Reliability Engineering and System Safety 165 (2017) 376–394

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

Efficient Monte Carlo methods for estimating failure probabilities

Andres Alban[a], Hardik A. Darji[a,b], Atsuki Imamura[c], Marvin K. Nakayama[c,*]

[a] Mathematical Sciences Dept., New Jersey Institute of Technology, Newark, NJ 07102, USA
[b] Mechanical Engineering Dept., New Jersey Institute of Technology, Newark, NJ 07102, USA
[c] Computer Science Dept., New Jersey Institute of Technology, Newark, NJ 07102, USA

ARTICLE INFO

Keywords:
Probabilistic safety assessment
Risk analysis
Structural reliability
Uncertainty
Monte Carlo
Variance reduction
Confidence intervals
Nuclear regulation
Risk-informed safety-margin characterization

ABSTRACT

We develop efficient Monte Carlo methods for estimating the failure probability of a system. An example of the problem comes from an approach for probabilistic safety assessment of nuclear power plants known as risk-informed safety-margin characterization, but it also arises in other contexts, e.g., structural reliability, catastrophe modeling, and finance. We estimate the failure probability using different combinations of simulation methodologies, including stratified sampling (SS), (replicated) Latin hypercube sampling (LHS), and conditional Monte Carlo (CMC). We prove theorems establishing that the combination SS+LHS (resp., SS +CMC+LHS) has smaller asymptotic variance than SS (resp., SS+LHS). We also devise asymptotically valid (as the overall sample size grows large) upper confidence bounds for the failure probability for the methods considered. The confidence bounds may be employed to perform an asymptotically valid probabilistic safety assessment. We present numerical results demonstrating that the combination SS+CMC+LHS can result in substantial variance reductions compared to stratified sampling alone.

1. Introduction

Consider a stochastic model of the behavior of a system (e.g., a structure, such as a building or ship). The system's uncertainties may include random loads and capacities, environmental conditions, material properties, etc. The system fails under specified conditions—e.g., a critical subset of components fails—and the goal is to determine if the failure probability θ is acceptably small. As the system's complexity renders computing θ as intractable, we instead apply Monte Carlo simulation to estimate θ. To account for the sampling error of the simulation estimates, a confidence interval for θ is also needed. Running the simulation model may be expensive, so we want to reduce the sampling error. In this paper, we combine different variance-reduction techniques (VRTs) to estimate θ.

A motivating example comes from a framework for probabilistic safety assessments (PSAs) of nuclear power plants (NPPs) known as *risk-informed safety-margin characterization* (RISMC), which was proposed by an international effort of the Nuclear Energy Agency Committee on the Safety of Nuclear Installations [1]. The purpose of RISMC is to address recent changes in NPPs. For example, many NPPs in the current U.S. fleet are aging past their original 40-year operating licenses, with owners applying for lengthy extensions [2]. Also, for economic reasons, plants are sometimes run at higher output levels, known as power uprates [3]. These and other factors can lead to degradations in safety margins previously deemed acceptable, and RISMC aims to better understand their impacts.

Current NPP PSAs compare a random load to a *fixed* capacity. For example, the U.S. Nuclear Regulatory Commission (NRC) [4, paragraph 50.46(b)(1)], specifies that the 0.95-quantile of the peak cladding temperature (PCT) during a hypothesized loss-of-coolant accident (LOCA) must not exceed 2200° F. The NRC permits a plant licensee to demonstrate its facility's compliance with federal regulations using a 95/95 analysis with Monte Carlo simulation; see Section 24.9 of [5] and [6]. This entails running a computer code [7] modeling the postulated event multiple times with randomly generated inputs to explore a range of uncertainties [8], using the code's outputs to construct a 95% *upper confidence bound* (UCB) for the 0.95-quantile of the PCT, and verifying that the UCB lies below the fixed threshold 2200° F. The UCB accounts for the statistical error of the Monte Carlo estimate due to sampling variability, and the difference between the fixed threshold and the UCB provides a type of safety margin.

Several papers have developed Monte Carlo methods for performing a 95/95 analysis. For example, [9] and [10] apply an approach of [11] based on order statistics, which is valid when employing *simple random sampling* (SRS). However, SRS can produce unusably noisy estimates of the 0.95-quantile, so [12] incorporates VRTs, including antithetic variates (AV) and Latin hypercube sampling (LHS), to obtain more statistically efficient quantile estimators; see Chapter 4 of [13]

* Corresponding author.
E-mail addresses: aa655@njit.edu (A. Alban), had8@njit.edu (H.A. Darji), ai64@njit.edu (A. Imamura), marvin@njit.edu (M.K. Nakayama).

http://dx.doi.org/10.1016/j.ress.2017.04.001
Received 25 July 2016; Received in revised form 16 March 2017; Accepted 11 April 2017
Available online 18 April 2017
0951-8320/

and Chapter V of [14] for overviews of these and other VRTs for estimating a mean. In addition, [12] provides a UCB for the quantile using a finite-difference approach of [15] and [16]. Utilizing VRTs in nuclear PSAs is especially important because each simulation run may be computationally expensive, e.g., it may require numerically solving systems of differential equations.

In addition to the condition on the PCT, the NRC further imposes requirements on other *criteria*—the core-wide oxidation (CWO<1%) and maximum local oxidation (MLO<17%)—for a PSA of a hypothesized NPP LOCA; see paragraph 50.46(b) of [4]. (The NRC is currently considering replacing MLO with another criterion, the integral time at temperature [17].) The papers [9] and [10] describe approaches to extend the 95/95 analysis based on an SRS quantile estimator for a single criterion to handle multiple criteria with *fixed* capacities.

RISMC differs from a 95/95 analysis in several important ways. First, RISMC assumes each criterion's capacity is *random* rather than fixed. Moreover, instead of examining a quantile, as in a 95/95 study, RISMC measures a safety margin through the *failure probability* θ that any random load exceeds its random capacity, where θ should be smaller than a given threshold θ_0, which may be specified by a regulator. RISMC also decomposes a postulated event into scenarios via an *event tree* [18], as in a probabilistic risk analysis (PRA), with uncertainties in each scenario, and Monte Carlo is employed to estimate each scenario's failure probability. To account for the statistical sampling error of the Monte Carlo estimate of the failure probability θ, one should further provide a UCB (or a two-sided confidence interval) for θ, and check if the UCB lies below θ_0. The pilot RISMC studies in [3] and [19] consider only a *single* criterion, PCT, and assume its capacity follows a triangular distribution. These papers apply a combination of stratified sampling (SS) and LHS, but they do not describe how to build a UCB for θ. (Other issues investigated in [3] and [19] include exploring the impact on θ from altering the distributions of input random variables and from operational changes, e.g., power uprates. A RISMC evaluation may further want to determine the core-damage frequency κ, which can be estimated by multiplying an estimator of θ with the (known) frequency of the postulated event.)

Our paper devises Monte Carlo methods to analyze a broad class of systems (not only for RISMC), with uncertainties encapsulated in a *basic random object*. The basic random object can be a random vector, as in structural reliability, where its entries are called *basic variables* [20, Section 1.5], which may be dependent and can represent, e.g., random loads, environmental factors, and material properties. But the basic random object may be more general, e.g., a stochastic process, as in time-dependent reliability [20, Chapter 6]. For example, the load and capacity on a system may vary randomly over time, which is modeled as a stochastic process. We also specify system failure in a general way, as a given binary-valued function of the basic random object. An example is a *series system*, where the basic random object is a random vector of the loads and capacities for a fixed number q of (dependent) criteria, and the system fails when any criterion's load exceeds its capacity, but our framework allows for many other possibilities.

We consider applying combinations of SS, LHS, and conditional Monte Carlo (CMC) to estimate θ. We formally prove that SS+LHS (resp., SS+CMC+LHS) has smaller asymptotic variance than SS (resp., SS+LHS). We also use replicated LHS (rLHS) [21] to construct UCBs for θ, and we prove the UCBs' asymptotic validity (as the total sample size grows large, with the number of replicates fixed). (Although we focus on providing UCBs, our methods can be easily modified to produce a lower confidence bound or two-sided confidence interval.) The combination of SS, CMC, and LHS can be especially effective in reducing variance, as we show through numerical experiments.

We also give insight into why SS+CMC+LHS can work so well. As shown in [22, Section 10.3], LHS substantially reduces variance when the response whose expectation we are estimating is well approximated by an additive function of the input random variables. As we are estimating a probability, the response without CMC is an indicator, which has a poor additive approximation. Thus, LHS by itself may not reduce variance much. In contrast, the conditioning of CMC produces a "smoother" response (in general, no longer binary-valued) with a better additive fit. Hence, LHS can yield much smaller variance when combined with CMC; [23] observes similar synergies between CMC and LHS.

In addition to [3] and [19], other related work includes [24], which provides UCBs for the *single-criterion* (i.e., q=1) failure probability when applying combinations of SS, CMC, and rLHS, but the paper does not give proofs of the UCBs' asymptotic validity nor for the reduction in asymptotic variance. The paper [25] estimates the single-criterion failure probability using SS and CMC, but leaves out LHS. Also, [23] provides a theoretical analysis of integrated VRTs (CMC, control variates, and correlation-induction schemes, including LHS and AV) in a general setting, but it does not include SS, which plays a key role in the RISMC framework, nor does the paper provide UCBs, as we do here. The paper [26] combines CMC with AV for estimating the failure probability in structural reliability, but it does not incorporate SS. As explained in [23], AV and LHS can be viewed as special cases of correlation-induction methods, but AV is often outclassed by LHS, especially in combination with CMC.

The methodologies developed in our current paper not only apply for nuclear PSAs, but also can be employed in a wide spectrum of other fields. Civil engineers need to compute the failure probability θ of a structure (e.g., a building, bridge, or dam) [20]. Insurance companies work with catastrophe models to determine the likelihood of infrastructure failures when subjected to hurricanes, floods, and earthquakes [27]. The Basel Accords [28] specify capital requirements (i.e., capacities) for financial institutions to ensure that they can absorb reasonable losses (i.e., loads). Other examples arise in the assessment of safety and reliability of nuclear weapons [29–31] and the disposal of radioactive waste [32–34].

Compared to SRS, LHS tries to sample more evenly from the sample space, which can produce less-variable estimators of measures of central tendency, e.g., the mean. Combining LHS with SS and CMC can lead to further substantial improvements, as our numerical results show, but perhaps not enough when estimating rare-event probabilities, smaller than say 10^{-4}. In such cases, other Monte Carlo methods may be more effective. In [35] and [36], the authors develop parametric approximations to the failure probability in terms of less-rare events, which are easier to estimate accurately via Monte Carlo. (In contrast, rather than devising an approximation, the methods in our paper work with the exact model.) The paper [37] employs the method of [35] to estimate the failure probability of a ship hull girder, modeled as a series system. In [38] the authors consider a method they call subset simulation, also known as splitting [14, Section VI.9]; the idea is to contain the rare failure event in a sequence of successively larger, less-rare events, and the failure probability is estimated as a product of estimates of conditional probabilities based on successive events. In [39] the author combines a type of CMC called directional simulation with importance sampling.

To summarize, the main contributions of the current paper are as follows:

- We develop combinations of SS, CMC, and LHS to efficiently estimate a failure probability θ. Our framework allows for a failure to be defined quite generally as a function of a basic random object, which may be a random vector or something more general, e.g., a stochastic process or random field. The combination of the three VRTs can produce substantial variance reduction, as our numerical results in Section 7 indicate. We also establish theory (Theorems 1 and 3) and provide insight (Sections 5.1 and 6.1) into why this is so.
- When applying SS+rLHS or SS+CMC+rLHS, we derive UCBs for θ, which are crucial to account for the statistical error arising from

sampling variability of the Monte Carlo estimators. We formally prove the UCBs' asymptotic validity; see Theorems 2 and 4.

The remaining sections unfold as follows. Section 2 develops the problem's mathematical framework in a general setting, with later sections adding extra structure as needed. Section 2.1 specializes the setup to define a system failure in terms of the loads and capacities of a fixed number of criteria, although later sections mostly do not require this additional assumption. Section 3 reviews the use of simple random sampling to estimate and produce a UCB for θ, and Section 4 does the same for stratified sampling. We combine SS with (replicated) Latin hypercube sampling in Section 5, and further incorporate conditional Monte Carlo in Section 6. In particular, Section 6.3 derives the combination of VRTs for the special setting of Section 2.1 (failure defined in terms of loads and capacities), exploiting the added structure. Section 7 presents numerical results from artificial RISMC experiments, and we give concluding remarks in Section 8. Appendices contain most of the proofs of the theorems. A (much shorter) conference paper [40] outlines some of the material from the current paper restricted to a series system, omitting Theorems 1, 3, and 5 and all of the proofs; moreover, the present paper contains additional numerical results. Throughout the paper, all vectors are column vectors, although we often write them as row vectors to save space.

List of Acronyms and Abbreviations

AHW	average half-width	AV	antithetic variates
CDF	cumulative distribution function	CLT	central limit theorem
CMC	conditional Monte Carlo	CV	coefficient of variation
CWO	core-wide oxidation	d.f.	degrees of freedom
i.i.d.	independent and identically distributed	KS	Kolmogorov-Smirnov
LHS	Latin hypercube sampling	LOCA	loss-of-coolant accident
MLE	maximum likelihood estimate	MLO	maximum local oxidation
NPP	nuclear power plant	NRC	Nuclear Regulatory Commission
PCD	probability of correct decision	PCT	peak cladding temperature
PSA	probabilistic safety assessment	RISMC	risk-informed safety-margin characterization
rLHS	replicated LHS	SCL	SS+CMC+LHS
SCrL	SS+CMC+rLHS	SRS	simple random sampling
UCB	upper confidence bound	VRF	variance-reduction factor
VRT	variance-reduction technique		

2. Mathematical framework

Before developing the general framework, we first motivate the problem by considering a simple model. The example actually does not require simulation as it can be solved numerically, but we use it to illustrate the ideas and notation. We will return to the example in later sections to introduce concepts for our various Monte Carlo methods.

Example 1. Consider a system with q=1 component, which experiences a (single) random *load* (stress or demand) L and has a (single) random *capacity* (resistance or strength) C to withstand the load. The system fails if and only if $L \geq C$. Let $\boldsymbol{X} = (L, C)$ be a random vector taking values in $\mathfrak{S} = \mathfrak{R}^2$, with $\mathfrak{R}$ the set of real numbers. Define the *failure function* $\mathfrak{d}(\boldsymbol{X}) = I(L \geq C)$ for $I(\cdot)$ the *indicator function*, which equals 1 (resp., 0) when its argument is true (resp., false). Thus, the system fails if and only if the *failure indicator* $\mathcal{I} \equiv \mathfrak{d}(\boldsymbol{X}) = 1$. Let H be the joint cumulative distribution function (CDF) of $\boldsymbol{X}$, denoted by $\boldsymbol{X} \sim H$, so $H(\boldsymbol{x}) = P(L \leq x_1, C \leq x_2)$ for a constant $\boldsymbol{x} = (x_1, x_2) \in \mathfrak{R}^2$. (Here L and C may be dependent, but later in Section 6 we will assume they are independent in this example.) The *failure probability* is then $\theta = E[\mathcal{I}] = P(L \geq C) = \int \mathfrak{d}(\boldsymbol{x})\, \mathrm{d}H(\boldsymbol{x})$, an integral of dimension $2q = 2$, where E denotes the expectation operator induced by H. Now assume that H is a *mixture* (e.g., Section 4.9 of [22]) of $s_0 = 4$ bivariate normal CDFs. Specifically, let λ_s, $1 \leq s \leq s_0$, be *known* positive constants summing to 1. For each s, let $H_{(s)}$ be the joint CDF of $\mathcal{N}_2(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s)$, a bivariate normal with mean vector $\boldsymbol{\mu}_s = (\mu_{L,s}, \mu_{C,s})$ and covariance matrix $\boldsymbol{\Sigma}_s \in \mathfrak{R}^{2\times 2}$ having diagonal entries $\sigma^2_{L,s}$ and $\sigma^2_{C,s}$, the marginal variances of L and C, respectively. Then $H(\boldsymbol{x}) = \sum_{s=1}^{s_0} \lambda_s H_{(s)}(\boldsymbol{x})$ for each $\boldsymbol{x} \in \mathfrak{R}^2$; i.e., $\boldsymbol{X} \sim H_{(s)}$ with probability λ_s. Let F (resp., G) be the marginal CDF of L (resp., C); e.g., G is the CDF of $\mathcal{N}(\sum_{s=1}^{s_0} \lambda_s \mu_{C,s}, (\sum_{s=1}^{s_0} \lambda_s(\sigma^2_{C,s} + \mu^2_{C,s})) - (\sum_{s=1}^{s_0} \lambda_s \mu_{C,s})^2)$, where $\mathcal{N}(a, b^2)$ is a normal random variable with mean a and variance b^2.

