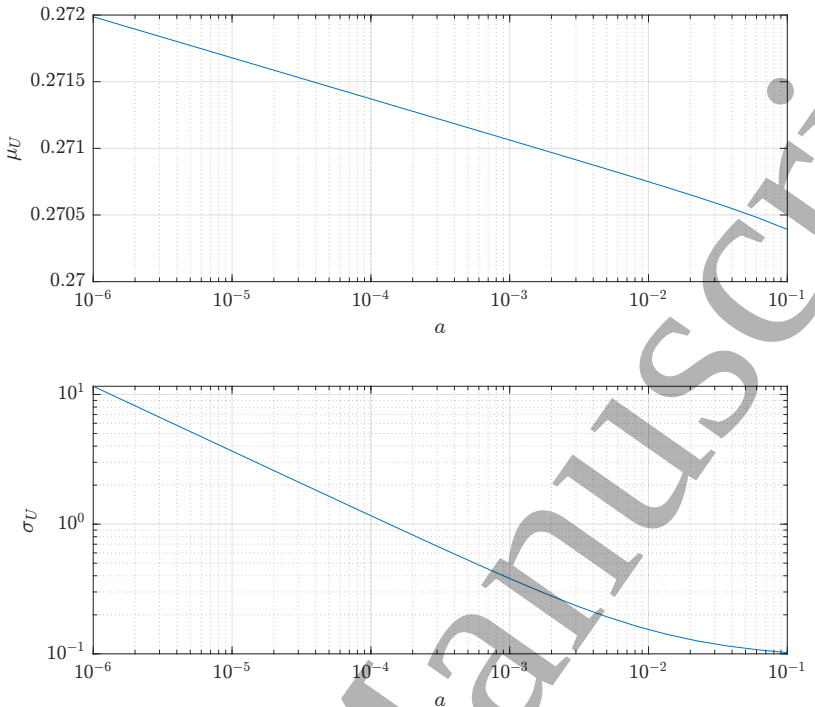

Figure 4. Mean value μ_U and standard deviation σ_U of the displacement of the spring when the stiffness follows a truncated normal distribution with parameters $\mu = 4$, $\sigma = 1$ and $a \in [10^{-6}, 10^{-1}]$.

standard deviation σ_U does not converge even after drawing $N = 10^8$ samples. In fact, it is observed that when $N \approx 3 \times 10^6$, there is an abrupt jump in the estimators of both μ_U and σ_U. The reason is that a value of the stiffness k which is close to the lower bound $a = 10^{-6}$ is sampled, thus leading to a huge sample of the displacement that explains the jump. Such an issue only highlights again that the selection of the truncation parameter for a truncated normal distribution plays a major role. Hence, extreme care must be applied when considering a truncated normal distribution, especially with respect to its truncation parameter.

2.3.3. Mean and Variance of the Response – Case of Stiffness Following a Lognormal Distribution The third case under consideration corresponds to the stiffness following a lognormal distribution with mean μ and standard deviation σ, whose probability density function $f_{K_L}(k)$ is defined in Eqs. (7)-(9). In such a case, the mean μ_U and standard deviation σ_U of the displacement in Eq. (4) can be calculated in closed-form, as shown below.

$$\mu_U = \int_0^\infty \frac{1}{k} f_{K_L}(k)\, dk = \frac{1}{\mu}\sqrt{1 + \frac{\sigma^2}{\mu^2}} \tag{19}$$

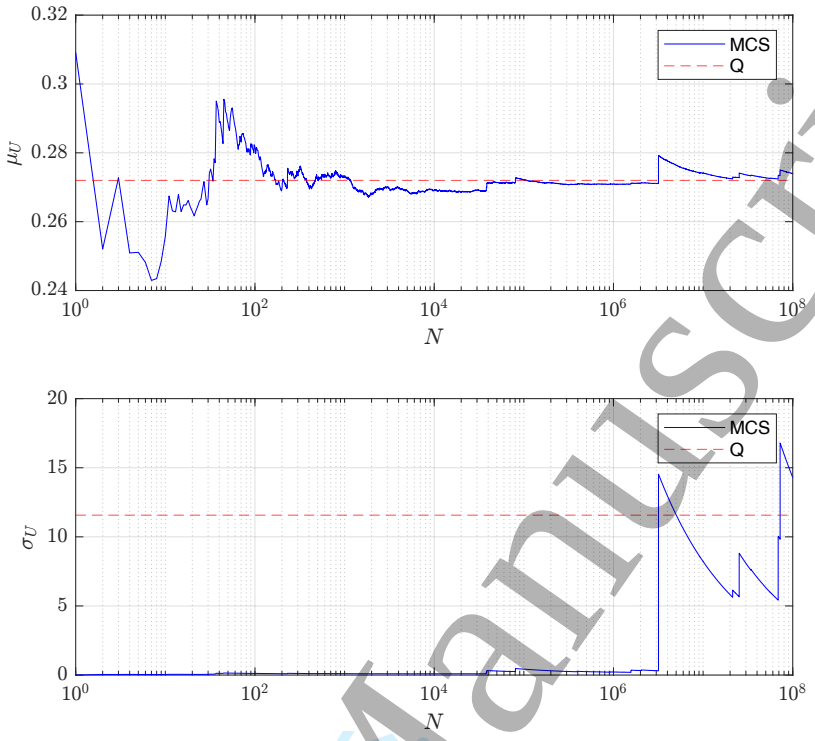

Figure 5. Estimators for the mean value μ_U and standard deviation σ_U of the displacement of the spring as a function of the number of samples N when applying Monte Carlo simulation (MCS). The dashed lines show the reference results obtained through numerical quadrature (Q).

$$\sigma_U = \sqrt{\int_0^\infty \left(\frac{1}{k} - \mu_U\right)^2 f_{K_L}(k)\, dk} = \sqrt{\frac{1}{\mu^2}\left(1 + \frac{\sigma^2}{\mu^2}\right) - \frac{1}{\mu^2}\sqrt{1 + \frac{\sigma^2}{\mu^2}}} \tag{20}$$

As the lognormal distribution ensures that the stiffness k is strictly positive, both values of the second-order statistics of the displacement, as shown in Eqs. (19) and (20) , are well defined for different combinations of the mean $\mu > 0$ and standard deviation $\sigma > 0$ of the stiffness.

2.4. Probability of Exceedance of the Response

The probability of exceedance is a key metric in assessing the performance of systems under uncertainty. This probability quantifies the chances that a given output response variable of a system will exceed a specified threshold, which is crucial for evaluating both serviceability and ultimate limit states. In the following analysis, attention is focused on the exceedance probability associated with the tip displacement of the spring as described in Eq. (4) when the uncertainty of the stiffness is modeled as a random variable following normal, truncated normal, and lognormal

distributions. The probability of exceedance is defined as the chances that the random variable U associated with the displacement exceeds a threshold u, that is:

$$P\left[U > u\right] = 1 - P\left[U \leq u\right] = 1 - F_U(u) \tag{21}$$

where $P[\cdot]$ denotes probability of the argument and F_U is the CDF associated with the random variable U. Hence, calculating the exceedance probability is equivalent to calculating one minus the CDF. In the following, this CDF is deduced for the three specific distributions considered to characterize the uncertainty of the stiffness. These expressions are then compared to each other in Section 2.4.4.