For the general mathematical framework, consider a stochastic model of the behavior of a system (e.g., a structure, such as a building or nuclear power plant) over a particular time period, possibly random. The system has uncertainties, e.g., loads and capacities of the system's components, environmental factors, material properties, etc., which may be observed at multiple times. We encapsulate all of the system's (aleatory and epistemic) uncertainties over the specified time horizon in a *basic random object* $\boldsymbol{X}$, which may be a random vector (Example 1 has $\boldsymbol{X} = (L, C) \in \mathfrak{R}^2$) or a more general random object, such as a stochastic process or random field. Specifically, let $(\Omega, \mathcal{F}, P)$ be the underlying probability space, where Ω is a sample space, $\mathcal{F}$ is a σ-field of subsets of Ω, and P is a probability measure on $\mathcal{F}$. We assume that the basic random object $\boldsymbol{X}$ takes values in a (complete, separable) metric space $(\mathfrak{S}, \mathcal{S})$, where $\mathfrak{S}$ is a set of possible objects $\boldsymbol{x}$ and $\mathcal{S}$ is a σ-field of subsets of $\mathfrak{S}$ [41, Appendix M], so $\boldsymbol{X}$: $\Omega \to \mathfrak{S}$. In other words, for each outcome $\omega \in \Omega$, we have that $\boldsymbol{X}(\omega) = \boldsymbol{x}$ for some $\boldsymbol{x} \in \mathfrak{S}$. Also, if $S \in \mathcal{S}$ is a subset of $\mathfrak{S}$, then $P(\boldsymbol{X} \in S) = P(\{\omega: \boldsymbol{X}(\omega) \in S\})$ is the probability that $\boldsymbol{X}$ lies in S. For example, if $\mathfrak{S} = \mathfrak{R}^d$, then $\boldsymbol{X} = (X^{[1]}, X^{[2]}, \ldots, X^{[d]})$ is a d-dimensional random vector (d=2 in Example 1), and the d entries in $\boldsymbol{X}$ are called *basic variables* in structural reliability [20, Section 1.5]. In this case, for each outcome $\omega \in \Omega$, we have $\boldsymbol{X}(\omega) = (X^{[1]}(\omega), X^{[2]}(\omega), \ldots, X^{[d]}(\omega)) = \boldsymbol{x}$ for some $\boldsymbol{x} = (x_1, x_2, \ldots, x_d) \in \mathfrak{R}^d$; also, for the set $S = \{(x_1, x_2, \ldots, x_d) \in \mathfrak{R}^d: x_k \leq y_k, k = 1, 2, \ldots, d\} \in \mathcal{S}$ for a fixed $(y_1, y_2, \ldots, y_d) \in \mathfrak{R}^d$, we have $P(\boldsymbol{X} \in S) = P(X^{[1]} \leq y_1, X^{[2]} \leq y_2, \ldots, X^{[d]} \leq y_d)$. The probability P may specify dependencies within the basic random object $\boldsymbol{X}$, e.g., loads of different components may be dependent. Section 7.2 describes how we can use copulas [42] to incorporate dependence.

We allow $\mathfrak{S}$ to be more general than $\mathfrak{R}^d$. The space $(\mathfrak{S}, \mathcal{S})$ may be infinite-dimensional, as can occur when uncertainties are modeled through a stochastic process. For example, over a continuous time interval $[0, t_0]$, a structure may be subjected to random events (e.g., earthquakes or storms) arising according to a Poisson process [43, Section 23], and each event has a random intensity; as the number of events occurring by t_0 is unbounded, the vector $\boldsymbol{X}$ of the event times and intensities then requires $\mathfrak{S}$ to have infinite dimension. Also, as in time-dependent reliability [20, Chapter 6], the system's *state* (which may include, e.g., the loads, capacities, environmental conditions) evolves randomly within a *state space* $\mathfrak{E}$ (e.g., $\mathfrak{E} \subseteq \mathfrak{R}^d$) over continuous time. In this case, the basic random object is a continuous-time, $\mathfrak{E}$-valued stochastic process $\boldsymbol{X} = [X(t): t \geq 0]$, where $X(t)$ is the system's state in $\mathfrak{E}$ at time t, so $(\mathfrak{S}, \mathcal{S})$ is an appropriate space of $\mathfrak{E}$-valued functions on $[0, \infty)$ [41, Chapters 2 and 3]. For example, consider a system with $p \geq 1$ components, where each component $1 \leq k \leq p$ has load $L^{[k]}(t)$ and capacity $C^{[k]}(t)$ at each time $t \leq t_0$. Then we may define $\boldsymbol{X} = [X(t): 0 \leq t \leq t_0]$ with $\mathfrak{E} = \mathfrak{R}^{2p}$ and

$$X(t) = (L^{[1]}(t), L^{[|2|]}(t), \ldots, L^{[p]}(t), C^{[1]}(t), C^{[|2|]}(t), \ldots, C^{[p]}(t)). \quad (1)$$

Next define a binary-valued *failure function* $\mathfrak{d}$: $\mathfrak{S} \to \{0, 1\}$ such

that if $\boldsymbol{X} = \boldsymbol{x}$ with $\boldsymbol{x} \in \mathfrak{S}$ and $\mathfrak{d}(\boldsymbol{x}) = 1$ (resp., 0), the system fails (resp., does not fail) within the given time horizon. Applying $\mathfrak{d}$ to the basic random object $\boldsymbol{X}$ yields

$$\mathcal{I} \equiv \mathcal{I}(\boldsymbol{X}) = \mathfrak{d}(\boldsymbol{X}) \tag{2}$$

as the *failure indicator*, which is a Bernoulli random variable that equals 1 (resp., 0) if the system eventually fails (resp., does not fail). Computing $\mathfrak{d}$ and $\mathcal{I}$ can be quite complicated. For example, in a PSA of a nuclear power plant, $\mathfrak{d}$ is based on a detailed computer code [7] that models the progression of a hypothesized accident, taking as input a random vector $\boldsymbol{X} = (X^{[1]}, X^{[2]}, \ldots, X^{[d]}) \in \mathfrak{S} = \mathfrak{R}^d$ representing aleatory and epistemic uncertainties with a specified joint CDF. The entries of $\boldsymbol{X}$, which may determine the timing, location and size of events during the accident, may be dependent and may have different marginal CDFs. The computer code can be computationally expensive to run for each $\boldsymbol{x} \in \mathfrak{S}$, as it may numerically solve systems of differential equations.

In structural reliability [20], the failure function $\mathfrak{d}$ in (2) is often defined in terms of a *limit state function* $\mathfrak{g}: \mathfrak{S} \to \mathfrak{R}$, where $\mathfrak{g}(\boldsymbol{X}) \le 0$ (resp., >0) denotes that the system fails (resp., does not fail). In this case, $\mathfrak{d}(\boldsymbol{X}) = I(\mathfrak{g}(\boldsymbol{X}) \le 0)$. Example 1 has $\mathfrak{g}(L, C) = C - L$, which is a *safety margin*, and $\mathfrak{d}(L, C) = I(L \ge C)$. Section 2.1 will provide other examples illustrating how $\mathfrak{d}$ may be defined for common types of systems in terms of multiple loads and capacities.

When the basic random object $\boldsymbol{X}$ is a continuous-time stochastic process $[X(t): t \ge 0]$, as in time-dependent reliability analysis, we may define the failure function $\mathfrak{d}$ in (2) in many ways. For example, consider $\boldsymbol{X}$ with $X(t)$ in (1), and let t_0 be a fixed time horizon. Then we may define $\mathfrak{d}(\boldsymbol{X}) = I(L^{[k]}(t) \ge C^{[k]}(t) \text{ for some } 1 \le k \le p \text{ and some } 0 \le t \le t_0)$, so the system fails if and only if some component's load exceeds its capacity before the time horizon. For another example, let $\tau > 0$ be a fixed constant, and set $\mathfrak{d}(\boldsymbol{X}) = I\left(\int_0^{t_0} I(L^{[k]}(t) \ge C^{[k]}(t))\, \mathrm{d}t \ge \tau \text{ for some } 1 \le k \le p\right)$, so the system fails if and only if for some component k, the amount of time its load exceeds its capacity is at least τ.

We define the system's *failure probability* (also known as the *unreliability*) as

$$\theta = P(\mathcal{I} = 1) = E[\mathcal{I}] = \int \mathcal{I}(\boldsymbol{X})\, \mathrm{d}P = \int \mathfrak{d}(\boldsymbol{X})\, \mathrm{d}P, \tag{3}$$

where E is the expectation operator induced by P. The integrals in (3) have the dimension (possibly infinite) of the space $\mathfrak{S}$. In the third expression of (3), we call the quantity inside the expectation a *response* (*function*), and we want to compute its expectation θ, which is an unknown constant. To avoid trivialities, we assume that $0 < \theta < 1$.

The system is deemed to be acceptably safe if the failure probability is sufficiently small, i.e., $\theta < \theta_0$ for some given constant θ_0. We assume the stochastic model is too complex to analytically or numerically compute θ, but it can be simulated on a computer (which may, e.g., require discretizing time). Thus, we will develop (Monte Carlo) simulation methods to estimate θ. The corresponding simulation model may be a computer code that generates an observation of $\boldsymbol{X}$ according to P, and eventually outputs $\mathcal{I}$ using (2). The resulting simulation estimators have statistical error because of variability in the sampled observations of $\boldsymbol{X}$, which we account for by further requiring the following:

$$\begin{aligned}&\text{given constants } 0 < \theta_0 < 1 \text{ and } 0 < \gamma < 1,\\ &\text{determine with confidence level } \gamma \text{ if } \theta < \theta_0.\end{aligned} \tag{4}$$

The constants θ_0 and γ may be specified by a regulator. For example, if $\theta_0 = 0.05$ and $\gamma = 0.95$, we want to determine with 95% confidence if $\theta < 0.05$; when there is only a single criterion with constant capacity, this is equivalent to a 95/95 requirement.

We can satisfy the requirement (4) through a hypothesis test. Define the null hypothesis $\mathcal{H}_0: \theta \ge \theta_0$ and alternative hypothesis $\mathcal{H}_1: \theta < \theta_0$. We perform a hypothesis test at significance level $\alpha = 1 - \gamma$ by first carrying out a total of n simulation runs. We then use the output data to construct a (statistical) *point estimator* $\hat{\theta}(n)$ of θ, along with a γ-level *upper confidence bound* (UCB) $B(n)$ for θ; i.e., $P(B(n) \ge \theta) = \gamma$. (We define the phrase "point estimator" or "point estimate" in the statistical sense, as a sampling-based estimator or estimate of a parameter—such as an expectation or probability—whose true (constant) value is unknown. This is in contrast to the meaning sometimes adopted in nuclear safety, where a "point estimate" may denote a single conservative value for a load or capacity, instead of assuming probability distributions for them, as when incorporating uncertainties.) The exact form of $B(n)$ depends on the simulation method applied and how $\hat{\theta}(n)$ is constructed. But without imposing unreasonable parametric assumptions (e.g., normally distributed outputs), it is typically not possible to construct a UCB $B(n)$ for which $P(B(n) \ge \theta) = \gamma$ exactly holds for a fixed sample size n. Instead, we desire $B(n)$ to be an *asymptotic* γ-level UCB, satisfying

$$\lim_{n \to \infty} P(B(n) \ge \theta) = \gamma, \tag{5}$$

where $B(n)$ is often derived from a central limit theorem (CLT) for $\hat{\theta}(n)$. We then arrive at the following *decision rule*:

$$\text{Reject } \mathcal{H}_0 \text{ if and only if } B(n) < \theta_0, \tag{6}$$

which asymptotically fulfills (4) as $n \to \infty$ when (5) holds. Starting in Section 3 we consider specific simulation methodologies to construct asymptotic UCBs achieving (5). (See [44] for a thorough treatment of asymptotic statistical methods.) While the current paper focuses on UCBs, the approaches presented can be easily modified to also provide lower confidence bounds or two-sided confidence intervals. Generating the basic random object $\boldsymbol{X}$ or computing the failure function $\mathfrak{d}$ in (2) may be quite costly, making it important to apply VRTs to reduce the sample size n necessary to obtain sufficiently accurate results.

2.1. Examples of defining the failure function in terms of loads and capacities

We now present some common ways of defining the failure function $\mathfrak{d}$ in (2). Our examples specify system failure in terms of a *fixed* number $q \ge 1$ of *criteria* (or *failure modes*), where each has a random *load* and a random *capacity*. Specifically, let $\boldsymbol{L} \equiv \boldsymbol{L}(\boldsymbol{X}) = (L^{[1]}, L^{[2]}, \ldots, L^{[q]})$ and $\boldsymbol{C} \equiv \boldsymbol{C}(\boldsymbol{X}) = (C^{[1]}, C^{[2]}, \ldots, C^{[q]})$ be random vectors computed from the basic random object $\boldsymbol{X}$, with $L^{[k]}$ and $C^{[k]}$ as the load and capacity, respectively, for *criterion* $k = 1, 2, \ldots, q$. For example, each of the q criteria may correspond to a component in the system, and a component fails when its load exceeds its capacity; the system then fails when specified combinations of components fail. (In Example 1, there is only q=1 component, so $\boldsymbol{L} = L = L^{[1]}$ and $\boldsymbol{C} = C = C^{[1]}$.) The load and capacity of a component i may be observed at multiple times $t_{i,1} < t_{i,2} < \cdots < t_{i,m_i}$, and each pair of component i and time $t_{i,j}$ could correspond to a criterion. A nuclear PSA may have q=3 criteria, and $L^{[1]}$, $L^{[2]}$, and $L^{[3]}$ represent the loads for criteria PCT, CWO, and MLO, respectively, during a hypothesized loss-of-coolant accident; see paragraph 50.46(b) of [4]. The initial RISMC studies [3,19] consider only q=1 criterion, PCT, and assume a triangular distribution for its capacity, but we allow for $q > 1$. Throughout the paper, we permit for dependence across criteria (e.g., different components have related loads, or capacities at different times are correlated), and except for Section 6.3, $\boldsymbol{L}$ and $\boldsymbol{C}$ can be dependent.

When defining the failure function $\mathfrak{d}$ in (2) in terms of $\boldsymbol{L}$ and $\boldsymbol{C}$, we will assume the following.

Assumption A1. There are deterministic functions $\mathfrak{v}: \mathfrak{S} \to \mathfrak{R}^{2q}$ and $\mathfrak{d}': \mathfrak{R}^{2q} \to \{0, 1\}$ such that $(\boldsymbol{L}, \boldsymbol{C}) = \mathfrak{v}(\boldsymbol{X})$ and $\mathfrak{d}(\boldsymbol{x}) = \mathfrak{d}'(\mathfrak{v}(\boldsymbol{x}))$ for all $\boldsymbol{x} \in \mathfrak{S}$.

In Assumption A1, the function $\mathfrak{v}$ computes the load and capacity vectors $\boldsymbol{L}$ and $\boldsymbol{C}$ from the basic random object $\boldsymbol{X}$. For example, $\boldsymbol{L}$ may be a function of random stresses, which are included in $\boldsymbol{X}$, and $\boldsymbol{C}$ may be computed from uncertain environmental conditions and material

properties specified in $\boldsymbol{X}$. Moreover, A1 further stipulates that the failure function $\mathfrak{d}$ computed in terms of $\boldsymbol{X}$ in (2) can be equivalently expressed as the function $\mathfrak{d}'$ of only the loads $\boldsymbol{L}$ and capacities $\boldsymbol{C}$ because $\mathfrak{d}(\boldsymbol{X}) = \mathfrak{d}'(\mathfrak{v}(\boldsymbol{X})) = \mathfrak{d}'(\boldsymbol{L}, \boldsymbol{C})$. Example 1 already has $\boldsymbol{X} = (L, C)$, so in this case, $\mathfrak{v}$ is simply the identity mapping, and $\mathfrak{d}(\boldsymbol{X}) = \mathfrak{d}'(L, C) = I(L \geq C)$. We now show how to define the failure function $\mathfrak{d}'$ of the $q \geq 1$ loads and capacities for some common types of systems.

Example 2. (Series system; e.g., Section 5.2.3.1 of [20]). Suppose a system with q criteria fails if and only if at least one criterion's load exceeds its capacity. Then by the inclusion-exclusion principle, we can define $\mathfrak{d}'$ in A1 as

$$\begin{aligned}\mathfrak{d}'(\boldsymbol{L}, \boldsymbol{C}) &= I(\cup_{k=1}^{q}\{L^{[k]} \geq C^{[k]}\}) \\ &= \sum_{p=1}^{q}(-1)^{p-1}\sum_{1\leq k_1<k_2<\cdots<k_p\leq q} I(\cap_{l=1}^{p}\{L^{[k_l]} \geq C^{[k_l]}\}).\end{aligned} \tag{7}$$

The paper [45] defines failure as in (7) for a nuclear PSA with q=3 criteria (PCT, CWO, MLO) and fixed capacities.

Example 3. (Parallel system; e.g., Section 5.2.3.2 of [20]). Suppose a system with q criteria fails if and only if each criterion's load exceeds its capacity. Then the failure function $\mathfrak{d}'$ in Assumption A1 is given by

$$\mathfrak{d}'(\boldsymbol{L}, \boldsymbol{C}) = I(\cap_{k=1}^{q}\{L^{[k]} \geq C^{[k]}\}). \tag{8}$$

Example 4. (K-out-of-N:F system). Suppose a system with $\mathrm{N} = q$ criteria fails if and only if there exist (at least) K criteria $k_1 < k_2 < \cdots < k_{\mathrm{K}}$ such that $L^{[k_l]} \geq C^{[k_l]}$ for all $l = 1, 2, \ldots, \mathrm{K}$. Then we define the failure function $\mathfrak{d}'$ in A1 as

$$\begin{aligned}\mathfrak{d}'(\boldsymbol{L}, \boldsymbol{C}) &= I\left(\cup_{p=\mathrm{K}}^{q}\cup_{1\leq k_1<k_2<\cdots<k_p\leq q}\cap_{l=1}^{p}\{L^{[k_l]} \geq C^{[k_l]}\}\right) \\ &= \sum_{p=\mathrm{K}}^{q}(-1)^{p-\mathrm{K}}\binom{p-1}{\mathrm{K}-1}\sum_{1\leq k_1<k_2<\cdots<k_p\leq q} I(\cap_{l=1}^{p}\{L^{[k_l]} \geq C^{[k_l]}\}),\end{aligned} \tag{9}$$

where $\binom{m}{j} = m!/(j!(m-j)!)$ and the second equality can be shown as in [46]. A series system (Example 2) is a special case with $\mathrm{K} = 1$. Another special case is a parallel system (Example 3), which has $\mathrm{K} = \mathrm{N} = q$.