2.4.1. Cumulative Distribution Function – Case of Stiffness Following Normal Distribution To calculate the sought CDF in case where the uncertainty in the stiffness is described by a normal distribution, it is necessary to solve:

$$F_U(u) = P\left[U \leq u\right] = P\left[\frac{1}{K} \leq u\right]. \tag{22}$$

As the random variable K admits negative and positive values for the stiffness, the inequality $\frac{1}{K} \leq u$ must be solved taking into account the cases where $u < 0$ and $u \geq 0$. Starting with the case $u < 0$, the solution of the inequality $\frac{1}{K} \leq u$ becomes $\frac{1}{u} \leq K \leq 0$. Then, when $u < 0$, the CDF is:

$$F_U(u) = P\left[\frac{1}{u} \leq K \leq 0\right]. \tag{23}$$

Recalling that K follows a normal distribution with mean μ and standard deviation σ, it is found that:

$$F_U(u) = \Phi\left(\frac{-\mu}{\sigma}\right) - \Phi\left(\frac{\frac{1}{u} - \mu}{\sigma}\right) \tag{24}$$

Now for the case where $u \geq 0$, the inequality $\frac{1}{K} \leq u$ implies that either $K \leq 0$ or $K \geq \frac{1}{u}$. Therefore:

$$F_U(u) = P\left[K \leq 0\right] + P\left[K \geq \frac{1}{u}\right]. \tag{25}$$

Recalling again that K follows a normal distribution, the last expression simplifies to:

$$F_U(u) = \Phi\left(-\frac{\mu}{\sigma}\right) + \left(1 - \Phi\left(\frac{\frac{1}{u} - \mu}{\sigma}\right)\right). \tag{26}$$

Combining both cases analyzed above, it is found that the sought CDF $F_U(u)$ is:

$$F_U(u) = \begin{cases} \Phi\left(\frac{-\mu}{\sigma}\right) - \Phi\left(\frac{\frac{1}{u} - \mu}{\sigma}\right), & \text{if } u < 0, \\ \Phi\left(-\frac{\mu}{\sigma}\right) + \left(1 - \Phi\left(\frac{\frac{1}{u} - \mu}{\sigma}\right)\right), & \text{if } u \geq 0. \end{cases} \tag{27}$$

It is important to note that although the displacement does not possess mean or variance, its CDF does exist. Such characteristic is also observed in other well-known cases such as, e.g. the Cauchy distribution, and is typical of uncertain quantities that exhibit a heavy-tailed behavior.

2.4.2. Cumulative Distribution Function – Case of Stiffness Following Truncated Normal Distribution In this case, the stiffness adopts values belonging to the interval $[a \geq 0, \infty)$. To calculate the sought CDF, the starting point is again:

$$F_U(u) = P\left[U \leq u\right] = P\left[\frac{1}{K} \leq u\right]. \tag{28}$$

Three cases need to be distinguished when solving the inequality $\frac{1}{K} \leq u$. First, u cannot be negative, as the stiffness is always positive. Therefore, $F_U(u) = 0$ whenever $u < 0$. The second case is that the upper bound for the displacement is $\frac{1}{a}$, as the smallest value that the stiffness may assume is a. Thus, $F_U(u) = 1$ whenever $u > 1/a$. The third case is considering that $0 \leq u \leq \frac{1}{a}$, which is fulfilled whenever $K \geq \frac{1}{u}$. Thus:

$$F_U(u) = P\left[K \geq \frac{1}{u}\right]. \tag{29}$$

For a truncated normal distribution, this probability is:

$$F_U(u) = \frac{1 - \Phi\left(\frac{\frac{1}{u}-\mu}{\sigma}\right)}{1 - \Phi\left(\frac{a-\mu}{\sigma}\right)}. \tag{30}$$

Combining the three cases described above, the complete CDF of U is:

$$F_U(u) = \begin{cases} 0, & \text{if } u \leq 0, \\ \frac{1-\Phi\left(\frac{\frac{1}{u}-\mu}{\sigma}\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)}, & \text{if } 0 < u \leq \frac{1}{a}, \\ 1, & \text{if } u > \frac{1}{a}. \end{cases} \tag{31}$$

2.4.3. Cumulative Distribution Function – Case of Stiffness Following Lognormal Distribution When the uncertainty associated with the stiffness is modeled as a lognormal random variable, the stiffness adopts values belonging to the interval $(0, \infty)$. To calculate the CDF of the displacement, it is noted that u cannot be negative, as the stiffness is always positive and thus, $F_U(u) = 0$ whenever $u \leq 0$. For the case where $u > 0$:

$$F_U(u) = P\left[U \leq u\right] = P\left[\frac{1}{K} \leq u\right] = P\left[K \geq \frac{1}{u}\right]. \tag{32}$$

Recalling that $\ln K$ follows a normal distribution with mean μ_G and standard deviation σ_G (see Eqs. (8) and (9)), the last expression is equal to:

$$F_U(u) = P\left[\ln K \geq \ln\left(\frac{1}{u}\right)\right] = 1 - \Phi\left(\frac{\ln\left(\frac{1}{u}\right) - \mu_G}{\sigma_G}\right). \tag{33}$$

In summary, the sought CDF is in this case equal to:

$$F_U(u) = \begin{cases} 0, & \text{if } u \leq 0, \\ 1 - \Phi\left(\frac{\ln\left(\frac{1}{u}\right) - \mu_G}{\sigma_G}\right), & \text{if } u > 0. \end{cases} \tag{34}$$

2.4.4. Probability of Exceedance – Comparison Between Different Distribution Types Associated with Stiffness Now that the CDFs associated with the displacement U for the three different distribution types for the stiffness have been determined, it is possible to calculate the probability of exceedance and also compare the results obtained with each of these three different models. For such purpose, the parameters μ and σ for the three distributions (normal, truncated normal and lognormal) are chosen as $\mu = 5$ and $\sigma = 1$, while the truncation parameter for the truncated normal is selected as either $a = 0.1$ or $a = 0.2$. The results obtained are depicted in Figure 6. From these results, it is observed that the four cases analyzed provide similar probabilities of exceedance up to the threshold level of about $u \approx 0.25$. However, for larger threshold levels, the results associated with the normal and truncated normal distribution differ significantly with respect to those of the lognormal distribution. Furthermore, there are also differences (albeit less pronounced) between the normal and truncated normal cases. The differences between normal/truncated normal and lognormal cases can be explained as follows. In both normal/truncated normal cases, the probability density associated with small values of the stiffness is too large, at least when compared to the lognormal case. Assigning more probability density to those small values implies that the probability distribution associated with the displacement becomes heavy-tailed, leading to values of the probability of exceedance which are orders of magnitude larger than those associated with the lognormal distribution. Such a behavior is not surprising. It had already been shown previously that when considering a normal distribution for modeling the stiffness, the displacement does not possess second-order statistics (see e.g. Figure 2). And for the case of a truncated normal distribution, the variance is highly sensitive to the truncation parameter (see e.g. Figure 4). Such behavior is typical of heavy-tailed distributions (such as the Cauchy distribution) and is clearly observed in Figure 6.

All of the observations described above suggest that the normal or truncated normal distributions may lead to overly conservative values of the probability of exceedance. Such overly conservative results may be undesirable, as they defeat one

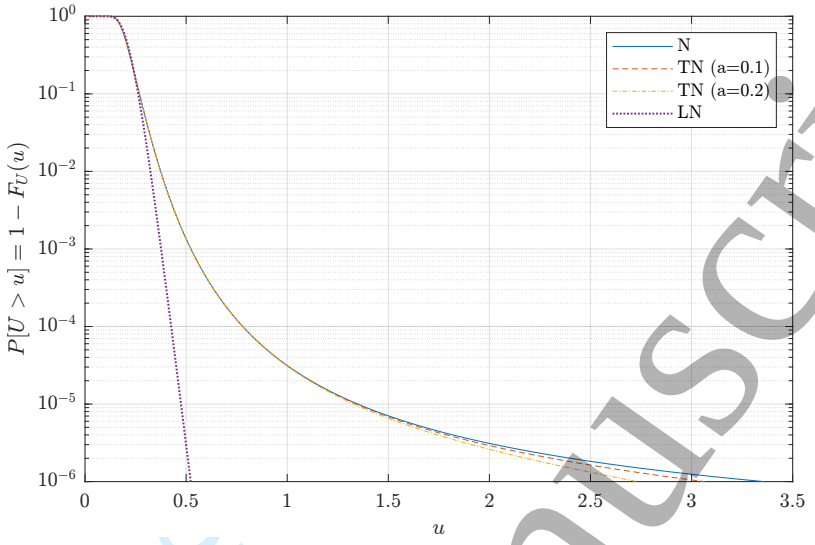

Figure 6. Probability that the displacement U exceeds a threshold level u. The random variable models considered for the stiffness are normal (N), truncated normal (TN), and lognormal (LN) distributions.

of the primary purposes of structural reliability, which is to provide a rational tool for decision-making under uncertainty. Indeed, while an overly conservative estimate of a failure probability would be on the safe side, it could nevertheless be harmful when weighted against, e.g. construction costs of a system.

3. Interval analysis via sampling

3.1. Interval analysis

As mentioned in the Introduction, the second part of this paper deals with a common pitfall in interval analysis, namely the use of sampling(-like) schemes to calculate the bounds of a response u_1 of interest. The main goal of performing interval analysis is to examine the full potential range of values that u_i might assume, while accounting for the epistemic uncertainty that is present in the model input parameters $\boldsymbol{\theta}$. Since we are looking for extremes in the response, it is important to ensure that our estimates are conservative as to avoid making unsafe decisions.

3.1.1. Formal definitions In interval analysis, we consider the epistemic uncertainty in a single parameter θ to be bounded by an interval scalar $\theta^I \subset \mathbb{R}$, which is defined as

$$\theta^I = \left[\underline{\theta},\ \overline{\theta}\right] = \left\{\theta \in \mathbb{R} | \underline{\theta} \leq \theta \leq \overline{\theta}\right\}, \tag{35}$$

where $\underline{\theta}$ and $\overline{\theta}$, with $\underline{\theta} \leq \overline{\theta}$, are bounds between which the unknown values of the uncertain parameter θ are deemed to lie. Similarly, when multiple parameters are

jointly uncertain, an interval vector $\boldsymbol{\theta}^I \subset \mathbb{R}^{d_{n_\theta}}$ is defined by the Cartesian product of n_θ interval scalars: $\boldsymbol{\theta}^I = \theta_1^I \times \theta_2^I \times \ldots \times \theta_{n_\theta}^I$, with $\times$ denoting the Cartesian product operator. As a result, the interval scalars $\theta_i^I,\ i = 1, \ldots, n_\theta$ are independent and as such describe a hyper-rectangular polytope in $\mathbb{R}^{n_\theta}$. The main point is that we neither have sufficient information to define a precise value for θ, nor to characterize a *crisp* probability distribution $F_\theta(\theta)$. The bounds as such represent our 'honest' worst-case estimation of the values the parameter θ could take. Intervals and interval vectors can be fit on data using methods such as scenario optimization [44], Bayesian extreme value methods [45], or Chebyshev's inequality [46].

In the context of engineering analysis, $\boldsymbol{\theta}^I$ effectively represents an n_θ dimensional hyper-rectangle describing the epistemic uncertainty we have as analysts on the *true* value of $\boldsymbol{\theta}$. The main goal of interval analysis is in this context to evaluate the worst and best possible behaviour of Eq. (1)-(2), given the fact that we are not able to exactly quantify $\boldsymbol{\theta}$. One particular way to describe the results of an interval analysis is to consider that the hyper-cube $\boldsymbol{\theta}^I$ is processed through a potentially nonlinear map to a set of possible responses $\mathfrak{U}$:

$$\mathfrak{U} = \left\{ \boldsymbol{u} \mid \mathcal{N}_x(\boldsymbol{u}; \boldsymbol{\theta}) = \boldsymbol{f}, \boldsymbol{x} \in D, \mathcal{B}_x(\boldsymbol{u}; \boldsymbol{\theta}) = \boldsymbol{b}, \boldsymbol{x} \in \Gamma, \boldsymbol{\theta} \in \boldsymbol{\theta}^I \right\}. \tag{36}$$

3.1.2. Propagation of intervals This set $\mathfrak{U}$ effectively contains all possible physical responses of the system that are consistent with the description of the epistemic uncertainty. In other words, it describes how the system under consideration could potentially react in correspondence with our lack of knowledge. In engineering decision-making, we are usually interested in the worst- and best-case behavior of the system. The main issue here is that finding the exact set $\mathfrak{U}$ is computationally intractable within finite time. This means that we want to find $\underline{u}_i$ and $\overline{u}_i$, with $\underline{u}_i \leq \overline{u}_i$ for every response $i, i = 1, ..., n_u$ that is compatible with both Eq. (1)-(2) and $\boldsymbol{\theta}^I$:

$$\begin{aligned} \underline{u}_i &= \min_{\boldsymbol{\theta} \in \boldsymbol{\theta}^I} \mathfrak{U}(i), \quad i = 1, ..., n_u \\ \overline{u}_i &= \max_{\boldsymbol{\theta} \in \boldsymbol{\theta}^I} \mathfrak{U}(i), \quad i = 1, ..., n_u, \end{aligned} \tag{37}$$

where the notation $\mathfrak{U}(i)$ is used to denote that we consider only the i^{th} response of the solution set.

All interval propagation methodologies discussed in the Introduction either deal with solving the min-max optimization problem formulated in Eq. (37) directly, or aim at effectively providing a minimum-encompassing (convex) representation of $\mathfrak{U}$ by explicitly considering Eq. (36). As discussed in the Introduction, these methods include techniques based on optimization [15, 16, 17], perturbation analysis [18, 19, 20], interval arithmetic [47, 21], affine arithmetic [22], improved interval analysis [48, 24], surrogate modeling [25, 26] and Bayesian cubature [27].

Also, sampling-based optimization methods, such as sequential Monte Carlo [49] can safely be used to solve Eq.(37).

3.1.3. Sampling from intervals Since we are interested in the bounds of the response of the structure, the intervals should be propagated in a conservative way. This means that our estimated bounds should be at least as wide as the true bounds from a numerical standpoint. Despite the wide availability of approaches to propagate intervals in such a conservative way, many researchers seem to follow a sampling-based approach that roughly looks like this schemata‡:

(i) Define the interval uncertainty as $\boldsymbol{\theta}^I = \theta_1^I \times \theta_2^I \times \ldots \times \theta_{n_\theta}^I$

(ii) Represent $\boldsymbol{\theta}^I$ as an n_θ-dimensional random variable $\tilde{\boldsymbol{\theta}}$ following a uniform distribution $\mathcal{U}_{n_\theta}\left(\underline{\boldsymbol{\theta}}, \overline{\boldsymbol{\theta}}\right)$

(iii) Generate a large space-filling design containing N samples $\left\{\boldsymbol{\theta}^j, \boldsymbol{u}_i^j\right\}$, with $j = 1, ..., N$, based on uniformly distributed random variables and compute the corresponding responses.

(iv) Determine the bounds $\underline{u}_i^s$ and $\overline{u}_i^s$ based on two sampling estimators:

$$\begin{aligned} \underline{u}_i^s &= \min_{j=1,...,N} u_i^j, \quad i = 1, ..., n_u \\ \overline{u}_i^s &= \max_{j=1,...,N} u_i^j, \quad i = 1, ..., n_u. \end{aligned} \tag{38}$$

A rough inspection of this approach reveals two main assumptions underlying it. First, authors resorting to this approach assume that the uniform distribution is a good tool to represent the uncertainty that is actually present in $\boldsymbol{\theta}$. Second, it is assumed that when N is taken sufficiently large, $\underline{u}_i^s \approx \underline{u}_i$, $\forall i = 1, ..., n_u$ and $\overline{u}_i^s \approx \overline{u}_i$, $\forall i = 1, ..., n_u$ with high accuracy. As we will show in the following subsections, neither assumption can be proven to hold in general, as there are both practical and theoretical objections.

As a small final comment, we note that some authors do not assume a uniform distribution here, but rather *discretise* the interval in equidistant sampling points. Even though this approach is similar to the sampling scheme delineated here before, subtle yet important differences exist with respect to sampling schemes. In essence, there exists no theoretical objections against this method as no probabilistic information is formally imposed on the uncertain parameter θ. Therefore, the interval paradigm is not violated. However, still, from a practical and computational viewpoint, such approach is quite sub-optimal, as will be explained in Section 3.3.