Example 5. (Series-parallel system). Suppose the q criteria are partitioned into $v \geq 1$ subsets (or subsystems), where subset $m = 1, 2, \ldots, v$, comprises criteria $k = q_m, q_m + 1, \ldots, q_{m+1} - 1$, and $q_1 = 1 < q_2 < \cdots < q_v < q_{v+1} \equiv q + 1$. Suppose further that the system fails if and only if for at least one subset m, the load exceeds the capacity for each of the subset's criteria. Then in Assumption A1 the failure function $\mathfrak{d}'$ is defined as

$$\begin{aligned}\mathfrak{d}'(\boldsymbol{L}, \boldsymbol{C}) &= I\left(\cup_{m=1}^{v}\cap_{k=q_m}^{q_{m+1}-1}\{L^{[k]} \geq C^{[k]}\}\right) \\ &= \sum_{p=1}^{v}(-1)^{p-1}\sum_{1\leq m_1<m_2<\cdots<m_p\leq v} I\left(\cap_{l=1}^{p}\cap_{k=q_{m_l}}^{q_{m_l+1}-1}\{L^{[k]} \geq C^{[k]}\}\right).\end{aligned} \tag{10}$$

Some of the q criteria may be the same, i.e., there may exist criteria $k_1 \neq k_2$ such that $P(L^{[k_1]} = L^{[k_2]}, C^{[k_1]} = C^{[k_2]}) = 1$. This allows the same criterion to be parts of different subsets. A series system (Example 2) is a special case with v=q subsets, where each subset m contains only criterion m; i.e., $q_m = m$ for each $m = 1, 2, \ldots, q + 1$. A parallel system (Example 3) is another special case with only v=1 subset having all q criteria; i.e., $q_1 = 1$ and $q_2 = q + 1$.

Example 6. (Parallel-series system). Suppose the q criteria are again partitioned into v subsets as in Example 5. Suppose further that the system fails if and only if for each subset $m = 1, 2, \ldots, v$, the load exceeds the capacity for at least one of the subset's criteria. Then the failure function $\mathfrak{d}'$ in Assumption A1 is defined as

$$\begin{aligned}\mathfrak{d}'(\boldsymbol{L}, \boldsymbol{C}) &= I\left(\cap_{m=1}^{v}\cup_{k=q_m}^{q_{m+1}-1}\{L^{[k]} \geq C^{[k]}\}\right) \\ &= \prod_{m=1}^{v}\sum_{p=1}^{q_{m+1}-q_m}(-1)^{p-1}\sum_{q_m\leq k_1<k_2<\cdots<k_p\leq q_{m+1}-1} I(\cap_{l=1}^{p}\{L^{[k_l]} \geq C^{[k_l]}\}).\end{aligned} \tag{11}$$

As in Example 5, we allow that some of the q criteria may be the same. A series system (Example 2) is a special case with v=1 subset having all q criteria; i.e., $q_1 = 1$ and $q_2 = q + 1$. Another special case is a parallel system (Example 3), which has v=q subsets, where each subset m contains only criterion m; i.e., $q_m = m$ for each $m = 1, 2, \ldots, q + 1$.

Let H be the joint CDF of $(\boldsymbol{L}, \boldsymbol{C}) \in \Re^{2q}$; i.e., $(\boldsymbol{L}, \boldsymbol{C}) \sim H$, and $H(\boldsymbol{x}, \boldsymbol{y}) = P(\boldsymbol{L} \leq \boldsymbol{x}, \boldsymbol{C} \leq \boldsymbol{y})$ for constant q-vectors $\boldsymbol{x}$ and $\boldsymbol{y}$. Under Assumption A1, we can rewrite (3) as $\theta = \int_{\Re^{2q}} \mathfrak{d}'(\boldsymbol{x}, \boldsymbol{y})\, \mathrm{d}H(\boldsymbol{x}, \boldsymbol{y})$. We assume that computing θ is intractable, but there is a simulation model that produces an observation of $\boldsymbol{X}$ so that $(\boldsymbol{L}, \boldsymbol{C}) = \mathfrak{v}(\boldsymbol{X}) \sim H$. Let F (resp., G) denote the marginal joint CDF of the random vector $\boldsymbol{L}$ (resp., $\boldsymbol{C}$). The setting of fixed capacities is a special case with degenerate G. Example 1 has $(L, C) \in \Re^{2q}$ with q=1 and H a mixture of bivariate normals, so F and G are mixtures of univariate normals. In general, we initially make no assumptions about the joint CDF H; e.g., we do not require that H is a particular multivariate parametric distribution, such as Gaussian or lognormal. Also, H may incorporate dependence among the $2q$ elements of $(\boldsymbol{L}, \boldsymbol{C})$. Later sections impose conditions as needed when introducing different simulation methodologies, which will increase statistical efficiency by exploiting the resulting additional problem structure (e.g., $\boldsymbol{L}$ is independent of $\boldsymbol{C}$ in Section 6.3, but each vector may still have dependencies).

While the framework of Assumption A1 is quite broad, it does restrict the definition of failure to be in terms of a finite number of criteria rather than a possibly infinite-dimensional random object. The methods we will develop in later sections allow for the more general definition of failure in (2) and as such will not require A1. However, incorporating A1 will sometimes enable us to design more efficient techniques, e.g., as in Section 6.3. Throughout the paper we will clearly identify when A1 (and other assumptions) are needed.

Assumption A1 may not hold when the basic random object is a continuous-time stochastic process, as in time-dependent reliability [20, Chapter 6]. For example, [31] defines a "loss of assured safety" as a system failure, as we now explain. The system has p components, where the loads and capacities may change randomly over continuous time. Let $L^{[k]}(t)$ and $C^{[k]}(t)$ be the load and capacity, respectively, of component k at time t, and let $\boldsymbol{X} = [X(t)\colon t \geq 0]$ be the continuous-time stochastic process with $X(t)$ from (1). Define the failure time of component k as $T^{[k]} = \inf\{t \geq 0\colon L^{[k]}(t) \geq C^{[k]}(t)\}$, which is the first time the load of component k exceeds its capacity. Then partitioning the p components into two nonempty sets, D_s ("strong links") and D_w ("weak links"), [31] provides several possible ways to define system failure in terms of the p components' failure times. One definition is $\mathfrak{d}(\boldsymbol{X}) = I(\sup_{k\in D_s} T^{[k]} \leq \inf_{k\in D_w} T^{[k]})$, so the system fails if and only if all strong links fail before any weak link fails. Under the assumption that the p components are mutually independent, [31] provides a representation for the failure probability $\theta = E[\mathfrak{d}(\boldsymbol{X})]$ in terms of the marginal CDFs of the component failure times, and examines different ways of computing θ. The Monte Carlo methods we develop allow for dependence of the components. In its full generality, the model just described does not satisfy Assumption A1 because the continuous-time stochastic process $\boldsymbol{X}$ cannot be reduced to a finite number of scalar loads and capacities (but it still fits in the setting of Section 2). (As noted in the previous paragraph, the Monte Carlo methods we develop in the later sections do not require A1.) But we can develop a discrete-time version over a finite time horizon for which A1 does hold by defining a criterion for each component at each discrete time step.

3. Simple random sampling

We first consider estimating θ with simple random sampling, which is also known as *naive simulation*, *standard simulation*, or *crude Monte Carlo*. The approach to construct a UCB satisfying (5) when applying SRS, which does not require Assumption A1 from Section 2.1, is well known (e.g., Example 3.2.4 of [44]), but we review the details here as analogous arguments will be employed for the other Monte Carlo methods in later sections.

Example 1. (Continued). We now explain how to sample $\boldsymbol{X} = (L, C) \sim H$, where we recall H is a mixture of the joint CDFs $H_{(s)}$ of $\mathcal{N}_2(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s)$, $1 \le s \le s_0$, with known mixing weights λ_s. First generate an observation of a discrete random variable A with $P(A = s) = \lambda_s$, $s = 1, 2, \ldots, s_0$. If A=s, let $\boldsymbol{X} = \boldsymbol{\mu}_s + \boldsymbol{\Gamma}_s \boldsymbol{Z}$, where $\boldsymbol{Z} = (Z_1, Z_2)$ is a (column) vector of i.i.d. $\mathcal{N}(0, 1)$ random variables, $\boldsymbol{\Gamma}_s$ is a 2×2 matrix satisfying $\boldsymbol{\Gamma}_s \boldsymbol{\Gamma}_s^{\top} = \boldsymbol{\Sigma}_s$ (e.g., $\boldsymbol{\Gamma}_s$ may be a Cholesky factor of $\boldsymbol{\Sigma}_s$ [13, pp. 72–74]), and the superscript $\top$ denotes transpose. Thus, with probability λ_s, we have that $\boldsymbol{X} \sim H_{(s)}$, as required. We then estimate $\theta = E[\mathcal{I}]$ by averaging independent and identically distributed (i.i.d.) copies of $\mathcal{I} = \mathfrak{d}(\boldsymbol{X}) = I(L \ge C)$.

More generally, to implement SRS, first generate a sample of n i.i.d. copies $\boldsymbol{X}_i$, $i = 1, 2, \ldots, n$, of the basic random object $\boldsymbol{X}$ in (2), and set $\mathcal{I}_i = \mathfrak{d}(\boldsymbol{X}_i)$. Then the SRS point estimator of θ in (3) is $\hat{\theta}_{\mathrm{SRS}}(n) = (1/n) \sum_{i=1}^{n} \mathcal{I}_i$, which is *unbiased*; i.e., $E[\hat{\theta}_{\mathrm{SRS}}(n)] = \theta$. The SRS point estimator averages i.i.d. copies of a bounded response, so

$$\frac{\sqrt{n}}{\hat{\sigma}_{\mathrm{SRS}}(n)} \left[\hat{\theta}_{\mathrm{SRS}}(n) - \theta\right] \Rightarrow \mathcal{N}(0, 1) \quad \text{as } n \to \infty, \tag{12}$$

by the ordinary CLT [43, Theorem 27.1], where $\Rightarrow$ denotes convergence in distribution (i.e., at each continuity point of the limit's CDF, the CDF of the left side of (12) converges to the limit's CDF [43, Chapter 5]), and $\hat{\sigma}^2_{\mathrm{SRS}}(n) = \hat{\theta}_{\mathrm{SRS}}(n)[1 - \hat{\theta}_{\mathrm{SRS}}(n)]$ *consistently* estimates $\sigma^2_{\mathrm{SRS}} = \theta[1 - \theta]$; i.e, $\hat{\sigma}^2_{\mathrm{SRS}}(n) \Rightarrow \sigma^2_{\mathrm{SRS}}$ as $n \to \infty$. Because the normal limit in (12) has a continuous distribution, the portmanteau theorem (e.g., Theorem 29.1 of [43]) ensures that

$$P\left(\frac{\sqrt{n}}{\hat{\sigma}_{\mathrm{SRS}}(n)} [\hat{\theta}_{\mathrm{SRS}}(n) - \theta] \ge x\right) \to P(\mathcal{N}(0, 1) \ge x) \quad \text{as } n \to \infty \tag{13}$$

for each real-valued constant x. Let z_γ be the γ-level upper one-sided critical point of a $\mathcal{N}(0, 1)$ random variable, which satisfies $P(\mathcal{N}(0, 1) \le z_\gamma) = P(\mathcal{N}(0, 1) \ge -z_\gamma) = \gamma$, e.g., $z_{0.95} = 1.645$. Hence,

$$B_{\mathrm{SRS}}(n) \equiv \hat{\theta}_{\mathrm{SRS}}(n) + z_\gamma \frac{\hat{\sigma}_{\mathrm{SRS}}(n)}{\sqrt{n}} \tag{14}$$

is an asymptotic γ-level UCB satisfying (5) because $P(B_{\mathrm{SRS}}(n) \ge \theta) = P(\sqrt{n}[\hat{\theta}_{\mathrm{SRS}}(n) - \theta]/\hat{\sigma}_{\mathrm{SRS}}(n) \ge -z_\gamma) \to P(\mathcal{N}(0, 1) \ge -z_\gamma) = \gamma$ as $n \to \infty$ by (13) with $x = -z_\gamma$. Computing $B_{\mathrm{SRS}}(n)$ uses only the observed simulation outputs and known constants, and does not require any unknown parameters. A simple modification of the well-known sign test (e.g., Example 3.2.4 of [44]), this approach has been specialized to multi-criteria nuclear PSAs with fixed capacities in Section 4.2 of [47].

We can easily modify the UCB $B_{\mathrm{SRS}}(n)$ to obtain an asymptotic γ-level lower confidence bound for θ by changing the plus in (14) to a minus. An asymptotic γ-level two-sided confidence interval for θ is simply $(\hat{\theta}_{\mathrm{SRS}}(n) \pm z_{1-(1-\gamma)/2} \hat{\sigma}_{\mathrm{SRS}}(n)/\sqrt{n})$. For the UCBs in the rest of the paper, we can similarly modify them to obtain lower confidence bounds and two-sided confidence intervals, so we omit the details.

If the failure function $\mathfrak{d}$ is defined as in Assumption A1 in terms of a function $\mathfrak{d}'$ of q loads and capacities, we apply SRS with sample size n by letting $(\boldsymbol{L}_i, \boldsymbol{C}_i) = \mathfrak{v}(\boldsymbol{X}_i)$, $i = 1, 2, \ldots, n$, using $\mathfrak{v}$ from A1, where $\boldsymbol{L}_i = (L_i^{[1]}, L_i^{[2]}, \ldots, L_i^{[q]})$, $\boldsymbol{C}_i = (C_i^{[1]}, C_i^{[2]}, \ldots, C_i^{[q]})$, and $\boldsymbol{X}_i$ is the ith i.i.d. copy of the basic random object $\boldsymbol{X}$, as before. Then estimator $\hat{\theta}_{\mathrm{SRS}}(n)$ averages the $\mathcal{I}_i = \mathfrak{d}'(\boldsymbol{L}_i, \boldsymbol{C}_i)$, $1 \le i \le n$.

In a RISMC study of a postulated NPP event (e.g., a LOCA), it is typically assumed that a failure, as in (7), results in damage to the core.

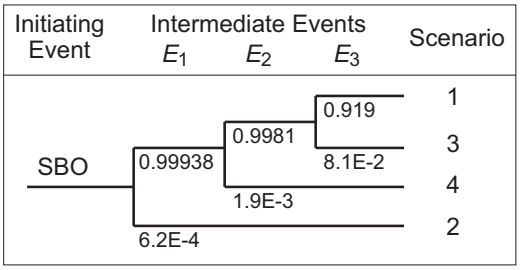

Fig. 1. An event tree for a hypothesized station blackout (SBO).

One is then sometimes also interested in estimating the *core-damage frequency* $\kappa \equiv \nu\theta$, where ν is the (known) frequency of the postulated event. When applying SRS, we can then estimate κ by $\hat{\kappa}_{\mathrm{SRS}}(n) = \nu\hat{\theta}_{\mathrm{SRS}}(n)$, and an asymptotic γ-level UCB for κ is $\hat{\kappa}_{\mathrm{SRS}}(n) + z_\gamma \nu \hat{\sigma}_{\mathrm{SRS}}(n)/\sqrt{n}$. For the other Monte Carlo methods considered in the rest of the paper, we can similarly estimate and construct confidence bounds for κ; we omit the details.

4. Stratified sampling

A key aspect of the RISMC approach is decomposing the sample space of a hypothesized event into $s_0 \ge 1$ *scenarios* through an *event tree* [18]. For example, Fig. 1 depicts an event tree from [3] with $s_0 = 4$ scenarios of a postulated station blackout. The intermediate events E_1, E_2, E_3 determine how the accident progresses; e.g., the lower (resp., upper) branch of E_2 is the event that a safety relief valve is stuck open (resp., properly closes). The branching probabilities of the intermediate events are assumed known, as shown in the figure, from, e.g., previous risk studies or historical data. Following a path from left to right through the event tree leads to a scenario $s = 1, 2, \ldots, s_0$, and the probability λ_s of scenario s occurring is computed by multiplying the branching probabilities of the intermediate events along the path; e.g., scenario 4 occurs with probability $\lambda_4 = 0.99938 \times 0.0019$. But the failure probability $\theta_{(s)}$ of each scenario s is unknown, and must be estimated via some form of simulation.

Example 1. (Continued). In Fig. 1, let A be a discrete random variable denoting which scenario s occurs, so $P(A = s) = \lambda_s$, $1 \le s \le s_0$, where the λ_s are known. When A=s, assume that $\boldsymbol{X} = (L, C)$ has the $\mathcal{N}_2(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s)$ CDF $H_{(s)}$, so we can set $\boldsymbol{X}$ equal to a random vector $\boldsymbol{X}_{(s)} \sim H_{(s)}$. Thus, as before, the unconditional CDF H of $\boldsymbol{X}$ is a mixture of CDFs $H_{(s)}$, $1 \le s \le s_0$, with mixing weights λ_s. Also, $\boldsymbol{X}_{(s)} = (L_{(s)}, C_{(s)})$ has the conditional distribution of $\boldsymbol{X} = (L, C)$ given A=s. Let $F_{(s)}$ (resp., $G_{(s)}$) be the conditional marginal CDF of $L_{(s)}$ (resp., $C_{(s)}$); e.g., $G_{(s)}$ is the $\mathcal{N}(\mu_{C,s}, \sigma^2_{C,s})$ CDF. Given scenario s occurs, the failure indicator $\mathcal{I}_{(s)} = \mathfrak{d}(\boldsymbol{X}_{(s)}) = I(L_{(s)} \ge C_{(s)})$ has the conditional distribution of $\mathcal{I}$ in (2), and the (conditional) failure probability is then $\theta_{(s)} = E[\mathcal{I}_{(s)}] = P(\mathcal{I}_{(s)} = 1) = P(L_{(s)} \ge C_{(s)}) = \int_{\Re^2} \mathfrak{d}(\boldsymbol{x})\, dH_{(s)}(\boldsymbol{x})$. We can then estimate $\theta_{(s)}$ by averaging i.i.d. copies of $\mathcal{I}_{(s)} = \mathfrak{d}(\boldsymbol{X}_{(s)}) = I(L_{(s)} \ge C_{(s)})$.