‡ We do not cite papers containing these pitfalls on purpose. Our intention is to be educative, not provocative.

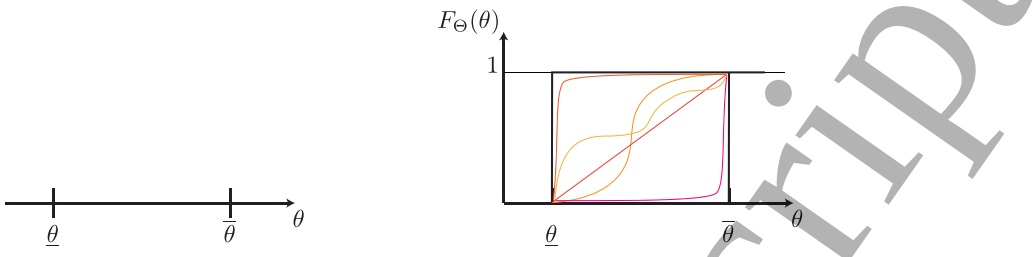

Figure 7. Graphical representations of intervals. Left: an interval as a bounded segment of the real line. Right: an interval as a p-box.

3.1.4. The special case of dependent intervals In case of dependent intervals, the hyper-rectangle $\boldsymbol{\theta}^I$ is tightened either into an admissible set [50], parallelepiped [51], or (hyper)-ellipsoid [52] model. In case of spatial or time dependence, also interval processes [18] and -fields [53, 54] have been introduced. In all these models of dependent intervals, this can be recast into a series of (linear)-inequalities via the Minkowsky-Weyl theorem [50], potentially after a discretisation step. Since the analyst still has no information about the exact value of $\boldsymbol{\theta}$; all we know is that the exact value is bounded by some hyper-planes as described by Minkowski-Weyl. As such, the optimization problem that was introduced in Eq. (37) can be recast into a very similar constrained optimization problem. Therefore, since optimization still lies at the core, the same arguments that were laid out in subsections 3.1.2 and 3.1.3 also hold in case dependence between the interval scalars is introduced in the analysis.

3.2. Theoretical objections to sampling from intervals

It appears to the authors that many of the sampling-based approaches stem from a fundamental misunderstanding of what an interval represents. Indeed, it is very tempting to resort to the assumption that the uncertainty in $\boldsymbol{\theta}$, when it is represented by interval bounds $\underline{\boldsymbol{\theta}}$ and $\overline{\boldsymbol{\theta}}$, can just as well be described by a uniform distribution. After all, also $\mathcal{U}_{n_\theta}\left(\underline{\boldsymbol{\theta}}, \overline{\boldsymbol{\theta}}\right)$ is defined by a lower and upper bound on $\boldsymbol{\theta}$, and a straightforward application of the maximum entropy principle based on just bounds yields a uniform distribution. However, such an assumption grossly neglects the original idea of interval analysis, namely that we do not have enough knowledge to specify anything else than the bounds $\underline{\boldsymbol{\theta}}$ and $\overline{\boldsymbol{\theta}}$. Indeed, when we do assume $\boldsymbol{\theta}$ to follow a uniform distribution, we include extra information in the analysis that was not there to begin with; namely, that each $\boldsymbol{\theta}_i \in \boldsymbol{\theta}^I$ is equally likely to occur. This knowledge is not available from the data we have and is, therefore, inherently subjective and potentially wrong. In essence, there are two compatible representations of the lack of knowledge on the uncertain-but-bounded uncertainty in $\boldsymbol{\theta}$. These are illustrated in Figure 7.

The first interpretation, as shown on the left-hand side of Figure 7, shows θ^I as a bounded segment of the real line $\mathbb{R}$. In essence, this corresponds to stating that one has no knowledge about the real value of θ, other than that it is bounded. Note that one does not make a statement here about the *true* nature of θ, as the parameter modeled by θ might be a deterministic value, but also an aleatory uncertain quantity where we do not have sufficient knowledge to build a probabilistic model. Note furthermore that by specifying θ^I, one does not imply that $\underline{\theta}$ and $\overline{\theta}$ are perfectly bounding all possible realizations of θ in real life. Indeed, several works have explored the reliability of such bounds (see e.g, [45], [30], or [44]). It is noteworthy that, depending on the nature of θ, determining such bounds might be impossible (due to the rarity of finding the corresponding samples experimentally), or undesirable (since they might be extremely wide, to the point that the analysis becomes uninformative).

The second possible interpretation of an interval, as shown on the right-hand side of Figure 7, can be explained when resorting to the theory of p-boxes [3, 4]. The main idea of a p-box is that there exists an unknown CDF F_Θ of the uncertain quantity θ for which only bounds can be provided. Thus, a p-box is described by a lower CDF $\underline{F}_\Theta \in \mathbb{F}$ and an upper CDF $\overline{F}_\Theta \in \mathbb{F}$, where $\mathbb{F}$ expresses the set of all CDFs on $D_\Theta \subseteq \mathbb{R}$. These CDFs are collected as a pair $\left[\underline{F}_\Theta, \overline{F}_\Theta\right]$ which yields a set of possible CDFs $\{F_\Theta \in \mathbb{F} \mid \underline{F}_\Theta(\theta) \leq F_\Theta(\theta) \leq \overline{F}_\Theta(\theta),\ \theta \in D_\Theta\}$ for the unknown CDF of Θ. The definition of a p-box corresponds to defining a lower probability $\underline{P}$ and upper probability $\overline{P}$ on events $\{\Theta \leq \theta\} = (-\infty, \theta] \cap D_\Theta$, i.e., $\underline{P}(\Theta \leq \theta) = \underline{F}_\Theta(\theta)$ and $\overline{P}(\Theta \leq \theta) = \overline{F}_\Theta(\theta)$ for $\theta \in D_\Theta$, which yields a credal set of probability measures. Now, in case we have absolute uncertainty about the real value of θ, we might state that it could belong to every possible distribution with support $D_\Theta = [\underline{\theta}, \overline{\theta}]$. This interpretation covers both the scenario where we have a lack of knowledge about a deterministic quantity θ and the situation where θ has a random nature which is elusive due to the lack of sufficient data. Note that the p-box has a fundamentally different interpretation in both cases, and also here care should be taken on how the analysis ensues. Such a p-box corresponds to defining the lower CDF and upper CDF, respectively, as $\overline{F}_\Theta = H(\theta - \underline{\theta})$ and $\underline{F}_\Theta = H(\theta - \overline{\theta})$, with $H(\bullet)$ the Heaviside function, as also illustrated on the right-hand side of Figure 7.

It can be observed that both interpretations contain some sort of duality: one can specify that there is a lack of knowledge of the *true* value of θ within the interval bounds by either providing only the bounds, or treating θ as if it could belong to any possible distribution F_Θ that is bounded on $D_\Theta = \theta^I$. Neither interpretation allows for treating θ as a uniformly distributed random variable.

3.3. Computational objections against sampling from intervals

Next to the theoretical objections raised in Section 3.2, there are very compelling computational arguments not to solve the interval propagation problem using Eq. (38). Just to specify, we criticize the use of *plain* MCS in the context of interval propagation.

3.3.1. Convergence issues Normally, in uncertainty quantification and reliability engineering, MCS is used to solve integral problems, such as those related to computing the probability of failure p_f (see Eq. (3)), which is approximated by using MCS via the estimator $\hat{p}_f$:

$$\hat{p}_f = \frac{1}{N} \sum_{i=1}^{N} I(g(\theta_i)). \tag{39}$$

It is well-known that the number of samples N required to build a $(1-\alpha)$ confidence interval around p_f with an accuracy level ϵ is bounded from below by:

$$N \geq \left(\frac{1}{\epsilon} z_{\alpha/2} \sqrt{p_f(1-p_f)} \right), \tag{40}$$

with $z_{\alpha/2}$ the z-score related to the confidence level α. In practice, ϵ is chosen to be at least $\epsilon \leq 0.1 p_f$ for accuracy reasons. To solve Eq. (38) in a similar way, one can argue that it is required that $g_i(\theta) \to \overline{u}_i$ to determine $\overline{u}_i$ and $g_i(\theta) \to \underline{u}_i$ to determine $\underline{u}_i$, with g_i the performance function only considering the i^{th} response. Indeed, to have an accurate estimator of the response bounds, the limit state surfaces of each response individually must tend towards those bounds. It is easy to see that the corresponding p_f values therefore must tend towards zero. This, in turn, allows the following analysis for the sample size required to compute $\overline{u}_i$, after re-arranging a few terms in Eq. (40):

$$N_{\overline{u}_i} \geq \lim_{p_f \to 0} z_{\alpha/2} \left(\frac{1}{\sqrt{p_f}} \right) = +\infty. \tag{41}$$

This result shows that one indeed needs a sample size $N \to +\infty$ to find the correct value of $\overline{u}_i$, regardless of the corresponding confidence level. As a sidenote, p_f can in this context be viewed as the probability of sampling a point that gives an extreme (near-bound) response, which grows vanishingly small as that point becomes more extreme. Of course, one can build a traceable argument that it is sufficient to find the interval bound with a precision level ϱ. In this case, the equivalent limit state function for the upper bound becomes $g_i(\theta) = \overline{u}_i - \varrho$. The number of required samples, in this case, can be shown to be bounded as:

$$N_{\overline{u}_i} \geq \left(\frac{1}{\epsilon} z_{\alpha/2} \sqrt{\varsigma(1-\varsigma)} \right), \tag{42}$$

with ς the value of ϱ normalized with respect to the width of the output interval u_i^I. Whereas in this case the required set of samples is not infinite, it will still be prohibitively large when one wants to get reasonably accurate estimators of the bounds. Take, for instance, the case where we want to thus find the upper bound $\overline{u}_i$ with a precision of $\varrho = 1 \cdot 10^{-06}$, while having a confidence level of 99% and a precision on the MCS estimator of $\epsilon = 0.1 * \varsigma$. In case we have an interval with unit-width (i.e, $|\overline{u}_i - \underline{u}_i| = 1$), the sample size must be $N_{\overline{u}_i} \geq 6.635 \cdot 10^{08}$.

Furthermore, as these analyses also show, the sample sizes are obtained while letting $g_i(\theta) \to \overline{u}_i$. This means that one irrevocably obtains an inner approximation of the interval boundaries. Indeed, since the problem is solved from the 'inside' towards the bounds of the response, while requiring an infinite sample size to converge, MCS gives inherently an inner interval approximation. This means that no conservatism can be proven. A fully analogous analysis can be made for the calculation of the lower bound $\underline{u}_i$.

Concerning discretisation schemes, the approach boils down to selecting a structured set of propagation points, consisting of both the interval bounds and a selection of inner points from the input space. This method aims to approximate the worst- and best-case responses without requiring a full MCS. While it is computationally more efficient, it fundamentally suffers from the same limitations as MCS in determining strict response bounds. The key issue is that the worst- and best-case responses may not necessarily occur at the selected propagation points. If the extreme responses lie between the chosen inner points due to non-monotonicity in the response function, they will be missed entirely. This results in an inner approximation of the response interval, just as in the MCS approach. To ensure finding the correct bounds, an infinitely fine discretisation would be required. As a final note, we want to acknowledge that this approach has merit in case the analyst is more interested in learning the nature of the mapping between the in- and outputs of the model under consideration (e.g., smoothness, convexity, continuity). Therefore, it is not a pitfall 'per se', but rather a method that should be used with proper consideration and care.

3.3.2. Dimensionality issues The problematic convergence of finding the bounds of the system's response subject to interval uncertainty when using MCS was illustrated in the last paragraphs. In this context, an additional problem shows up. Whereas MCS, when used to determine p_f is insensitive to the dimension of the input vector $\boldsymbol{\theta}$, this is not the case when looking for extremal values in the system's response. Let's just use a very intuitive example to illustrate this. Consider the input to be defined by the hyper-cube $\boldsymbol{\theta} = [0, 1]^{n_\theta}$. In the case of $n_\theta = 1$, when drawing 10 Latin Hypercube samples from $\mathcal{U}_1(0, 1)$, on average each sample will cover 10% of the total space. However, when $n_\theta = 2$, each of those 10 Latin Hypercube samples

will only still cover 1% of the total sample space. In general, when drawing N Latin Hypercube samples over the n_θ dimensional hyper-cube, each of those samples will on average cover a fraction of $(1/N)^{n_\theta}$ of the total hyper-cubic input space. Now, we have to realize that we are trying to find the realizations $\overline{\boldsymbol{\theta}_i}^* \in \boldsymbol{\theta}^I$ and $\underline{\boldsymbol{\theta}_i}^* \in \boldsymbol{\theta}^I$ that solve Eq. (38) for every $i = 1, ..., n_u$. As explained, we can allow for a precision level ς, meaning we are essentially trying to randomly find two ς dimensional hyper-cubes in the n_θ dimensional input space. This means, that to find either of the bounds, we would need, on average, $\frac{1}{\varsigma^{n_\theta}}$ samples, assuming a unit-hypercube. Assuming a problem that we wish to solve with $n_\theta = 20$ input variables and a precision of $\varsigma = 1 \cdot 10^{-06}$, already would cause us to need $1 \cdot 10^{120}$ samples. Just as a comparison to position this order of magnitude: there are roughly speaking $1 \cdot 10^{80}$ atoms in the observable universe.

It should be pointed out that these arguments can be one-on-one transferred to the case of discretising the interval. Indeed, the number of inner points required to sufficiently approximate the true interval bounds, resp. learn how the mapping looks like, grows rapidly with the dimensionality of the problem.

3.4. Practical examples of sampling from intervals

In this section, we present simple examples to illustrate the under-estimation and inner-approximation phenomena, as well as the slow convergence of the Monte Carlo estimator for reaching the bounds of the distribution.

3.4.1. The Ishigami function As a first example, consider the 2-dimensional Ishigami function, which is defined as:

$$u = \sin\theta_1 + 7\sin{\theta_2}^2 + 0.1\theta_3 \sin\theta_1. \tag{43}$$

The values of the parameters $\theta_i, i = 1, 2, 3$ are considered to be unknown, but bounded by the cube $[-\pi, \pi]^3$. The corresponding bounds on u, namely $\underline{u}$ and $\overline{u}$ are found following two approaches:

- assuming a uniform distribution in $[-\pi, \pi]^3$. Following Eq. (38), this produces the estimators $\underline{u}^s$ and $\overline{u}^s$. To study convergence, we increase the sample size N in 500 steps from 10 to $5 \cdot 10^{08}$.
- performing a brute-force particle swarm optimization that is bounded on $[-\pi, \pi]^3$. For both bounds, the built-in `Matlab` tool `particleswarm` requires roughly 1000 samples to converge.

We perform brute-force particle swarm optimization as opposed to the more sophisticated techniques that are available to strengthen our point that even brute-force optimization is more efficient and accurate than sampling between the bounds.

It is pointed out that the application of sophisticated interval propagation techniques such as those based on active learning Bayesian Optimization can still provide several gains of computational efficiency. Assuming the result of the global optimization to be the correct one, we compute the relative error on the lower and upper bound, as obtained by Monte Carlo sampling, as follows:

$$\begin{aligned} \epsilon_l &= \left(\underline{u}^s - \underline{u}\right)/\underline{u} \\ \epsilon_u &= \left(\overline{u} - \overline{u}^s\right)/\overline{u} \end{aligned} \tag{44}$$

Note that these errors are designed such that they become negative when the sampling-based method becomes more 'conservative' than the optimization approach.