The event-tree framework is well suited for estimating θ via stratified sampling, with each scenario corresponding to a stratum. In general, SS partitions the sample space into $s_0 \ge 1$ strata, where each stratum has a known probability of occurring; see Section 4.3 of [13] for an overview of SS. This is done through a *stratification variable* A, which is a random variable A: $\Omega \to R$ defined on the same probability space $(\Omega, \mathcal{F}, P)$ as the basic random object $\boldsymbol{X}$. Partition the support R of A into s_0 subsets $R_1, R_2, \ldots, R_{s_0}$, with $R = \cup_{s=1}^{s_0} R_s$ and $R_s \cap R_{s'} = \varnothing$ for $s \ne s'$. We call each R_s a *stratum*, and its index s a *scenario*. Let $\lambda = (\lambda_s: s = 1, 2, \ldots, s_0)$, where each stratum probability $\lambda_s = P(A \in R_s)$ is assumed known and positive. (We could have defined A to be a more general random object, but for simplicity we only consider A as just a random variable.) In Example 1 the stratification variable A is the randomly chosen scenario, its support $R = \{1, 2, 3, 4\}$ decomposes into strata $R_s = \{s\}$ for $s = 1, 2, 3, 4$, and, e.g., $\lambda_4 = 0.99938 \times 0.0019$. Let

$\boldsymbol{X}_{(s)}$ be an $\mathfrak{S}$-valued random object having the conditional distribution of the basic random object $\boldsymbol{X}$ in (2) given $A \in R_s$. Assume that for each scenario s, we can generate observations of $\boldsymbol{X}_{(s)}$, e.g., via a computer code with random inputs. Then for the failure function $\mathfrak{d}$ in (2), let

$$\mathcal{I}_{(s)} = \mathfrak{d}(\boldsymbol{X}_{(s)}), \tag{15}$$

which has the conditional distribution of the failure indicator $\mathcal{I}$ in (2) given $A \in R_s$. The failure probability in (3) satisfies

$$\theta = \sum_{s=1}^{s_0} P(A \in R_s)P(\mathcal{I} = 1 | A \in R_s) = \sum_{s=1}^{s_0} \lambda_s \theta_{(s)} \tag{16}$$

by the law of total probability, where

$$\theta_{(s)} = P(\mathcal{I}_{(s)} = 1) = E[\mathcal{I}_{(s)}] \tag{17}$$

is the (conditional) failure probability for scenario s. In (16), each λ_s is known but $\theta_{(s)}$ is not, and we use some form of simulation to estimate each $\theta_{(s)}$, which are combined as in (16) to obtain an estimate of θ. (A RISMC study may also be interested in the value of each $\theta_{(s)}$.)

Specifically, we implement SS with overall sample size n by letting $n_s = \eta_s n$ be the sample size allocated to scenario s, where η_s, $s = 1, 2, \ldots, s_0$, are user-specified positive constants satisfying $\sum_{s=1}^{s_0} \eta_s = 1$. We call $\eta = (\eta_s: s = 1, 2, \ldots, s_0)$ the SS *allocation*. For simplicity, we assume that n_s is an integer; otherwise, let $n_s = \lfloor \eta_s n \rfloor$, where $\lfloor \cdot \rfloor$ denotes the floor function. For each scenario $s = 1, 2, \ldots, s_0$, let $\boldsymbol{X}_{(s),i}$, $i = 1, 2, \ldots, n_s$, be a sample of n_s i.i.d. copies of $\boldsymbol{X}_{(s)}$ in (15), and let $\mathcal{I}_{(s),i} = \mathfrak{d}(\boldsymbol{X}_{(s),i})$ for $\mathfrak{d}$ in (2). Then we estimate the expectation $\theta_{(s)}$ of the response $\mathcal{I}_{(s)}$ by

$$\widehat{\theta}_{(s),\mathrm{SS},\eta}(n) = \frac{1}{n_s} \sum_{i=1}^{n_s} \mathcal{I}_{(s),i} \tag{18}$$

where the subscript SS denotes stratified sampling with simple random sampling applied within each stratum. Then

$$\widehat{\theta}_{\mathrm{SS},\eta}(n) = \sum_{s=1}^{s_0} \lambda_s \widehat{\theta}_{(s),\mathrm{SS},\eta}(n) \tag{19}$$

is the SS estimator of $\theta = \sum_{s=1}^{s_0} \lambda_s \theta_{(s)}$. For each scenario $s = 1, 2, \ldots, s_0$, the estimator $\widehat{\theta}_{(s),\mathrm{SS},\eta}(n)$ satisfies a CLT

$$\frac{\sqrt{n}}{\widehat{\sigma}_{(s),\mathrm{SS},\eta}(n)/\sqrt{\eta_s}} \left[\widehat{\theta}_{(s),\mathrm{SS},\eta}(n) - \theta_{(s)}\right] \Rightarrow \mathcal{N}(0, 1) \quad \text{as } n \to \infty, \tag{20}$$

where $\widehat{\sigma}^2_{(s),\mathrm{SS},\eta}(n) \equiv \widehat{\theta}_{(s),\mathrm{SS},\eta}(n)[1 - \widehat{\theta}_{(s),\mathrm{SS},\eta}(n)]$ consistently estimates

$$\sigma^2_{(s),\mathrm{SS}} \equiv \mathrm{Var}[\mathcal{I}_{(s)}] = \theta_{(s)}[1 - \theta_{(s)}]. \tag{21}$$

(The extra factor $\sqrt{\eta_s}$ appears on the left side of (20) because the scaling is $\sqrt{n}$ but the estimator $\widehat{\theta}_{(s),\mathrm{SS},\eta}(n)$ is based on a sample size of $n_s = \eta_s n$.) Assuming that the s_0 scenarios for SS are simulated independently, we then have that (20) jointly holds for $s = 1, 2, \ldots, s_0$, by Problem 29.2 of [43], so the SS estimator $\widehat{\theta}_{\mathrm{SS},\eta}(n)$ of the overall failure probability θ satisfies the CLT $\sqrt{n}[\widehat{\theta}_{\mathrm{SS},\eta}(n) - \theta]/\widehat{\sigma}_{\mathrm{SS},\eta}(n) \Rightarrow \mathcal{N}(0, 1)$ as $n \to \infty$, where $\widehat{\sigma}^2_{\mathrm{SS},\eta}(n) \equiv \sum_{s=1}^{s_0} \lambda_s^2 \widehat{\sigma}^2_{(s),\mathrm{SS},\eta}(n)/\eta_s$ consistently estimates

$$\sigma^2_{\mathrm{SS},\eta} \equiv \sum_{s=1}^{s_0} \frac{\lambda_s^2 \sigma^2_{(s),\mathrm{SS}}}{\eta_s}; \tag{22}$$

e.g., see p. 215 of [13]. Finally, an asymptotic γ-level UCB for θ when using SS is

$$B_{\mathrm{SS},\eta}(n) = \widehat{\theta}_{\mathrm{SS},\eta}(n) + z_\gamma \frac{\widehat{\sigma}_{\mathrm{SS},\eta}(n)}{\sqrt{n}}, \tag{23}$$

which satisfies (5) (without requiring Assumption A1). If there is only $s_0 = 1$ scenario, then SS reduces to SRS.

When the SS allocation η is chosen so that $\eta = \lambda$, which is called the *proportional allocation*, it is well known (e.g., p. 217 of [13]) that SS guarantees a variance reduction compared to SRS; i.e., $\sigma^2_{\mathrm{SS},\lambda} \le \sigma^2_{\mathrm{SRS}}$. The optimal choice of η to minimize $\sigma^2_{\mathrm{SS},\eta}$ is $\eta^* = (\eta_s^*: s = 1, 2, \ldots, s_0)$ with $\eta_s^* = (\lambda_s \sigma_{(s),\mathrm{SS}})/\sum_{s'=1}^{s_0} (\lambda_{s'} \sigma_{(s'),\mathrm{SS}})$. Note that η^* depends on the unknown $\sigma_{(s),\mathrm{SS}}$ in (21), making it impossible to directly implement the optimal allocation. But one could instead apply a two-stage procedure: the first stage estimates the $\sigma_{(s),\mathrm{SS}}$, which are used to estimate η^*, and the second stage then collects additional observations from the scenarios using the estimated optimal allocation.

When Assumption A1 holds (i.e., failure is defined in terms of q loads and capacities), for the function $\mathfrak{v}$ in A1, let $(\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)}) = \mathfrak{v}(\boldsymbol{X}_{(s)})$, which has the conditional distribution of $(\boldsymbol{L}, \boldsymbol{C})$ given $A \in R_s$, where $\boldsymbol{L}_{(s)} = (L^{[1]}_{(s)}, L^{[2]}_{(s)}, \ldots, L^{[q]}_{(s)})$ and $\boldsymbol{C}_{(s)} = (C^{[1]}_{(s)}, C^{[2]}_{(s)}, \ldots, C^{[q]}_{(s)})$. Let $H_{(s)}$ denote the joint CDF of $(\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)})$ in scenario s, and let $F_{(s)}$ (resp., $G_{(s)}$) be the marginal CDF of the load vector $\boldsymbol{L}_{(s)}$ (resp., capacity vector $\boldsymbol{C}_{(s)}$). Then the failure indicator $\mathcal{I}_{(s)}$ in (15) becomes

$$\mathcal{I}_{(s)} = \mathfrak{d}'(\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)}) \tag{24}$$

for failure function $\mathfrak{d}'$ in A1. To construct the SS estimator in (18) of $\theta_{(s)}$, we use the same i.i.d. observations $\boldsymbol{X}_{(s),i}$ $1 \le i \le n_s$, of $\boldsymbol{X}_{(s)}$ from before. Then set $(\boldsymbol{L}_{(s),i}, \boldsymbol{C}_{(s),i}) = \mathfrak{v}(\boldsymbol{X}_{(s),i}) \sim H_{(s)}$, where $\boldsymbol{L}_{(s),i} = (L^{[1]}_{(s),i}, L^{[2]}_{(s),i}, \ldots, L^{[q]}_{(s),i})$ (resp., $\boldsymbol{C}_{(s),i} = (C^{[1]}_{(s),i}, C^{[2]}_{(s),i}, \ldots, C^{[q]}_{(s),i})$) is an observation of the q loads $(L^{[1]}, L^{[2]}, \ldots, L^{[q]})$ (resp., capacities $(C^{[1]}, C^{[2]}, \ldots, C^{[q]})$) given $A \in R_s$. The estimator $\widehat{\theta}_{(s),\mathrm{SS},\eta}(n)$ in (18) of $\theta_{(s)} = E[\mathcal{I}_{(s)}]$ averages $\mathcal{I}_{(s),i} = \mathfrak{d}'(\boldsymbol{L}_{(s),i}, \boldsymbol{C}_{(s),i})$, $1 \le i \le n_s$, which are i.i.d.

5. Combined SS and Latin hypercube sampling

We now consider combining SS with Latin hypercube sampling to estimate θ. Originally proposed by [48], LHS can be thought of as an efficient way to simultaneously stratify multiple input coordinates, and it reduces variance by producing negatively correlated outputs (under certain regularity conditions). The LHS estimator of the mean of a bounded response function satisfies a CLT [49], and [50] extends this to the setting when the response has a finite absolute third moment. Due to its ease of implementation and broad applicability, LHS is widely used in practice; e.g., citing over 300 references, [51] surveys LHS works, with an emphasis on nuclear engineering.

Applying LHS requires additional problem structure. Let $\mathcal{U}[0, 1]$ denote a uniform distribution on [0, 1], and we will assume that the basic random object can be generated using a fixed number of i.i.d. $\mathcal{U}[0, 1]$ random numbers.

Example 1. (Continued). Let U_1 and U_2 be i.i.d. $\mathcal{U}[0, 1]$. We can then generate $\boldsymbol{X}_{(s)} \sim \mathcal{N}_2(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s)$ through a function $\mathfrak{w}_{(s)}$ of the d_s=2 uniforms that operates as follows. First build vector $\boldsymbol{Z} = (Z_1, Z_2)$ with $Z_i = \Phi^{-1}(U_i)$, where Φ is the $\mathcal{N}(0, 1)$ CDF. Then $\mathfrak{w}_{(s)}(U_1, U_2)$ returns $\boldsymbol{X}_{(s)} = \boldsymbol{\mu}_s + \boldsymbol{\Gamma}_s \boldsymbol{Z}$, which lies in the set $\mathfrak{S} = \Re^2$, where $\boldsymbol{\Gamma}_s \boldsymbol{\Gamma}_s^\top = \boldsymbol{\Sigma}_s$.

More generally, we impose the following condition, where $\overset{\mathcal{D}}{=}$ denotes equality in distribution.

Assumption A2. For each scenario $s = 1, 2, \ldots, s_0$, there is a deterministic function $\mathfrak{w}_{(s)}: [0, 1]^{d_s} \to \mathfrak{S}$ such that if $U_1, U_2, \ldots, U_{d_s}$ are d_s i.i.d. $\mathcal{U}[0, 1]$ random variables, then

$$\mathfrak{w}_{(s)}(U_1, U_2, \ldots, U_{d_s}) \overset{\mathcal{D}}{=} \boldsymbol{X}_{(s)}. \tag{25}$$

In other words, the function $\mathfrak{w}_{(s)}$ takes a fixed number d_s of i.i.d. $\mathcal{U}[0, 1]$ random numbers as inputs, and transforms them into an observation of the basic random object $\boldsymbol{X}_{(s)} \in \mathfrak{S}$ having the correct distribution for scenario s. Consider the special case when $\boldsymbol{X}_{(s)} = (X^{[1]}_{(s)}, X^{[2]}_{(s)}, \ldots, X^{[d_s]}_{(s)})$ is a random vector with independent (but not necessarily identically distributed) entries, where each $X^{[j]}_{(s)}$ has marginal CDF $Q_{(s),j}$. Then we set $\mathfrak{w}_{(s)}(U_1, U_2, \ldots, U_{d_s}) = (Q^{-1}_{(s),1}(U_1), Q^{-1}_{(s),2}(U_2), \ldots, Q^{-1}_{(s),d_s}(U_{d_s}))$, assuming $Q^{-1}_{(s),j}$ can be computed. Assumption A2 also permits the random vector $(X^{[1]}_{(s)}, X^{[2]}_{(s)}, \ldots, X^{[d_s]}_{(s)})$ (or a more general basic random object $\boldsymbol{X}_{(s)}$) to have a dependence structure, and Section 7.2 shows how copulas [42] can be used to model the dependence. While the framework of Assumption A2 is quite

general, it precludes applying variate-generation methods such as acceptance-rejection (e.g., Section II.2b of [14]), in which the number of uniforms used to generate $\boldsymbol{X}_{(s)}$ is random and unbounded. Also, A2 does not allow the space $\mathfrak{S}$ in which $\boldsymbol{X}_{(s)}$ lives to be infinite dimensional, as is the case for some stochastic processes; see Section 2.

Before describing how to employ LHS under Assumption A2 (but without requiring A1), we first explain the case of applying SRS within each scenario s to obtain n_s i.i.d. failure indicators. Arrange i.i.d. $\mathcal{U}[0, 1]$ random numbers $U_{(s),i,j}$, $1 \le i \le n_s$, $1 \le j \le d_s$, in an $n_s \times d_s$ grid, and apply the composition $\mathfrak{d}(\mathfrak{w}_{(s)}(\cdot))$ of (15) and (25) to each row:

$$\begin{array}{lcccccc} \mathcal{I}_{(s),1} & = & \mathfrak{d}(\mathfrak{w}_{(s)}(U_{(s),1,1}, & U_{(s),1,2}, & \ldots, & U_{(s),1,d_s})), \\ \mathcal{I}_{(s),2} & = & \mathfrak{d}(\mathfrak{w}_{(s)}(U_{(s),2,1}, & U_{(s),2,2}, & \ldots, & U_{(s),2,d_s})), \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathcal{I}_{(s),n_s} & = & \mathfrak{d}(\mathfrak{w}_{(s)}(U_{(s),n_s,1}, & U_{(s),n_s,2}, & \ldots, & U_{(s),n_s,d_s})), \end{array} \tag{26}$$

where for each $i = 1, 2, \ldots, n_s$, we have that $\mathcal{I}_{(s),i}$ is the failure indicator for the ith run of scenario s. Because each row $i = 1, 2, \ldots, n_s$, of (26) has d_s i.i.d. $\mathcal{U}[0, 1]$ random numbers, $\mathcal{I}_{(s),i}$ has the correct distribution by (15) and (25). Also, as the rows of uniforms in (26) are independent, we have that $\mathcal{I}_{(s),i}$, $i = 1, 2, \ldots, n_s$, are i.i.d. We can then use them to construct the estimator $\widehat{\theta}_{(s),\mathrm{SS},\eta}(n)$ in (18). Moreover, assuming the grids in (26) across scenarios $s = 1, 2, \ldots, s_0$, are independent, we then apply (19) to obtain the asymptotic UCB in (23).

We now explain how to implement LHS under the framework of Assumption A2 (without requiring A1) to obtain a *dependent* sample of n_s failure indicators. For each scenario $s = 1, 2, \ldots, s_0$, and each input coordinate $j = 1, 2, \ldots, d_s$, in (25), let $\boldsymbol{\pi}_{(s),j} = (\pi_{(s),j}(1), \pi_{(s),j}(2), \ldots, \pi_{(s),j}(n_s))$ be a random permutation of $(1, 2, \ldots, n_s)$; i.e., each of the $n_s!$ permutations of $(1, 2, \ldots, n_s)$ is equally likely, and $\pi_{(s),j}(i)$ is the number to which i is mapped in permutation $\boldsymbol{\pi}_{(s),j}$. Also, let $\boldsymbol{\pi}_{(s),j}$, $j = 1, 2, \ldots, d_s$, be independent random permutations, independent of the i.i.d. $U_{(s),i,j}$ in (26). Then let

$$V_{(s),i,j} = \frac{\pi_{(s),j}(i) - 1 + U_{(s),i,j}}{n_s}, \tag{27}$$

which is easily shown to be $\mathcal{U}[0, 1]$. Next arrange $V_{(s),i,j}$, $1 \le i \le n_s$, $1 \le j \le d_s$, into an $n_s \times d_s$ array, which we call an *LHS grid* (of uniforms), and apply the composition $\mathfrak{d}(\mathfrak{w}_{(s)}(\cdot))$ of (15) and (25) to each row to get

$$\begin{array}{lcccccc} \mathcal{I}_{(s),1} & = & \mathfrak{d}(\mathfrak{w}_{(s)}(V_{(s),1,1}, & V_{(s),1,2}, & \ldots, & V_{(s),1,d_s})), \\ \mathcal{I}_{(s),2} & = & \mathfrak{d}(\mathfrak{w}_{(s)}(V_{(s),2,1}, & V_{(s),2,2}, & \ldots, & V_{(s),2,d_s})), \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathcal{I}_{(s),n_s} & = & \mathfrak{d}(\mathfrak{w}_{(s)}(V_{(s),n_s,1}, & V_{(s),n_s,2}, & \ldots, & V_{(s),n_s,d_s})). \end{array} \tag{28}$$

It is straightforward to show that $V_{(s),i,1}, V_{(s),i,2}, \ldots, V_{(s),i,d_s}$ in each row i of (28) are d_s i.i.d. $\mathcal{U}[0, 1]$ random numbers, so the failure indicator $\mathcal{I}_{(s),i}$ has the correct distribution for scenario s by (15) and (25). But all entries in each column j of the $V_{(s),i,j}$ in (28) share the same permutation $\boldsymbol{\pi}_{(s),j}$, so the n_s responses $\mathcal{I}_{(s),i}$, $i = 1, 2, \ldots, n_s$, in (28) are *dependent*.

Each column $j = 1, 2, \ldots, d_s$, of the $V_{(s),i,j}$ in (28) forms a stratified sample of size n_s of the unit interval partitioned into n_s subintervals, each of length $1/n_s$. Specifically, for each column j, there is exactly one value $V_{(s),i,j}$ in the stratum $[(a-1)/n_s, a/n_s)$ for each $a = 1, 2, \ldots, n_s$, as can be easily seen through (27).

Let SS+LHS denote stratified sampling with LHS used within each scenario. (If SS has only $s_0 = 1$ scenario, then SS+LHS reduces to LHS.) For $\mathcal{I}_{(s),i}$, $i = 1, 2, \ldots, n_s$, from (28), the SS+LHS estimator of $\theta_{(s)}$ in (17) is

$$\widehat{\theta}_{(s),\mathrm{SS+LHS},\eta}(n) = \frac{1}{n_s}\sum_{i=1}^{n_s} \mathcal{I}_{(s),i}, \tag{29}$$

which averages *dependent* terms but is still unbiased because of the linearity of expectation and each $\mathcal{I}_{(s),i} \overset{\mathcal{D}}{=} \mathcal{I}_{(s)}$. The SS+LHS estimator of the overall failure probability $\theta = \sum_{s=1}^{s_0} \lambda_s \theta_{(s)}$ is $\widehat{\theta}_{\mathrm{SS+LHS},\eta}(n) = \sum_{s=1}^{s_0} \lambda_s \widehat{\theta}_{(s),\mathrm{SS+LHS},\eta}(n)$.

As the response $\mathcal{I}_{(s),i}$ in (29) is bounded, the LHS CLT of [49] implies the estimator $\widehat{\theta}_{(s),\mathrm{SS+LHS},\eta}(n)$ of $\theta_{(s)}$ satisfies

$$\frac{\sqrt{n}}{\sigma_{(s),\mathrm{SS+LHS}}/\sqrt{\eta_s}}\left[\widehat{\theta}_{(s),\mathrm{SS+LHS},\eta}(n) - \theta_{(s)}\right] \Rightarrow \mathcal{N}(0, 1) \quad \text{as } n \to \infty \tag{30}$$

for each scenario s, where we will give an expression for the asymptotic variance constant $\sigma^2_{(s),\mathrm{SS+LHS}}$ in Section 5.1. We assume the s_0 scenarios for SS are simulated independently, so (30) jointly holds for $s = 1, 2, \ldots, s_0$, by Problem 29.2 of [43]. Hence, the overall SS+LHS estimator of θ obeys the CLT $\sqrt{n}[\widehat{\theta}_{\mathrm{SS+LHS},\eta}(n) - \theta]/\sigma_{\mathrm{SS+LHS},\eta} \Rightarrow \mathcal{N}(0, 1)$ as $n \to \infty$, where the asymptotic variance is

$$\sigma^2_{\mathrm{SS+LHS},\eta} = \sum_{s=1}^{s_0} \frac{\lambda_s^2 \sigma^2_{(s),\mathrm{SS+LHS}}}{\eta_s}. \tag{31}$$

If Assumption A1 (i.e., failure is defined in terms of q loads and capacities) holds in addition to A2, then we can sample the load and capacity vectors for scenario s through a deterministic function $\mathfrak{w}'_{(s)}$: $[0, 1]^{d_s} \to \mathfrak{R}^{2q}$ defined as

$$\mathfrak{w}'_{(s)}(U_1, U_2, \ldots, U_{d_s}) \equiv \mathfrak{v}(\mathfrak{w}_{(s)}(U_1, U_2, \ldots, U_{d_s})) = (\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)}) \sim H_{(s)} \tag{32}$$

wtih $U_1, U_2, \ldots, U_{d_s}$ as i.i.d. $\mathcal{U}[0, 1]$ random variables, where $\mathfrak{v}$ is from A1, $\mathfrak{w}_{(s)}$ is from A2, and we recall that $H_{(s)}$ is the joint CDF of $(\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)})$. In other words, the function $\mathfrak{w}'_{(s)}$ takes d_s i.i.d. $\mathcal{U}[0, 1]$ random numbers as inputs, and transforms them into an observation of the load and capacity vectors having the correct joint CDF $H_{(s)}$. Specifically, $\mathfrak{w}'_{(s)}$ accomplishes this by first using $\mathfrak{w}_{(s)}$ to generate the basic random object $\boldsymbol{X}_{(s)}$, and then feeds $\boldsymbol{X}_{(s)}$ into $\mathfrak{v}$ to obtain $(\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)})$. (In Example 1, because $\boldsymbol{X}_{(s)} = (L_{(s)}, C_{(s)})$, the function $\mathfrak{v}$ is the identity mapping, so $\mathfrak{w}_{(s)} = \mathfrak{w}'_{(s)}$.) Then for the failure function $\mathfrak{d}'$ in A1, we can generate the failure indicator in (15) and (24) as

$$\mathcal{I}_{(s)} = \mathfrak{d}(\mathfrak{w}_{(s)}(U_1, U_2, \ldots, U_{d_s})) = \mathfrak{d}'(\mathfrak{w}'_{(s)}(U_1, U_2, \ldots, U_{d_s})) = \mathfrak{d}'(\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)}). \tag{33}$$

Also, we compute each $\mathcal{I}_{(s),i}$ in (28) as $\mathcal{I}_{(s),i} = \mathfrak{d}'(\mathfrak{w}'_{(s)}(V_{(s),i,1}, V_{(s),i,2}, \ldots, V_{(s),i,d_s}))$ to obtain the estimator in (29).

5.1. Analyzing the effect of LHS

We now give an expression for $\sigma^2_{(s),\mathrm{SS+LHS}}$ in (30) and (31) under Assumption A2 (without A1). Let $\boldsymbol{U} = (U_1, U_2, \ldots, U_{d_s})$ be a vector of d_s i.i.d. $\mathcal{U}[0, 1]$ variables. The failure indicator in (15) satisfies

$$\mathcal{I}_{(s)} = \mathfrak{d}(\boldsymbol{X}_{(s)}) = \mathfrak{d}(\mathfrak{w}_{(s)}(\boldsymbol{U})) \equiv \mathcal{I}_{(s)}(\boldsymbol{U}) \tag{34}$$

by (25), where we write $\mathcal{I}_{(s)}(\boldsymbol{U})$ to emphasize the dependence of $\mathcal{I}_{(s)}$ on $\boldsymbol{U}$. Next use an analysis-of-variance (ANOVA) decomposition as in [52] to approximate the response function in (34) for scenario s by an *additive* function of the uniform inputs:

$$\mathcal{I}_{(s)}(\boldsymbol{U}) = \theta_{(s)} + \sum_{j=1}^{d_s} \phi_{(s),j}(U_j) + \epsilon_{(s)}(\boldsymbol{U}), \tag{35}$$

where the residual $\epsilon_{(s)}(\boldsymbol{U})$ is defined so the equality in (35) holds, and $\phi_{(s),j}(U_j) = E[\mathcal{I}_{(s)}(U_1, U_2, \ldots, U_{d_s}) | U_j] - \theta_{(s)}$ is a function of only the jth uniform input U_j as the other coordinates have been integrated out by the conditional expectation. Then for $\sigma^2_{(s),\mathrm{SS}}$ defined in (21), [52] and [49] show that the asymptotic variance in (30) satisfies

$$\sigma^2_{(s),\mathrm{SS+LHS}} = \mathrm{Var}[\epsilon_{(s)}(\boldsymbol{U})] = \sigma^2_{(s),\mathrm{SS}} - \sum_{j=1}^{d_s} \mathrm{Var}\left[\phi_{(s),j}(U_j)\right]. \tag{36}$$

Thus, $\sigma^2_{(s),\mathrm{SS+LHS}} \le \sigma^2_{(s),\mathrm{SS}}$ for each scenario s, so LHS removes the variability of the additive part of the response $\mathcal{I}_{(s)}(\boldsymbol{U})$. The next result, which then directly follows from putting (36) in (31) and comparing to (22), establishes that after combining across all scenarios, the asymptotic variance for SS+LHS is no larger than that for SS.

Theorem 1. Under Assumption A2, when both SS and SS+LHS use

the same stratification allocation η, we have that $\sigma^2_{\mathrm{SS+LHS},\eta} \le \sigma^2_{\mathrm{SS},\eta}$, with $\sigma^2_{\mathrm{SS},\eta} - \sigma^2_{\mathrm{SS+LHS},\eta} = \sum_{s=1}^{s_0} (\lambda_s^2/\eta_s) \sum_{j=1}^{d_s} \mathrm{Var}\left[\phi_{(s),j}(U_j)\right] \ge 0$.

In general, LHS can produce substantial variance reduction when the response whose expectation we are estimating is nearly an additive function of the input variables. To see why, we now adapt the explanation in [22, Section 10.3] to our setting, where the response for each scenario s is the indicator function on the left side of (35). Putting the additive representation (35) in (29) when the uniform inputs are now from the LHS grid of $V_{(s),i,j}$ in (28) leads to

$$\begin{aligned}\widehat{\theta}_{(s),\mathrm{SS+LHS},\eta}(n) &= \frac{1}{n_s}\sum_{i=1}^{n_s}\left[\theta_{(s)} + \sum_{j=1}^{d_s}\phi_{(s),j}(V_{(s),i,j}) + \epsilon_{(s)}(\boldsymbol{V}_{(s),i})\right] \\ &= \theta_{(s)} + \sum_{j=1}^{d_s}\frac{1}{n_s}\sum_{i=1}^{n_s}\phi_{(s),j}(V_{(s),i,j}) + \frac{1}{n_s}\sum_{i=1}^{n_s}\epsilon_{(s)}(\boldsymbol{V}_{(s),i}),\end{aligned} \tag{37}$$

where $\boldsymbol{V}_{(s),i} = (V_{(s),i,1}, V_{(s),i,2}, \ldots, V_{(s),i,d_s})$ is from the ith row of (28). Recall that for each input coordinate $j = 1, 2, \ldots, d_s$, the jth column of the LHS grid of inputs in (28) forms a stratified sample from the unit interval partitioned into n_s subintervals, each of length $1/n_s$. Thus, for each input coordinate j in the middle term on the right side of (37), the average $(1/n_s)\sum_{i=1}^{n_s}\phi_{(s),j}(V_{(s),i,j})$ corresponds to a (random) Reimann approximation to the one-dimensional integral $\alpha_{(s),j} \equiv \int_0^1 \phi_{(s),j}(u)\,\mathrm{d}u$. When each conditional expectation $\phi_{(s),j}(\cdot)$ is smooth enough, a Reimann approximation of the one-dimensional integral $\alpha_{(s),j}$ based on n_s points converges much more quickly than the standard Monte Carlo rate of $1/\sqrt{n_s}$ as $n_s \to \infty$. The last term in (37) has error decreasing at rate $1/\sqrt{n_s}$, which thus dominates the error of the middle term in (37). Hence, if the additive approximation in (35) is good in the sense that the residual has small variance, then LHS can do much better than SRS for each scenario. Finally, combining across scenarios leads to SS+LHS having substantially lower variance than SS when the response function for each scenario is nearly additive.

But the response function $\mathcal{I}_{(s)}(\boldsymbol{U})$ in (35) is an indicator, so the additive approximation in (35) may be a poor fit. Hence, SS+LHS may not reduce variance much compared to SS. In Section 6 we consider further applying conditional Monte Carlo, which produces a smoother response function. This can lead to a better additive fit and correspondingly smaller asymptotic variance when applying LHS.

5.2. Combined SS and replicated LHS to construct UCB

While it is possible (e.g., see [49]) to consistently estimate the asymptotic variance $\sigma^2_{(s),\mathrm{SS+LHS}}$ in (36) to construct a UCB for θ, we instead now apply *replicated LHS* [21]. Here we generate $b \ge 2$ independent LHS samples, e.g., b=10, where each of the b LHS samples are as in (28), but with size $m_s = \eta_s n/b$ (i.e., m_s rows) instead of n_s, and n is the overall sample size across all scenarios and replications. We assume for simplicity that each m_s is integer-valued; otherwise, let $m_s = \lfloor \eta_s n/b \rfloor$.

We now give the details of estimating θ by combining SS with rLHS, denoted by SS+rLHS, where we use rLHS within each scenario, and rLHS samples across scenarios are independent. For each scenario $s = 1, 2, \ldots, s_0$, and each replication $r = 1, 2, \ldots, b$, let $\boldsymbol{\pi}^{\langle r\rangle}_{(s),j} = (\pi^{\langle r\rangle}_{(s),j}(1), \pi^{\langle r\rangle}_{(s),j}(2), \ldots, \pi^{\langle r\rangle}_{(s),j}(m_s))$ be a random permutation of $(1, 2, \ldots, m_s)$ for input coordinate $j = 1, 2, \ldots, d_s$. Also, let $\boldsymbol{\pi}^{\langle r\rangle}_{(s),j}$, $1 \le j \le d_s, 1 \le r \le b, 1 \le s \le s_0$, be independent random permutations. Then independently of the permutations, define $m_s \times d_s$ i.i.d. $\mathcal{U}[0, 1]$ random numbers $U^{\langle r\rangle}_{(s),i,j}, 1 \le i \le m_s, 1 \le j \le d_s$, as in the array in (26) but with only m_s rows instead of n_s, where the b grids of i.i.d. uniforms for $r = 1, 2, \ldots, b$, are also independent. For each scenario $s = 1, 2, \ldots, s_0$, and each replication $r = 1, 2, \ldots, b$, let $V^{\langle r\rangle}_{(s),i,j} = [\pi^{\langle r\rangle}_{(s),j}(i) - 1 + U^{\langle r\rangle}_{(s),i,j}]/m_s$, $1 \le i \le m_s$, $1 \le j \le d_s$, which we arrange in an $m_s \times d_s$ LHS grid, and apply the composition $\mathfrak{d}(\mathfrak{w}_{(s)}(\cdot))$ of (15) and (25) to each row to obtain m_s identically distributed but dependent copies of $\mathcal{I}_{(s)}$:

$$\begin{aligned}
\mathcal{I}^{\langle r\rangle}_{(s),1} &= \mathfrak{d}(\mathfrak{w}_{(s)}(V^{\langle r\rangle}_{(s),1,1}, V^{\langle r\rangle}_{(s),1,2}, \ldots, V^{\langle r\rangle}_{(s),1,d_s})),\\
\mathcal{I}^{\langle r\rangle}_{(s),2} &= \mathfrak{d}(\mathfrak{w}_{(s)}(V^{\langle r\rangle}_{(s),2,1}, V^{\langle r\rangle}_{(s),2,2}, \ldots, V^{\langle r\rangle}_{(s),2,d_s})),\\
\vdots \quad &\quad \vdots \qquad\qquad \vdots \qquad\quad \ddots \quad \vdots\\
\mathcal{I}^{\langle r\rangle}_{(s),m_s} &= \mathfrak{d}(\mathfrak{w}_{(s)}(V^{\langle r\rangle}_{(s),m_s,1}, V^{\langle r\rangle}_{(s),m_s,2}, \ldots, V^{\langle r\rangle}_{(s),m_s,d_s})).
\end{aligned} \tag{38}$$

The estimator of the failure probability $\theta_{(s)}$ in (17) for scenario s from replication r is

$$\widehat{\theta}^{\langle r\rangle}_{(s)}(n) = \frac{1}{m_s}\sum_{i=1}^{m_s}\mathcal{I}^{\langle r\rangle}_{(s),1}, \tag{39}$$

and the estimator of the overall failure probability θ in (16) from replication r is

$$\widehat{\theta}^{\langle r\rangle}(n) = \sum_{s=1}^{s_0}\lambda_s\widehat{\theta}^{\langle r\rangle}_{(s)}(n). \tag{40}$$

The final SS+rLHS estimator of θ across all b replications is

$$\widehat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n) = \frac{1}{b}\sum_{r=1}^{b}\widehat{\theta}^{\langle r\rangle}(n). \tag{41}$$

To derive a UCB for θ when employing SS+rLHS with overall sample size n, compute the sample variance of $\widehat{\theta}^{\langle r\rangle}(n)$, $1 \le r \le b$, as $S^2_b(n) = \frac{1}{b-1}\sum_{r=1}^{b}\left[\widehat{\theta}^{\langle r\rangle}(n) - \widehat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n)\right]^2$. Let $\tau_{b-1,\gamma}$ be the upper one-sided γ-level critical point of a Student-t random variable T_{b-1} with $b-1$ degrees of freedom (d.f.); i.e., $\gamma = P(T_{b-1} \le \tau_{b-1,\gamma}) = P(T_{b-1} \ge -\tau_{b-1,\gamma})$ by the symmetry of the T_{b-1} density function. For example, $\tau_{9,0.95} = 1.833$. The next result, whose proof is in Appendix A, includes rLHS as the special case when $s_0 = 1$, which can provide a UCB for $\theta_{(s)}$ for a single scenario s.