Figure 8 illustrates the convergence of the sampling-based interval bound estimator in a relative sense to the global optimization procedure. Observing this figure, a few points need to be made. First, when comparing both propagation schemes for roughly the same computational cost, one can observe that even in this very simple 3-dimensional input case, an error of roughly 10% is made. Second, even when using a sample size as large as $N = 5 \cdot 10^{08}$, an error of roughly 0.01% is made. Both points illustrate the sheer inefficiency of propagating intervals using crude sampling methods, especially when we consider that (1) we study a very well-behaved function here and (2) the number of input parameters is very low. Third, it can be observed that both ϵ_l and ϵ_u are strictly positive. This, by construction, means that the sampling estimators inherently provide an inner estimation of the interval width of u. Indeed, this clearly shows that $\underline{u}^s$ consistently over-estimates the lower bound $\underline{u}$, just as $\overline{u}^s$ consistently under-estimates the upper bound $\overline{u}$. As a final note, we want to point out that the particle swarm optimizer is even easier to implement compared to the sampling approach (even though, admittedly, both are straightforward). As such, we see no reason to use sampling-based methods for the propagation of interval-valued uncertainty.

3.4.2. Rastrigin's function in 20 dimensions To illustrate that this issue is even more pronounced in high-dimensional problems, we test both approaches mentioned in Subsection 3.4.1 on the 20-dimensional Rastrigin function, which is generally defined as

$$u = 10n + \sum_{i=1}^{n} \left(\theta_i^2 - 10\cos(2\pi\theta_i)\right), \tag{45}$$

where we set $n = 20$. Note that Rastrigin's function is known for its many local minima, so searching for bounds is non-trivial. The correct bounds of Rastrigin's function were obtained using the same untuned particle swarm optimization algorithm as used in Subsection 3.4.1 using approximately 15000 samples.

Due to the very high dimensionality of this problem, as mentioned before, the probability of randomly finding a point that is even a good approximation of the

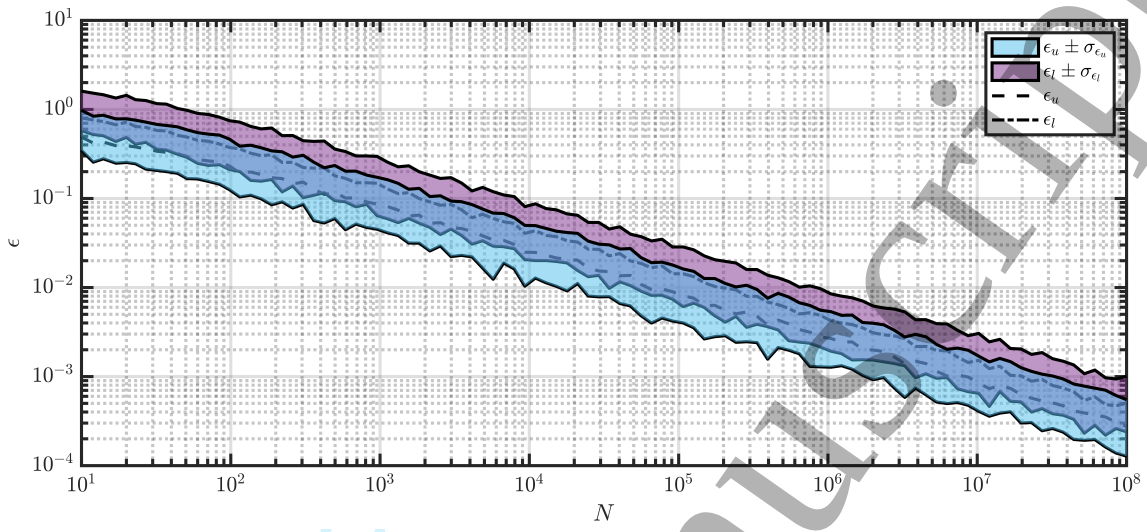

Figure 8. Convergence of ϵ_l and ϵ_u on the extremes of the Ishigami function with respect to the sample size

extreme responses of u drastically decreases, as we are with 10^{08} samples still very far away from the required number of $1 \cdot 10^{120}$. As can be noted, even with 10^{08} samples, the error that is made on the relatively easy-to-find upper bound is still approximately 10%. This indicates that for complicated, high-dimensional examples, sampling-based propagation schemes are a very bad choice for propagating intervals. To be fair, in this particular case, also the global optimization algorithm required a huge computational effort.

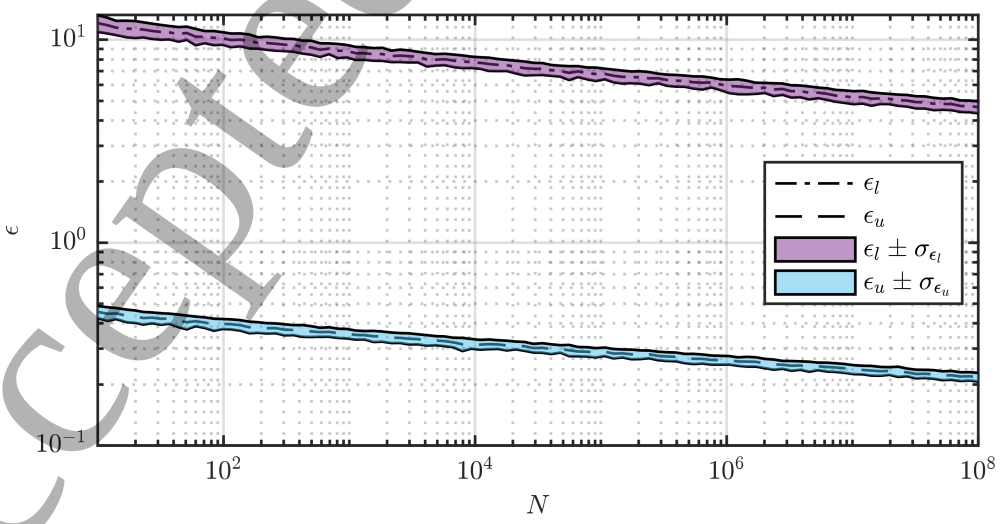

Figure 9. Convergence of ϵ_l and ϵ_u on the extremes of Rastrigins function with respect to the sample size

3.4.3. A two-bar truss structure As a final illustration, we include the analysis of a simple 2-bar truss structure. The structure consists of three nodes: Node 1 (location: $0, 0$ m) and Node 2 (location: $2, 0$ m) are fully fixed (both translational degrees of freedom restrained), and Node 3 (location: $0, 2$ m) is free. Two truss elements connect Node 3 to Nodes 1 and 2, respectively, forming a simple triangular arrangement. At Node 3, a vertical load is applied. Interval uncertainty is assumed on the load F, Young's modulus of the material E and the truss cross-sectional area A. These intervals are considered as $F^I = [1, 5]$ kN, $E^I = [200, 220]$ GPa and $A^I = [0.05, 0.15]\ m^2$. The quantity of interest is the vertical displacement of Node 3, u.

Since this is a small-displacement-based, linear finite element model, the relation between the inputs and the displacement of Node 3 is purely monotonic. By using engineering judgment, the exact bounds can be calculated using as few as 2 finite element model simulations. Indeed, $\overline{u}$ is related to the triplet $\left\{\overline{F}, \underline{E}, \underline{A}\right\}$, and was found to be $\overline{u} = 0,7$ mm. Conversely, $\underline{u}$ is related to $\left\{\underline{F}, \overline{E}, \overline{A}\right\}$, and was found to be $\overline{u} = 0,04$ mm. In addition, sampling is used to approximate $\underline{u}$ and $\overline{u}$ as resp. $\underline{u}^s$ and $\overline{u}^s$. These results are used in the calculation of the error ϵ, as illustrated in Eq. (44).

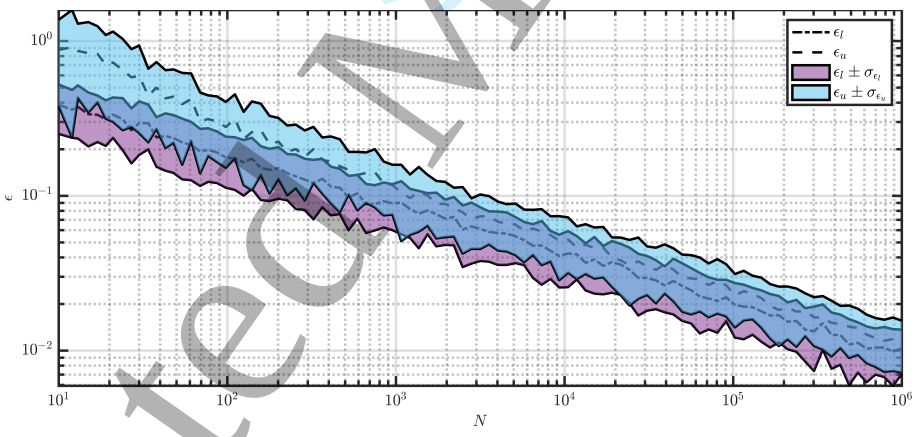

Figure 10. Convergence of ϵ_l and ϵ_u of the extreme truss responses with respect to the sample size

Figure 10 illustrates the convergence of ϵ_l and ϵ_u with respect to the sample size N when $\underline{u}$ and $\overline{u}$ are approximated by sampling as resp. $\underline{u}^s$ and $\overline{u}^s$. As can be seen, a sample size of roughly $N = 1 \cdot 10^{06}$ is required to achieve a relative error of 1%. At first sight, one may be tempted to say that this is acceptable. However, it should be kept in mind that by using proper interval techniques, an exact solution can be achieved in as little as 2 simulations.

4. Conclusions

The primary objective in uncertainty quantification, whether using probabilistic or non-traditional approaches, is to accurately capture and represent uncertainty. The ultimate goal of such practice is to help analysts and engineers make risk-informed decisions. Uncertainty quantification involves several key steps, which include data collection, mathematical modeling, uncertainty propagation, identifying critical outcomes, and making decisions under uncertainty. Each step introduces potential challenges and possible sources of error that should be minimized or eliminated.

In this study, we examine two common pitfalls in uncertainty quantification. The first issue concerns the assumption of normality for parameters that must be positive due to physical constraints. This pitfall is often justified by arguments such as the maximum entropy principle, the central limit theorem, or the fact that the probability of observing negative numbers is negligible. We illustrate that using a normal distribution can lead to outputs with undefined second-order statistics and overly conservative estimates. In other words, if a positive quantity is modeled with a normal distribution, the obtained results will be (at least) questionable or (plainly) wrong, even if the probability content associated to negative values is extremely small. While the truncated normal distribution can mitigate some of these issues, it remains highly sensitive to the truncation parameter. Hence, truncation of the negative values associated with a normal distribution which is used to model positive-valued properties may not suffice to ensure correctness and in fact, it may lead to gross errors. Consequently, modeling approaches should respect the inherent characteristics of the parameters involved to avoid misleading or invalid results. In a nutshell, the recommendations for modeling the uncertainty associated with positive-valued properties of a model are: (1) avoid the normal distribution and (2) use a truncated normal distribution only after investigating the effect of the truncation parameter.

The second issue, analyzed in this paper, is the inappropriate use of sampling-based methods, such as Monte Carlo simulation from a uniform distribution, for propagating interval uncertainties. While these methods are powerful for probabilistic analysis, their direct application to interval analysis is problematic. Due to their design, they impose unjustified uniform assumptions, leading to computationally expensive and inherently non-conservative approximations of the true response bounds. The key takeaway here is that rigorous interval analysis should rely on methods specifically designed for bounding problems, such as direct optimization, interval arithmetic, affine arithmetic, or advanced surrogate-assisted techniques. Otherwise, one risks conducting an expensive analysis only to obtain an interior approximation of the uncertainty.

Taken together, the findings of both parts of this paper highlight the need for a

critical evaluation of standard practices in uncertainty quantification. To ensure reliable results, it is essential to align modeling choices with the nature of the input data and apply the correct mathematical framework to represent uncertainty properly. By avoiding fundamental errors at each stage of the process, we can develop practical and robust decision-support tools for engineering applications.

References

[1] A. Tabandeh, G. Jia, and P. Gardoni, "A review and assessment of importance sampling methods for reliability analysis," *Struct. Saf.*, vol. 97, p. 102216, July 2022.

[2] M. Faes and D. Moens, "Recent trends in the modeling and quantification of non-probabilistic uncertainty," *Arch. Comput. Methods Eng.*, vol. 27, p. 633–671, July 2020.

[3] M. Beer, S. Ferson, and V. Kreinovich, "Imprecise probabilities in engineering analyses," *Mech. Syst. Sig. Process.*, vol. 37, no. 1–2, p. 4–29, 2013.

[4] M. Faes, M. Daub, S. Marelli, E. Patelli, and M. Beer, "Engineering analysis with probability boxes: A review on computational methods," *Struct. Saf.*, vol. 93, p. 102092, Nov. 2021.

[5] C. Bucher, *Computational Analysis of Randomness in Structural Mechanics*. CRC Press, 2009.

[6] O. Ditlevsen and H. O. Madsen, *Structural Reliability Methods*. John Wiley and Sons, 1996.

[7] R. E. Melchers and A. T. Beck, *Structural Reliability Analysis and Prediction*. John Wiley & Sons, 3rd ed., 2018.

[8] K. Breitung, "The return of the design points," *Reliab. Eng. Syst. Saf.*, vol. 247, p. 110103, July 2024.

[9] H. J. Pradlwarter, G. I. Schuëller, P. S. Koutsourelakis, and D. C. Charmpis, "Application of line sampling simulation method to reliability benchmark problems," *Struct. Saf.*, vol. 29, no. 3, p. 208–221, 2007.

[10] C. Dang, M. Valdebenito, M. Faes, J. Song, P. Wei, and M. Beer, "Structural reliability analysis by line sampling: A bayesian active learning treatment," *Struct. Saf.*, vol. 104, p. 102351, Sept. 2023.

[11] S. K. Au and J. L. Beck, "Estimation of small failure probabilities in high dimensions by subset simulation," *Probab. Eng. Mech.*, vol. 16, no. 4, p. 263–277, 2001.

[12] M. Rashki, "Sesc: A new subset simulation method for rare-events estimation," *Mech. Syst. Signal Process.*, vol. 150, p. 107139, Mar. 2021.

[13] C. Dang, M. Valdebenito, M. Faes, P. Wei, and M. Beer, "Structural reliability analysis: A Bayesian perspective," *Struct. Saf.*, vol. 99, p. 102259, Nov. 2022.

[14] M. A. Valdebenito, H. A. Jensen, G. I. Schuëller, and F. E. Caro, "Reliability sensitivity estimation of linear systems under stochastic excitation," *Comput. Struct.*, vol. 92–93, p. 257–268, 2012.

[15] F. Tonon and A. Bernardini, "A random set approach to the optimization of uncertain structures," *Comput. Struct.*, vol. 68, p. 583–600, Sept. 1998.

[16] Z. Qiu and I. Elishakoff, "Anti-optimization technique - a generalization of interval analysis for nonprobabilistic treatment of uncertainty," *Chaos, Solitons & Fractals*, vol. 12, no. 9, p. 1747–1759, 2001.

[17] D. Moens and D. Vandepitte, "An interval finite element approach for the calculation of envelope frequency response functions," *Int. J. Numer. Methods Eng.*, vol. 61, no. 14, p. 2480–2507, 2004.

[18] C. Jiang, B. Ni, N. Liu, X. Han, and J. Liu, "Interval process model and non-random vibration analysis," *J. Sound Vib.*, vol. 373, p. 104–131, July 2016.

[19] B. Wu, D. Wu, W. Gao, and C. Song, "Time-variant random interval natural frequency analysis of structures," *J. Sound Vib.*, vol. 414, p. 284–298, 2018.

[20] B. Ni and C. Jiang, "Interval field model and interval finite element analysis," *Comput. Methods Appl. Mech. Eng.*, vol. 360, p. 112713, Mar. 2020.

[21] H. Leng, "Real eigenvalue bounds of standard and generalized real interval eigenvalue problems," *Appl. Math. Comput.*, vol. 232, p. 164–171, 2014.