Theorem 2. Under Assumption A2, an SS+rLHS asymptotic γ-level UCB for θ is $B_{\mathrm{SS+rLHS},\eta,b}(n) = \widehat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n) + \tau_{b-1,\gamma}S_b(n)/\sqrt{b}$, i.e., $\lim_{n\to\infty}P(B_{\mathrm{SS+rLHS},\eta,b}(n) \ge \theta) = \gamma$ as in (5) for any fixed $b \ge 2$ and SS allocation η.

If Assumption A1 from Section 2.1 also holds (i.e., failure defined in terms of q loads and capacities), (38) has $\mathcal{I}^{\langle r\rangle}_{(s),i} = \mathfrak{d}'(\mathfrak{w}'_{(s)}(V^{\langle r\rangle}_{(s),i,1}, V^{\langle r\rangle}_{(s),i,2}, \ldots, V^{\langle r\rangle}_{(s),i,d_s}))$, $1 \le i \le m_s$, with functions $\mathfrak{d}'$ and $\mathfrak{w}'_{(s)}$ as in Assumption A1 and (32), respectively.

6. Combined SS, conditional Monte Carlo, and LHS

Conditional Monte Carlo (also known as *conditional expectation* or *conditioning*) reduces variance by analytically integrating out some of the variability; see Section V.4 of [14] for an overview of CMC.

Example 1. (Continued). Recall that because the basic random object in scenario s is $\boldsymbol{X}_{(s)} = (L_{(s)}, C_{(s)})$, the function $\mathfrak{w}_{(s)}$ in (25) generating $\boldsymbol{X}_{(s)}$ is the same as the function $\mathfrak{w}'_{(s)}$ in (32) producing the load and capacity. Now further suppose that $\mathfrak{w}_{(s)} = \mathfrak{w}'_{(s)}$ uses two other functions $\mathfrak{l}_{(s)}$: $[0, 1] \to \Re$ and $\mathfrak{c}_{(s)}$: $[0, 1] \to \Re$ to generate the load $L_{(s)}$ and capacity $C_{(s)}$, respectively, in $\boldsymbol{X}_{(s)} \sim \mathcal{N}_2(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s)$. Specifically, for U_1 and U_2 as d_s=2 i.i.d. $\mathcal{U}[0, 1]$ random numbers, we assume that $\mathfrak{w}_{(s)}(U_1, U_2) = (\mathfrak{l}_{(s)}(U_1), \mathfrak{c}_{(s)}(U_2))$, where

$$\begin{aligned}\mathfrak{l}_{(s)}(U_1) &= \mu_{L,s} + \sigma_{L,s}\Phi^{-1}(U_1) = L_{(s)} \sim F_{(s)} \quad \text{and}\\ \mathfrak{c}_{(s)}(U_2) &= \mu_{C,s} + \sigma_{C,s}\Phi^{-1}(U_2) = C_{(s)} \sim G_{(s)}\end{aligned} \tag{42}$$

are the load and capacity, respectively. Note that $\mathfrak{l}_{(s)}$ takes as input $d'_s = 1$ uniform U_1, and $\mathfrak{c}_{(s)}$ has the other $d_s - d'_s = 1$ uniform U_2 as input. As $\mathfrak{l}_{(s)}$ and $\mathfrak{c}_{(s)}$ use different inputs, which are independent, the load $L_{(s)} \sim \mathcal{N}(\mu_{L,s}, \sigma^2_{L,s})$ is *independent* of the capacity $C_{(s)} \sim \mathcal{N}(\mu_{C,s}, \sigma^2_{C,s})$. By taking iterated expectations (e.g., p. 448 of [43]) of the failure indicator $\mathcal{I}_{(s)} = I(L_{(s)} \ge C_{(s)})$, the failure probability in (17) for scenario s satisfies

$$\theta_{(s)} = E[\mathcal{I}_{(s)}] = E[E[\mathcal{I}_{(s)}|\, L_{(s)}]] = E[\mathcal{J}_{(s)}(L_{(s)})], \tag{43}$$

where $\mathcal{J}_{(s)}(L_{(s)}) = E[\mathcal{I}_{(s)} \mid L_{(s)}] = P(C_{(s)} \le L_{(s)} \mid L_{(s)}) = G_{(s)}(L_{(s)})$ (44)

by the independence of $L_{(s)}$ and $C_{(s)}$. Therefore, (43) implies that sampling identically distributed copies of $\mathcal{J}_{(s)}(L_{(s)})$ will produce an unbiased estimator of $\theta_{(s)}$. This is the key idea of CMC (when combined with SS). The advantage of this approach becomes apparent from a variance decomposition (e.g., p. 456 of [43]):

$$\begin{aligned}\mathrm{Var}[\mathcal{I}_{(s)}] &= \mathrm{Var}[E[\mathcal{I}_{(s)}|L_{(s)}]] + E[\mathrm{Var}[\mathcal{I}_{(s)}|L_{(s)}] \\ &\ge \mathrm{Var}[E[\mathcal{I}_{(s)}|L_{(s)}]] = \mathrm{Var}[\mathcal{J}_{(s)}(L_{(s)})].\end{aligned} \quad (45)$$

Thus, sampling $\mathcal{J}_{(s)}(L_{(s)})$ rather than $\mathcal{I}_{(s)}$ yields a variance reduction. But applying the combination SS+CMC here critically relies on being able to generate and compute $\mathcal{J}_{(s)}(L_{(s)})$, which is a function of only $L_{(s)}$ as $C_{(s)}$ has been integrated out by the conditional expectation in (44). Because $G_{(s)}$ is the $\mathcal{N}(\mu_{C,s}, \sigma^2_{C,s})$ CDF, we can use (42) to compute (44) as $\mathcal{J}_{(s)}(L_{(s)}) = \Phi([L_{(s)} - \mu_{C,s}]/\sigma_{C,s}) = \Phi([\sigma_{L,s}\Phi^{-1}(U_1) + \mu_{L,s} - \mu_{C,s}]/\sigma_{C,s})$, which is a function of $L_{(s)}$ (or U_1) but not $C_{(s)}$ (nor U_2). In contrast to the failure indicator $\mathcal{I}_{(s)} = I(L_{(s)} \ge C_{(s)})$, the response $\mathcal{J}_{(s)}(L_{(s)})$ is not binary-valued but rather lies in the interval (0, 1) because it is a conditional probability.

Returning to the general problem, we now explain how to combine CMC with SS to estimate $\theta_{(s)} = E[\mathcal{I}_{(s)}]$ in (17) for each scenario $s = 1, 2, \ldots, s_0$. (Later we will add LHS.) Let $\boldsymbol{M}_{(s)}$ be an $\mathfrak{S}'$-valued conditioning random object for CMC defined on the same probability space as $\boldsymbol{X}_{(s)}$, where $(\mathfrak{S}', \mathcal{S}')$ is a metric space, so each simulation run generates $\boldsymbol{M}_{(s)}$ along with $\boldsymbol{X}_{(s)}$; Example 1 has $\boldsymbol{M}_{(s)} = L_{(s)}$ and $\mathfrak{S}' = \Re$. Substituting $\boldsymbol{M}_{(s)}$ for $L_{(s)}$ in (43) then results in

$$\theta_{(s)} = E[\mathcal{I}_{(s)}] = E[E[\mathcal{I}_{(s)} \mid \boldsymbol{M}_{(s)}]] = E[\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})], \quad (46)$$

where the SS+CMC response function generalizing (44) is the conditional expectation (probability)

$$\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)}) = E[\mathcal{I}_{(s)} \mid \boldsymbol{M}_{(s)}] = P(\mathcal{I}_{(s)} = 1 \mid \boldsymbol{M}_{(s)}) = P(\mathfrak{d}(\boldsymbol{X}_{(s)}) = 1 \mid \boldsymbol{M}_{(s)}) \quad (47)$$

for $\mathfrak{d}$ in (2). By (46), SS+CMC produces an unbiased estimator of $\theta_{(s)}$ by averaging identically distributed copies of $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$, rather than copies of $\mathcal{I}_{(s)}$ as when applying only SS. The response $\mathcal{I}_{(s)}$ in (17) when not applying conditioning is binary, but in general the SS+CMC response $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$ in (47) is not binary but rather lies in [0, 1] because it is a conditional probability. As $\sigma^2_{(s),\mathrm{SS}} = \mathrm{Var}[\mathcal{I}_{(s)}]$ by (21), replacing $L_{(s)}$ with $\boldsymbol{M}_{(s)}$ in (45) leads to

$$\mathrm{Var}[\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})] = \sigma^2_{(s),\mathrm{SS}} - E[\mathrm{Var}[\mathcal{I}_{(s)} \mid \boldsymbol{M}_{(s)}]] \le \sigma^2_{(s),\mathrm{SS}}. \quad (48)$$

Thus, estimating $\theta_{(s)}$ by averaging identically distributed copies of the SS+CMC response $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$ from (47) rather than the SS response $\mathcal{I}_{(s)}$ shows that SS+CMC yields a variance reduction compared to SS for each scenario s.

Applying CMC crucially relies on being able to generate and compute the conditional expectation $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$ in (47), which depends on the particular forms of the basic random object $\boldsymbol{X}_{(s)}$, the failure function $\mathfrak{d}$ in (2), and the conditioning random object $\boldsymbol{M}_{(s)}$. Section 6.3 will show how to compute $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$ for a generalization of Example 1 with $q \ge 1$ loads and capacities when $\boldsymbol{M}_{(s)} = \boldsymbol{L}_{(s)}$, the vector of loads. Other CMC work also assumes the basic random object is a vector of basic variables. [26] conditions on various subsets of the basic variables, assumed to be mutually independent, and combines CMC with antithetic variates. When the basic variables are i.i.d. standard normals, [39] expresses the normal vector in polar coordinates and conditions on the angle, so the conditional expectation integrates out the radius, which is independent of the angle. CMC also applies for stochastic processes; e.g., to analyze a highly reliable system modeled as a continuous-time Markov chain, [53] conditions on its embedded discrete-time chain, and the conditional expectation replaces the exponential holding time in each state visited with its conditional mean.

To implement SS+CMC (which reduces to CMC when SS has only $s_0 = 1$ scenario), for each scenario $s = 1, 2, \ldots, s_0$, we generate $\boldsymbol{M}_{(s),i}$, $1 \le i \le n_s$, as a sample of n_s i.i.d. copies of $\boldsymbol{M}_{(s)}$. The SS+CMC estimator of the failure probability $\theta_{(s)}$ for scenario s is then $\widehat{\theta}_{(s),\mathrm{SS+CMC},\eta}(n) = (1/n_s)\sum_{i=1}^{n_s}\mathcal{J}_{(s)}(\boldsymbol{M}_{(s),i})$, and $\widehat{\theta}_{\mathrm{SS+CMC},\eta}(n) = \sum_{s=1}^{s_0}\lambda_s\widehat{\theta}_{(s),\mathrm{SS+CMC},\eta}(n)$ is the SS+CMC estimator of the overall failure probability θ in (16). To build a UCB for θ, let $\widehat{\sigma}^2_{(s),\mathrm{SS+CMC}}(n) = (1/(n_s - 1))\sum_{i=1}^{n_s}[\mathcal{J}_{(s)}(\boldsymbol{M}_{(s),i}) - \widehat{\theta}_{(s),\mathrm{SS+CMC},\eta}(n)]^2$ be a consistent estimator of $\sigma^2_{(s),\mathrm{SS+CMC}} = \mathrm{Var}[\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})]$, so $\widehat{\sigma}^2_{\mathrm{SS+CMC},\eta}(n) = \sum_{s=1}^{s_0}\lambda_s^2\widehat{\sigma}^2_{(s),\mathrm{SS+CMC}}(n)/\eta_s$ consistently estimates $\sigma^2_{\mathrm{SS+CMC},\eta} = \sum_{s=1}^{s_0}\lambda_s^2\sigma^2_{(s),\mathrm{SS+CMC}}/\eta_s$. Recalling z_γ satisfies $P(\mathcal{N}(0, 1) \le z_\gamma) = \gamma$, we obtain an SS+CMC asymptotic γ-level UCB for θ as $B_{\mathrm{SS+CMC},\eta}(n) = \widehat{\theta}_{\mathrm{SS+CMC},\eta}(n) + z_\gamma\widehat{\sigma}_{\mathrm{SS+CMC},\eta}(n)/\sqrt{n}$; i.e., $\lim_{n\to\infty}P(B_{\mathrm{SS+CMC},\eta}(n) \ge \theta) = \gamma$ as in (5) for any SS allocation η. The SS+CMC UCB requires neither Assumption A1 nor A2.

We now impose Assumption A2 so we can add LHS to SS+CMC. By the first equality in (33), (46) becomes

$$\theta_{(s)} = E[\mathcal{I}_{(s)}] = E[\mathfrak{d}(\mathfrak{w}_{(s)}(U_1, U_2, \ldots, U_{d_s}))] = E[E[\,\mathfrak{d}(\mathfrak{w}_{(s)}(U_1, U_2, \ldots, U_{d_s})) \mid \boldsymbol{M}_{(s)}\,]]. \quad (49)$$

We assume the conditioning random object $\boldsymbol{M}_{(s)}$ satisfies the following:

Assumption A3. For $\boldsymbol{M}_{(s)}$ an $\mathfrak{S}'$-valued conditioning random object, there exist functions $\mathfrak{w}^*_{(s)}$: $[0, 1]^{d_s} \to \mathfrak{S} \times \mathfrak{S}'$ and $\mathfrak{m}_{(s)}$: $[0, 1]^{d'_s} \to \mathfrak{S}'$ with $1 \le d'_s \le d_s$ such that for $\mathfrak{w}_{(s)}$: $[0, 1]^{d_s} \to \mathfrak{S}$ in (25) and $U_1, U_2, \ldots, U_{d_s}$ i.i.d. $\mathcal{U}[0, 1]$,

$$(\boldsymbol{X}_{(s)}, \boldsymbol{M}_{(s)}) \overset{\mathcal{D}}{=} \mathfrak{w}^*_{(s)}(U_1, U_2, \ldots, U_{d_s}) = (\mathfrak{w}_{(s)}(U_1, U_2, \ldots, U_{d_s}), \mathfrak{m}_{(s)}(U_1, U_2, \ldots, U_{d'_s})). \quad (50)$$

In other words, the function $\mathfrak{w}^*_{(s)}$ uses the *same* d_s i.i.d. uniforms to generate both the $\mathfrak{S}$-valued basic random object $\boldsymbol{X}_{(s)}$ and the $\mathfrak{S}'$-valued conditioning random object $\boldsymbol{M}_{(s)}$, so they are dependent. Moreover, while $\mathfrak{w}_{(s)}$ employs all d_s uniforms to output $\boldsymbol{X}_{(s)}$, function $\mathfrak{m}_{(s)}$ builds $\boldsymbol{M}_{(s)}$ from just the first $d'_s \le d_s$ of them, which may require relabeling the inputs. Example 1 considers q=1 criterion with $\boldsymbol{X}_{(s)} = (L_{(s)}, C_{(s)}) \in \mathfrak{S} = \Re^2$, so $\boldsymbol{X}_{(s)}$ is a random *vector*. The conditioning random object in Example 1 is $\boldsymbol{M}_{(s)} = L_{(s)}$, which is a random *variable* lying in $\mathfrak{S}' = \Re$ and is generated by function $\mathfrak{m}_{(s)} = \mathfrak{l}_{(s)}$ in (42) from just $d'_s = 1$ uniform U_1 out of the d_s=2 uniforms that $\mathfrak{w}_{(s)}$ uses to produce $\boldsymbol{X}_{(s)}$; the other uniform U_2 is used by function $\mathfrak{c}_{(s)}$ in (42) to sample $C_{(s)}$. As in Example 1, the spaces $\mathfrak{S}'$ and $\mathfrak{S}$ in A3 can differ in general, so $\boldsymbol{M}_{(s)}$ and $\boldsymbol{X}_{(s)}$ may be different types of random objects, e.g., random vectors of unequal dimensions.