[22] D. Degrauwe, G. Lombaert, and G. D. Roeck, "Improving interval analysis in finite element calculations by means of affine arithmetic," *Comput. Struct.*, vol. 88, no. 3–4, p. 247–254, 2010.

[23] G. Muscolino and A. Sofi, "Stochastic analysis of structures with uncertain-but-bounded parameters via improved interval analysis," *Probab. Eng. Mech.*, vol. 28, pp. 152–163, 2012.

[24] A. Sofi, E. Romeo, O. Barrera, and A. Cocks, "An interval finite element method for the analysis of structures with spatially varying uncertainties," *Adv. Eng. Software*, vol. 128, p. 1–19, 2019.

[25] T. Huang, X. Song, and M. Liu, "A kriging-based non-probability interval optimization of loading path in t-shape tube hydroforming," *Int. J. Adv. Manuf. Technol.*, vol. 85, p. 1615–1631, July 2016.

[26] A. Sofi and E. Romeo, "A unified response surface framework for the interval and stochastic finite element analysis of structures with uncertain parameters," *Probab. Eng. Mech.*, vol. 54, p. 25–36, 2018.

[27] C. Dang, P. Wei, M. Faes, M. Valdebenito, and M. Beer, "Interval uncertainty propagation by a parallel bayesian global optimization method," *Appl. Math. Model.*, vol. 108, p. 220–235, Aug. 2022.

[28] V. Kreinovich and S. Ferson, "A new cauchy-based black-box technique for uncertainty in risk analysis," *Reliab. Eng. Syst. Saf.*, vol. 85, p. 267–279, July 2004.

[29] R. R. Callens, M. G. Faes, and D. Moens, "Multilevel quasi-Monte Carlo for interval analysis," *Int. J. Uncertain. Quantif.*, vol. 12, no. 4, p. 1–19, 2022.

[30] R. Callens, D. Moens, and M. Faes, "Certified interval model updating using scenario optimization," *AIAA J.*, 2025, accepted.

[31] A. Beck and C. Ávila da Silva, "Strategies for finding the design point under bounded random variables," *Struct. Saf.*, vol. 58, pp. 79–93, 2016.

[32] E. Rosenblueth, Karmeshu, and H. Hong, "Maximum entropy and discretization of probability distributions," *Probab. Eng. Mech.*, vol. 2, no. 2, pp. 58–63, 1987.

[33] D. Ormoneit and H. White, "An efficient algorithm to compute maximum entropy densities," *Econom. Rev.*, vol. 18, pp. 127–140, 1999.

[34] A. Der Kiureghian, *Engineering Design Reliability Handbook*, ch. First- and Second-Order Reliability Methods. CRC Press, 2004.

[35] I. Babuška and P. Chatzipantelidis, "On solving elliptic stochastic partial differential equations," *Comput. Methods Appl. Mech. Eng.*, vol. 191, no. 37, pp. 4093–4122, 2002.

[36] I. Babuška, R. Tempone, and G. Zouraris, "Solving elliptic boundary value problems with uncertain coefficients by the finite element method: the stochastic formulation," *Comput. Methods Appl. Mech. Eng.*, vol. 194, no. 12, pp. 1251 – 1294, 2005. Special Issue on Computational Methods in Stochastic Mechanics and Reliability Analysis.

[37] C. da Silva Jr. and A. Beck, "Bending of stochastic Kirchhoff plates on Winkler foundations via the Galerkin method and the Askey-Wiener scheme," *Probab. Eng. Mech.*, vol. 25, no. 2, pp. 172–182, 2010.

[38] C. Ávila da Silva and A. Beck, "Chaos–Galerkin solution of stochastic Timoshenko bending

problems," *Comput. Struct.*, vol. 89, no. 7-8, pp. 599–611, 2011.

[39] F. Yamazaki, M. Shinozuka, and G. Dasgupta, "Neumann expansion for stochastic finite element analysis," *J. Eng. Mech.*, vol. 114, no. 8, pp. 1335–1354, 1988.

[40] X.-Y. Zhang, Y.-G. Zhao, and Z.-H. Lu, "Unified Hermite polynomial model and its application in estimating non-gaussian processes," *J. Eng. Mech.*, vol. 145, no. 3, p. 04019001, 2019.

[41] M. Schevenels, G. Lombaert, and G. Degrande, "Application of the stochastic finite element method for Gaussian and non-Gaussian systems," in *ISMA2004 International Conference on Noise and Vibration Engineering*, pp. 3299–3314, 2004.

[42] W. Piegorsch and G. Casella, "The existence of the first negative moment," *Am. Stat.*, vol. 39, no. 1, pp. 60–62, 1985.

[43] E. Lehmann and J. Shaffer, "Inverted distributions," *Selected Works of E.L. Lehmann*, pp. 833–836, 2012.

[44] L. G. Crespo, S. P. Kenny, and D. P. Giesy, "Interval predictor models with a linear parameter dependency," *J. Verif. Valid. Uncertain. Quantif.*, vol. 1, no. 2, p. 1–10, 2016.

[45] M. Imholz, M. Faes, D. Vandepitte, and D. Moens, "Robust uncertainty quantification in structural dynamics under scarse experimental modal data: A Bayesian-interval approach," *J. Sound Vib.*, vol. 467, p. 114983, Feb. 2020.

[46] I. Elishakoff and N. Sarlin, "Uncertainty quantification based on pillars of experiment, theory, and computation. part ii: Theory and computation," *Mech. Syst. Signal Process.2*, vol. 74, 2016.

[47] R. Moore, *Interval arithmetic and automatic error analysis in digital computing.* PhD thesis, Ph. D. Dissertation, Department of Mathematics, Stanford University, 1962.

[48] G. Muscolino and A. Sofi, "Bounds for the stationary stochastic response of truss structures with uncertain-but-bounded parameters," *Mech. Syst. Sig. Process.*, vol. 37, no. 1–2, p. 163–181, 2013.

[49] P. Del Moral, A. Doucet, and A. Jasra, "Sequential Monte Carlo samplers," *J. R. Stat. Soc. Ser. B Stat. Methodol.*, vol. 68, p. 411–436, June 2006.

[50] M. Faes and D. Moens, "Multivariate dependent interval finite element analysis via convex hull pair constructions and the extended transformation method," *Comput. Methods Appl. Mech. Eng.*, vol. 347, p. 85–102, Apr. 2019.

[51] C. Jiang, Q. F. Zhang, X. Han, J. Liu, and D. A. Hu, "Multidimensional parallelepiped model – a new type of non-probabilistic convex model for structural uncertainty analysis," *Int. J. Numer. Methods Eng.*, vol. 103, no. 1, p. 31–59, 2015.

[52] I. Elishakoff and F. Elletro, "Interval, ellipsoidal, and super-ellipsoidal calculi for experimental and theoretical treatment of uncertainty: Which one ought to be preferred?," *Int. J. Solids Struct.*, vol. 51, no. 7–8, p. 1576–1586, 2014.

[53] A. Sofi, "Structural response variability under spatially dependent uncertainty: Stochastic versus interval model," *Probab. Eng. Mech.*, vol. 42, pp. 78–86, 2015.

[54] W. Verhaeghe, W. Desmet, D. Vandepitte, and D. Moens, "Interval fields to represent uncertainty on the output side of a static FE analysis," *Comput. Methods Appl. Mech. Eng.*, vol. 260, no. 0, pp. 50–62, 2013. ISBN: 0045-7825 Publisher: Elsevier B.V.

Matthias Faes is full professor at the Chair for Reliability Engineering of TU Dortmund, Germany.

Alba Sofi is associate professor of mechanics of solids and structures at the University "Mediterranea" of Reggio Calabria, Italy.

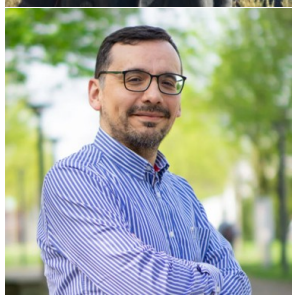

Marcos Valdebenito is chief engineer at the Chair for Reliability Engineering of TU Dortmund, Germany.