We next develop the estimator of θ when applying SS+CMC+LHS (abbreviated SCL). (If SS has only $s_0 = 1$ scenario, then SCL reduces to CMC+LHS.) By (46), the average of identically distributed copies of $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$, which is also the SS+CMC+LHS response, will be an unbiased estimator of $\theta_{(s)}$ in (17). Thus, we first generate an LHS grid of uniforms $V_{(s),i,j}$ (with dependent rows) as in (28) for each scenario $s = 1, 2, \ldots, s_0$, but with d'_s columns instead of d_s because computing (47) requires only the conditioning random object $\boldsymbol{M}_{(s)}$ in (50), which uses d'_s uniforms rather than d_s that $\mathcal{I}_{(s)} = \mathfrak{d}(\boldsymbol{X}_{(s)})$ does in (25). Then apply the function $\mathfrak{m}_{(s)}$ in (50) to each of the n_s rows of the LHS grid to obtain n_s dependent but identically distributed copies $\boldsymbol{M}_{(s),i} = \mathfrak{m}_{(s)}(V_{(s),i,1}, V_{(s),i,2}, \ldots, V_{(s),i,d'_s})$, $1 \le i \le n_s$, of the conditioning random object $\boldsymbol{M}_{(s)}$. Next compute an (unbiased) estimator of $\theta_{(s)}$ as

$$\widehat{\theta}_{(s),\mathrm{SCL},\eta}(n) = \frac{1}{n_s}\sum_{i=1}^{n_s}\mathcal{J}_{(s)}(\boldsymbol{M}_{(s),i}), \quad (51)$$

which is the average of dependent copies of non-binary values, in general. The SS+CMC+LHS estimator of the overall failure probability $\theta = \sum_{s=1}^{s_0}\lambda_s\theta_{(s)}$ is $\widehat{\theta}_{\mathrm{SCL},\eta}(n) = \sum_{s=1}^{s_0}\lambda_s\widehat{\theta}_{(s),\mathrm{SCL},\eta}(n)$.

As the response function $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$ in (47) inherits the boundedness of the binary-valued $\mathcal{I}_{(s)}$, the estimator $\widehat{\theta}_{(s),\mathrm{SCL},\eta}(n)$ of $\theta_{(s)}$ for each scenario s satisfies the LHS CLT of [49] when Assumptions A2 and A3 hold:

$$\frac{\sqrt{n}}{\sigma_{(s),\mathrm{SCL}}/\sqrt{\eta_s}}[\widehat{\theta}_{(s),\mathrm{SCL},\eta}(n) - \theta_{(s)}] \Rightarrow \mathcal{N}(0, 1) \quad \text{as } n \to \infty, \quad (52)$$

where we will give an expression for the asymptotic variance constant $\sigma^2_{(s),\mathrm{SCL}}$ in Section 6.1. We assume the s_0 scenarios for SS are simulated independently, so the overall SS+CMC+LHS estimator of θ satisfies the CLT $\sqrt{n}\,[\widehat{\theta}_{\mathrm{SCL},\eta}(n) - \theta]/\sigma_{\mathrm{SCL},\eta} \Rightarrow \mathcal{N}(0, 1)$ as $n \to \infty$, where the asymptotic variance is

$$\sigma^2_{\mathrm{SCL},\eta} = \sum_{s=1}^{s_0} \frac{\lambda_s^2 \sigma^2_{(s),\mathrm{SCL}}}{\eta_s}. \tag{53}$$

6.1. Analyzing the effect of CMC on LHS

Recall that when not applying CMC, (36) shows that for each scenario s, LHS removes the variability of the additive part of the SS indicator response in (35). When we further incorporate CMC, generating the conditioning random object $\boldsymbol{M}_{(s)}$ in (50) requires a vector $\boldsymbol{U}' = (U_1, U_2, \ldots, U_{d'_s})$ of d'_s i.i.d. $\mathcal{U}[0, 1]$ random numbers. To analyze the SS+CMC+LHS response $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$ in (47) for estimating $\theta_{(s)}$, we emphasize the dependence of $\boldsymbol{M}_{(s)}$ on the vector $\boldsymbol{U}'$ and approximate the response as an additive function of the uniform inputs $\boldsymbol{U}'$ based on an ANOVA decomposition:

$$\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)}(\boldsymbol{U}')) = \theta_{(s)} + \sum_{j=1}^{d'_s} \left(E[\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)}(\boldsymbol{U}'))|\, U_j] - \theta_{(s)}\right) + \bar{\epsilon}_{(s)}(\boldsymbol{U}'), \tag{54}$$

where the residual $\bar{\epsilon}_{(s)}(\boldsymbol{U}')$ is defined so that the above equality holds. Then as shown in [52] and [49], we have that the asymptotic variance in (52) satisfies $\sigma^2_{(s),\mathrm{SCL}} = \mathrm{Var}[\bar{\epsilon}_{(s)}(\boldsymbol{U}')]$. The following result, proven in Appendix B, establishes that the asymptotic variance $\sigma^2_{\mathrm{SCL},\eta}$ in (53) for SS+CMC+LHS is no greater than $\sigma^2_{\mathrm{SS+LHS},\eta}$ in (31) for SS+LHS.

Theorem 3. Under Assumptions A2 and A3, when both SCL and SS +LHS use the same stratification allocation η, we have that $\sigma^2_{\mathrm{SCL},\eta} \le \sigma^2_{\mathrm{SS+LHS},\eta}$, where, for $\epsilon_{(s)}(U_1, U_2, \ldots, U_{d_s})$ defined in (35),

$$\sigma^2_{\mathrm{SS+LHS},\eta} - \sigma^2_{\mathrm{SCL},\eta} = \sum_{s=1}^{s_0} \frac{\lambda_s^2}{\eta_s} E\Big[\mathrm{Var}[\,\epsilon_{(s)}(U_1, U_2, \ldots, U_{d_s})|\, \boldsymbol{M}_{(s)}(U_1, U_2, \ldots, U_{d'_s})]\Big] \ge 0. \tag{55}$$

The key point here is that the conditioning in (47) leads to $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)}(\boldsymbol{U}'))$ in (54) being a "smoother" response function than the indicator $\mathcal{I}_{(s)}(\boldsymbol{U})$ in (34). As a consequence, the additive approximation (54) can be a better fit than that in (35). This can result in SS+CMC +LHS having much lower asymptotic variance than SS+LHS, as we will see in the numerical results of Section 7. The idea of smoothing a response function through conditioning has also been fruitfully applied in other contexts, such as gradient estimation [54].

6.2. Combined SS+CMC+rLHS to construct UCB

We now provide details of the combined SS+CMC+rLHS approach to construct a UCB for θ under Assumptions A2 and A3. For each scenario $s = 1, 2, \ldots, s_0$, and each replication $r = 1, 2, \ldots, b$, with $b \ge 2$, let $V^{\langle r\rangle}_{(s),i,j}$, $1 \le i \le m_s$, $1 \le j \le d'_s$, be an $m_s \times d'_s$ LHS grid of uniforms (with dependent rows) as in (38), but with d'_s columns instead of d_s. LHS grids across the s_0 scenarios and b replications are independent. Then apply function $\mathfrak{m}_{(s)}$ from (50) to each row of the LHS grid to get m_s dependent but identically distributed copies of the conditioning random object:

$$\begin{aligned}
\boldsymbol{M}^{\langle r\rangle}_{(s),1} &= \mathfrak{m}_{(s)}(V^{\langle r\rangle}_{(s),1,1},\ V^{\langle r\rangle}_{(s),1,2},\ \ldots,\ V^{\langle r\rangle}_{(s),1,d'_s}),\\
\boldsymbol{M}^{\langle r\rangle}_{(s),2} &= \mathfrak{m}_{(s)}(V^{\langle r\rangle}_{(s),2,1},\ V^{\langle r\rangle}_{(s),2,2},\ \ldots,\ V^{\langle r\rangle}_{(s),2,d'_s}),\\
&\vdots\quad\ \vdots\qquad\ \vdots\qquad\ \ddots\quad\ \vdots\\
\boldsymbol{M}^{\langle r\rangle}_{(s),m_s} &= \mathfrak{m}_{(s)}(V^{\langle r\rangle}_{(s),m_s,1},\ V^{\langle r\rangle}_{(s),m_s,2},\ \ldots,\ V^{\langle r\rangle}_{(s),m_s,d'_s}).
\end{aligned} \tag{56}$$

For the rth replicated LHS for scenario s, define the estimator of the failure probability $\theta_{(s)}$ in (46) as

$$\widetilde{\theta}^{\langle r\rangle}_{(s)}(n) = \frac{1}{m_s}\sum_{i=1}^{m_s} \mathcal{J}_{(s)}(\boldsymbol{M}^{\langle r\rangle}_{(s),i}) \tag{57}$$

for the response function $\mathcal{J}_{(s)}$ in (47). Compared to the SS+rLHS estimator without CMC in (39), the SS+CMC+rLHS estimator $\widetilde{\theta}^{\langle r\rangle}_{(s)}(n)$ replaces the indicator in (39) with its conditional expectation $\mathcal{J}_{(s)}$. The estimator of the overall failure probability θ in (16) from replication r is $\widetilde{\theta}^{\langle r\rangle}(n) = \sum_{s=1}^{s_0} \lambda_s \widetilde{\theta}^{\langle r\rangle}_{(s)}(n)$. The final estimator of θ across all b replications is $\widehat{\theta}_{\mathrm{SCrL},\eta,b}(n) = (1/b)\sum_{r=1}^{b} \widetilde{\theta}^{\langle r\rangle}(n)$. Also, compute the sample variance of $\widetilde{\theta}^{\langle r\rangle}(n)$, $1 \le r \le b$, as $\widetilde{S}^2_b(n) = \frac{1}{b-1}\sum_{r=1}^{b}\left[\widetilde{\theta}^{\langle r\rangle}(n) - \widehat{\theta}_{\mathrm{SCrL},\eta,b}(n)\right]^2$. The next result, proven in Appendix C, provides an asymptotically valid SS+CMC +rLHS UCB, which reduces to a UCB for CMC+rLHS when SS has only $s_0 = 1$ scenario.

Theorem 4. Under Assumptions A2 and A3, an SS+CMC+rLHS asymptotic γ-level UCB for θ is $B_{\mathrm{SCrL},\eta,b}(n) = \widehat{\theta}_{\mathrm{SCrL},\eta,b}(n) + \tau_{b-1,\gamma}\,\widetilde{S}_b(n)/\sqrt{b}$, *i.e.*, $\lim_{n\to\infty} P(B_{\mathrm{SCrL},\eta,b}(n) \ge \theta) = \gamma$ as in (5) for any fixed $b \ge 2$ and SS allocation η.

6.3. SS+CMC+LHS for failures defined in terms of loads and capacities, with loads independent of capacities

As noted in Section 6, the key to applying CMC is being able to sample and compute the response $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$ in (47), which depends on how we define the basic random object $\boldsymbol{X}_{(s)}$, failure function $\mathfrak{d}$, and conditioning random object $\boldsymbol{M}_{(s)}$. We now show how to do this for a generalization of Example 1 when Assumption A1 holds (i.e., failure is defined as a function $\mathfrak{d}'$ of $q \ge 1$ loads and capacities) and $\boldsymbol{M}_{(s)}$ is the load vector. We first augment Assumptions A2 and A3.

Assumption A4. For each scenario $1 \le s \le s_0$, the conditioning random object in Assumption A3 is $\boldsymbol{M}_{(s)} = \boldsymbol{L}_{(s)}$. The function $\mathfrak{w}'_{(s)}\colon [0, 1]^{d_s} \to \Re^{2q}$ in (32) used to generate the q loads $\boldsymbol{L}_{(s)}$ and capacities $\boldsymbol{C}_{(s)}$ decomposes into functions $\mathfrak{l}_{(s)}\colon [0, 1]^{d_{s,L}} \to \Re^q$ generating loads and $\mathfrak{c}_{(s)}\colon [0, 1]^{d_{s,C}} \to \Re^q$ generating capacities such that $d_{s,L} + d_{s,C} = d_s$ and

$$\mathfrak{w}'_{(s)}(u_1, u_2, \ldots, u_{d_s}) = \left(\mathfrak{l}_{(s)}(u_1, u_2, \ldots, u_{d_{s,L}}),\ \mathfrak{c}_{(s)}(u_{d_{s,L}+1}, u_{d_{s,L}+2}, \ldots, u_{d_{s,L}+d_{s,C}})\right)$$

for $(u_1, u_2, \ldots, u_{d_s}) \in [0, 1]^{d_s}$.

Thus, Assumption A3 has $\mathfrak{m}_{(s)} = \mathfrak{l}_{(s)}$, so $d'_s = d_{s,L}$, $\mathfrak{S}' = \Re^q$, and $\mathfrak{w}^*_{(s)}\colon [0, 1]^{d_s} \to \mathfrak{S} \times \Re^q$ in (50) satisfies $\mathfrak{w}^*_{(s)}(u_1, u_2, \ldots, u_{d_s}) = (\mathfrak{w}_{(s)}(u_1, u_2, \ldots, u_{d_s}), \mathfrak{l}_{(s)}(u_1, u_2, \ldots, u_{d_{s,L}}))$ for $(u_1, u_2, \ldots, u_{d_s}) \in [0, 1]^{d_s}$ and $\mathfrak{w}_{(s)}$ in (25).

Assumption A4 specifies the conditioning random object $\boldsymbol{M}_{(s)}$ in A3 to be the vector $\boldsymbol{L}_{(s)}$ of the $q \ge 1$ criteria's loads. In the setting of A1, (33) specializes A2 to compute the failure indicator $\mathcal{I}_{(s)}$ in terms of $\boldsymbol{L}_{(s)}$ and $\boldsymbol{C}_{(s)}$ output by the function $\mathfrak{w}'_{(s)}$ in (32). Now A4 further stipulates that $\mathfrak{w}'_{(s)}$ decomposes into functions $\mathfrak{l}_{(s)}$ and $\mathfrak{c}_{(s)}$ operating on *disjoint* sets of inputs, where $\mathfrak{l}_{(s)}$ outputs the loads and $\mathfrak{c}_{(s)}$ the capacities. (Example 1 has q=1 and in (42), function $\mathfrak{l}_{(s)}$ has input U_1 and $\mathfrak{c}_{(s)}$ takes input U_2.) Thus, if $U_1, \ldots, U_{d_{s,L}}, U_{d_{s,L}+1}, \ldots, U_{d_{s,L}+d_{s,C}}$ are d_s i.i.d. $\mathcal{U}[0, 1]$ random variables, then

$$\begin{aligned}
\boldsymbol{L}_{(s)} &= \mathfrak{l}_{(s)}(U_1, U_2, \ldots, U_{d_{s,L}}) \sim F_{(s)} \quad \text{and}\\
\boldsymbol{C}_{(s)} &= \mathfrak{c}_{(s)}(U_{d_{s,L}+1}, U_{d_{s,L}+2}, \ldots, U_{d_{s,L}+d_{s,C}}) \sim G_{(s)},
\end{aligned} \tag{58}$$

where $\boldsymbol{L}_{(s)} = (L^{[1]}_{(s)}, L^{[2]}_{(s)}, \ldots, L^{[q]}_{(s)})$ and $\boldsymbol{C}_{(s)} = (C^{[1]}_{(s)}, C^{[2]}_{(s)}, \ldots, C^{[q]}_{(s)})$, and we recall that $F_{(s)}$ and $G_{(s)}$ were defined at the end of Section 4 as the marginal CDFs of the load and capacity vectors, respectively, for scenario s. Consequently, as the two functions $\mathfrak{l}_{(s)}$ and $\mathfrak{c}_{(s)}$ use disjoint sets of i.i.d. arguments, the following holds:

Proposition 1. Assumption A4 implies $\boldsymbol{L}_{(s)}$ is independent of $\boldsymbol{C}_{(s)}$.

Proposition 1 ensures that $\boldsymbol{L}_{(s)}$ is independent of $\boldsymbol{C}_{(s)}$, but each vector

may still have a dependence structure for its own q entries. Without A4, Assumption A2 and $\mathfrak{w}'_{(s)}$ in (32) allow $\boldsymbol{L}_{(s)}$ and $\boldsymbol{C}_{(s)}$ to be dependent.

In structural reliability, loads are sometimes (but not always) assumed to be independent of capacities [20, Chapter 1]. In many NPP PSAs, loads are determined by the way in which a hypothesized accident evolves, whereas the capacities depend on the material properties and manufacturing variability of the components. Hence, it may be reasonable to model loads as independent of capacities, as is assumed in the initial RISMC studies with a single criterion in [3] and [19]. Moreover, [3] and [19] take the capacity distribution $G_{(s)}$ to be the same for all scenarios s, although our development here does not require this.

We next further refine how a failure is defined.

Assumption A5. The failure function $\mathfrak{d}'$: $\mathfrak{R}^{2q} \to \{0, 1\}$ in Assumption A1 has the form

$$\mathfrak{d}'(\boldsymbol{L}, \boldsymbol{C}) = a^{[0]} + \sum_{p=1}^{q} \sum_{1 \le k_1 < k_2 < \cdots < k_p \le q} a^{[k_1,k_2,\ldots,k_p]} I(\cap_{l=1}^{p} \{L^{[k_l]} \ge C^{[k_l]}\}), \tag{59}$$

where $a^{[0]}$ and each $a^{[k_1,k_2,\ldots,k_p]}$ are constant coefficients.

Assumption A5 requires that the failure function $\mathfrak{d}'$ is a *linear* function of the indicators that the load exceeds capacity for each criterion in each nonempty subset of criteria. In Section 2.1, Examples 2–6, which define failure in terms of q loads and capacities, all satisfy Assumption A5. In Example 2 (series system), (7) equals $\mathfrak{d}'$ in (59) with $a^{[0]} = 0$ and $a^{[k_1,k_2,\ldots,k_p]} = (-1)^{p-1}$ for each $1 \le p \le q$. In Example 3 (parallel system), (8) corresponds to (59) with $a^{[0]} = 0$, $a^{[1,2,\ldots,q]} = 1$, and all other $a^{[k_1,k_2,\ldots,k_p]} = 0$. In Example 4 (K-out-of-N:F system), (9) fits into (59) with $a^{[0]} = a^{[k_1,k_2,\ldots,k_p]} = 0$ for each $1 \le p < \mathrm{K}$, and $a^{[k_1,k_2,\ldots,k_p]} = (-1)^{p-\mathrm{K}} \binom{p-1}{\mathrm{K}-1}$ for each $\mathrm{K} \le p \le q$. In Example 5 (series-parallel system), (10) matches (59) with $a^{[q_{m_1}, q_{m_1+1}, \ldots, q_{m_1+1-1}, q_{m_2}, q_{m_2+1}, \ldots, q_{m_2+1-1}, \ldots, q_{m_p}, q_{m_p+1}, \ldots, q_{m_p+1-1}]} = (-1)^{p-1}$ for each $1 \le p \le v$, and all other coefficients $a^{[\cdot]} = 0$. In Example 6, we can show that $\mathfrak{d}'$ in (11) has the form in (59) by expanding the product in (11) and repeatedly using the fact that $I(L^{[k]} \ge C^{[k]}) I(L^{[k']} \ge C^{[k']}) = I(\{L^{[k]} \ge C^{[k]}\} \cap \{L^{[k']} \ge C^{[k']}\})$.

We now exploit the added structure ensured by our assumptions to obtain another representation for the response $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s)})$ in (47), which will be computable after imposing one final condition (A6) below. By Assumption A4 the conditioning random object in (47) is $\boldsymbol{M}_{(s)} = \boldsymbol{L}_{(s)}$, which Proposition 1 guarantees is independent of $\boldsymbol{C}_{(s)}$. Assumptions A1 and A5 specify that $\mathcal{I}_{(s)} = \mathfrak{d}'(\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)})$ by (24) for $\mathfrak{d}'$ in (59), whose linearity implies that (47) becomes

$$\begin{aligned} \mathcal{J}_{(s)}(\boldsymbol{L}_{(s)}) &= E[\mathfrak{d}'(\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)}) | \boldsymbol{L}_{(s)}] \\ &= a^{[0]} + \sum_{p=1}^{q} \sum_{1 \le k_1 < k_2 < \cdots < k_p \le q} a^{[k_1,k_2,\ldots,k_p]} E[I(\cap_{l=1}^{p} \{L_{(s)}^{[k_l]} \ge C_{(s)}^{[k_l]}\}) | \boldsymbol{L}_{(s)}]. \end{aligned} \tag{60}$$

Now let $G_{(s)}^{[k_1,k_2,\ldots,k_p]}$ be the marginal joint CDF of the capacities $(C_{(s)}^{[k_1]}, C_{(s)}^{[k_2]}, \ldots, C_{(s)}^{[k_p]})$ for scenario $s = 1, 2, \ldots, s_0$; i.e.,

$$G_{(s)}^{[k_1,k_2,\ldots,k_p]}(x_1, x_2, \ldots, x_p) = P(C_{(s)}^{[k_1]} \le x_1, C_{(s)}^{[k_2]} \le x_2, \ldots, C_{(s)}^{[k_p]} \le x_p) \tag{61}$$

for any constants $x_1, x_2, \ldots, x_p$. (In the special case when capacities across criteria are independent, (61) simplifies to $\prod_{l=1}^{p} G_{(s)}^{[k_l]}(x_l)$, but we do not require this.) Then, analogously to how we obtained (44) in the special case when q=1 in Example 1, we can rewrite the conditional expectation in the right side of (60) as

$$\begin{aligned} E[I(\cap_{l=1}^{p} \{L_{(s)}^{[k_l]} \ge C_{(s)}^{[k_l]}\}) | \boldsymbol{L}_{(s)}] &= P(\cap_{l=1}^{p} \{C_{(s)}^{[k_l]} \le L_{(s)}^{[k_l]}\} | \boldsymbol{L}_{(s)}) \\ &= G_{(s)}^{[k_1,k_2,\ldots,k_p]}(L_{(s)}^{[k_1]}, L_{(s)}^{[k_2]}, \ldots, L_{(s)}^{[k_p]}) \end{aligned} \tag{62}$$

because $\boldsymbol{L}_{(s)}$ is independent of $\boldsymbol{C}_{(s)}$ by Proposition 1. We next require the following to ensure (62) is computable.

Assumption A6. The marginal capacity CDF $G_{(s)}^{[k_1,k_2,\ldots,k_p]}$ can be computed analytically or numerically for each $1 \le k_1 < k_2 < \cdots < k_p \le q$ such that $a^{[k_1,k_2,\ldots,k_p]} \ne 0$ in (59), each $1 \le p \le q$, and each scenario $s = 1, 2, \ldots, s_0$.

Finally putting (62) into (60) yields the following.

Theorem 5. If Assumptions A1–A6 hold, then the SS+CMC+LHS response function in (47) is computable as

$$\mathcal{J}_{(s)}(\boldsymbol{L}_{(s)}) = a^{[0]} + \sum_{p=1}^{q} \sum_{1 \le k_1 < k_2 < \cdots < k_p \le q} a^{[k_1,k_2,\ldots,k_p]} G_{(s)}^{[k_1,k_2,\ldots,k_p]}(L_{(s)}^{[k_1]}, L_{(s)}^{[k_2]}, \ldots, L_{(s)}^{[k_p]}). \tag{63}$$

To compute the estimator in (51), we simply replace $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s),i})$ with $\mathcal{J}_{(s)}(\boldsymbol{L}_{(s),i})$ using (63), where $\boldsymbol{L}_{(s),i} = \mathfrak{l}_{(s)}(V_{(s),i,1}, V_{(s),i,2}, \ldots V_{(s),i,d_{s,L}})$ with $V_{(s),i,j}$ from the LHS grid of uniforms in (28) but with $d_{s,L}$ columns instead of d_s, and function $\mathfrak{l}_{(s)}$ is from Assumption A4 and (58). Similarly, we compute the estimator in (57) by replacing $\mathcal{J}_{(s)}(\boldsymbol{M}_{(s),i}^{\langle r \rangle})$ with $\mathcal{J}_{(s)}(\boldsymbol{L}_{(s),i}^{\langle r \rangle})$ using (63) with $\boldsymbol{L}_{(s),i}^{\langle r \rangle} = \mathfrak{l}_{(s)}(V_{(s),i,1}^{\langle r \rangle}, V_{(s),i,2}^{\langle r \rangle}, \ldots V_{(s),i,d_{s,L}}^{\langle r \rangle})$, where each $V_{(s),i,j}^{\langle r \rangle}$ is from the LHS grid of uniforms in (56) with $d'_s = d_{s,L}$ columns.

The number of terms in (63) with nonzero coefficient $a^{[k_1,k_2,\ldots,k_p]}$ can grow exponentially in the number q of criteria. For example, a series system (Example 2 in Section 2.1) uses $\mathfrak{d}'$ in (7), so (63) has $2^q - 1$ nonzero summands; this can practically limit the number of criteria that can be considered. In contrast, a parallel system (Example 3) has $\mathfrak{d}'$ in (8), so (63) has only one nonzero summand; thus, as long as $G_{(s)}^{[1,2,\ldots,q]}$ can be computed efficiently, we are not restricted to small q. If the number of summands in (63) is overwhelming, we may instead obtain an estimator of an approximation to $\theta_{(s)}$ by truncating the first sum in (63) at some upper index $q' < q$; see [20, Section 5.6.2] for related ideas.

7. Numerical experiments

We next explore the effectiveness of the VRTs considered in this paper through numerical experiments with synthetic RISMC problems, which satisfy Assumptions A1–A6 so that the setting of Section 6.3 holds. Although particular aspects of the models we consider are motivated by previous work on nuclear PSAs, including [10,3,19], we acknowledge that our models are artificial. The purpose of our experiments is not to carry out an actual nuclear PSA, but rather to provide a proof of concept that our Monte Carlo methods can produce substantial variance reductions.

In an actual RISMC analysis of a hypothesized NPP accident, as in [3] and [19], a nuclear-specific computer code models the progression of the event, and each code run produces observations of the q criteria's loads and capacities. While [3] and [19] consider only q=1 criterion (PCT), our four synthetic models have q=3 criteria: PCT, CWO, and MLO (see Section 1). The results in Section 7.4 use the failure function $\mathfrak{d}'$ in Assumption A5 for the series system (7), so the system fails when any criterion's load exceeds its capacity, but Section 7.5 also considers a parallel system, with $\mathfrak{d}'$ in (8). Our Models 1–4 differ in the choice of PCT load distributions and the dependence structures among the output variables from each simulation run. We base our models on the event tree in Fig. 1, originally from [3], with $s_0 = 4$ scenarios, where the probability λ_s of each scenario $1 \le s \le 4$ occurring is given by $\lambda_1 = 0.99938 \times 0.9981 \times 0.919$, $\lambda_2 = 0.00062$, $\lambda_3 = 0.99938 \times 0.9981 \times 0.081$, and $\lambda_4 = 0.99938 \times 0.0019$. The basic random object for scenario s in (15) is $\boldsymbol{X}_{(s)} = (\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)})$. We next describe the marginals (Section 7.1) and dependence structure (Section 7.2) for the loads $(L_s^{[1]}, L_{(s)}^{[2]}, L_{(s)}^{[3]})$ and capacities $(C_s^{[1]}, C_{(s)}^{[2]}, C_{(s)}^{[3]})$, and how $\mathfrak{w}'_{(s)}$ in A4 generates them (Section 7.3).

7.1. Marginal distributions for loads and for capacities

Let $F_{(s)}^{[1]}, F_{(s)}^{[2]}, F_{(s)}^{[3]}$ be the marginal CDFs of the loads for PCT, CWO, MLO, respectively, in scenario s. To select the marginal CDFs, we extracted data from Figures 1 and 10 of [10], which are scatter plots of

these criteria's load data (sample sizes ≈180) output from a computer code of a hypothesized LOCA. Based on the shape of the histogram of each criterion's load data, we chose a lognormal for the marginal CDF of the PCT load, and marginal Weibull distributions for both CWO and MLO loads.

For Models 1 and 2, the lognormal PCT load CDF $F_{(s)}^{[1]}$ for scenario s = 1, 2, 3, 4, corresponds to exponentiating a normal random variable with mean $\mu_{(s)}^{[1]} = 7.35 + 0.01\,s$ and standard deviation $\sigma_{(s)}^{[1]} = 0.14 + 0.01\,s$, where the parameter values for scenario s=1 are the maximum likelihood estimates (MLEs; Section 7.1 of [44]) from the PCT load data of [10]. To assess the appropriateness of our choice $F_{(s)}^{[1]}$ for scenario s=1, we applied the Kolmogorov-Smirnov (KS) and chi-square goodness-of-fit tests (e.g., Section 5.7 of [44]) with the data from [10], which yielded p-values of 0.45 and 0.39, respectively; thus, at significance level 0.05, we do not reject the CDF $F_{(s)}^{[1]}$ for scenario s=1. For Models 3 and 4, we changed the parameters of the PCT load's lognormal CDF for scenario s to $\mu_{(s)}^{[1]} = 7.4 + 0.1\,s$ and $\sigma_{(s)}^{[1]} = 0.01 + 0.01\,s$, which are used in [24].

For Models 1–4, the CWO load for scenario s has marginal Weibull CDF $F_{(s)}^{[2]}(x) = 1 - \exp\{-(x/\alpha_{(s)}^{[2]})^{\beta_{(s)}^{[2]}}\}$ for $x \geq 0$, where $\alpha_{(s)}^{[2]} = 0.010 + 0.005\,s$ and $\beta_{(s)}^{[2]} = 0.90 + 0.02\,s$ are the scale and shape parameters, respectively. For scenario s=1, the parameter values are the MLEs from the CWO data in [10], and the KS and chi-square tests do not reject this choice of CDF (p-values 0.22 and 0.61, respectively).

For the MLO load, its marginal Weibull CDF $F_{(s)}^{[3]}$ for scenario s in Models 1–4 has parameters $\alpha_{(s)}^{[3]} = 0.35 + 0.3\,s$ and $\beta_{(s)}^{[3]} = 0.82 + 0.03\,s$. The parameter values for scenario s=1 are the MLEs from the MLO load data in [10], and our CDF choice for s=1 is not rejected by the KS and chi-square tests (p-values 0.82 and 0.08, respectively).

Although the load CDFs depend on the scenario s, we assume that the capacity CDFs remain the same for all s, as is the case in [3] and [19], so $G_{(s)}^{[k]} = G^{[k]}$ for each criterion k = 1, 2, 3.

As in [3,19], we take the PCT capacity $C^{[1]}$ to have a triangular CDF $G^{[1]}$ with support $[a^{[1]}, b^{[1]}] = [1800, 2600]$ and mode $c^{[1]} = 2200$. The mode $c^{[1]}$ is the fixed limit prescribed by the NRC [4] for the PCT criterion, and the support parameters are symmetric and separated from the mode by approximately 20%.

For the CWO and MLO capacities $C^{[2]}$ and $C^{[3]}$, respectively, we also assume marginal triangular distributions with parameters following a similar approach as the one used for PCT capacity. Because the NRC specifies a 1% fixed limit for CWO (see Section 1), we set the CWO capacity CDF $G^{[2]}$ as triangular with support $[a^{[2]}, b^{[2]}] = [0.8, 1.2]$ and mode $c^{[2]} = 1$. Finally, the NRC fixed limit for MLO is 17%, so the MLO capacity CDF $G^{[3]}$ has support $[a^{[3]}, b^{[3]}] = [13.6, 20.4]$ and mode $c^{[3]} = 17$.

7.2. Specifying dependence structures through copulas

So far we have discussed the marginal CDFs of loads and capacities, but we further need to specify their joint distributions $F_{(s)}$ and G. In an actual NPP PSA with a computer code, each code run produces a vector of the loads for PCT, CWO, and MLO, so the q=3 criteria's loads have a particular stochastic dependence. For example, if the PCT load is unusually high in one run, then the same is likely true for the CWO and MLO. But as our numerical experiments do not use such a computer code, we require a mechanism to generate observations of the load vector with a reasonable dependence structure. Moreover, the criteria's capacities in a run should also be statistically dependent. But as Assumption A4 implies by Proposition 1, we have that loads are stochastically independent of the capacities.

We specify the dependence among the loads and among the capacities via *copulas* [42]. While (Pearson) correlation measures only *linear* dependence, a copula fully characterizes the dependence structure for given marginal CDFs. To define a copula, we start with a random vector $\boldsymbol{W} = (W^{[1]}, W^{[2]}, \ldots, W^{[q]})$ in which each marginal is $\mathcal{U}[0, 1]$, where there may be dependence among the entries in $\boldsymbol{W}$. Then a copula K: $[0, 1]^q \to [0, 1]$ is the joint CDF of $\boldsymbol{W}$; i.e.,

$$K(w_1, w_2, \ldots, w_q) = P(W^{[1]} \leq w_1, W^{[2]} \leq w_2, \ldots, W^{[q]} \leq w_q), \quad \text{for } (w_1, w_w, \ldots, w_q) \in [0, 1]^q. \tag{64}$$

Now suppose we want to generate a load or capacity vector, generically denoted by $\boldsymbol{Y} = (Y^{[1]}, Y^{[2]}, \ldots, Y^{[q]})$, with joint CDF F_0, where each $Y^{[i]}$ has a marginal CDF F_i as specified in Section 7.1. By Sklar's theorem [55, Theorem 5.3], there exists a copula K such that $F_0(y_1, y_2, \ldots, y_q) = K(F_1(y_1), F_2(y_2), \ldots, F_q(y_q))$ and $K(w_1, w_2, \ldots, w_q) = F_0(F_1^{-1}(w_1), F_2^{-1}(w_2), \ldots, F_q^{-1}(w_q))$; K is unique when each marginal F_i is continuous. Thus, any joint CDF has a copula, and we may combine a copula and marginal CDFs to construct a joint CDF. Hence, to generate $\boldsymbol{Y}$ using $\boldsymbol{W} \sim K$ in (64), where K is the copula of F_0, note that because each $W^{[i]} \sim \mathcal{U}[0, 1]$, we simply set

$$Y^{[i]} = F_i^{-1}(W^{[i]}) \sim F_i, \text{ for each } i. \tag{65}$$

If the marginals F_i have been specified (e.g., as in Section 7.1), the next issue is how to construct a copula K (appropriate for the problem) and to sample $\boldsymbol{W} \sim K$. One approach begins with a CDF H_0 on $\Re^q$ having continuous marginals H_i. For a random vector $\boldsymbol{J} = (J^{[1]}, J^{[2]}, \ldots, J^{[q]}) \sim H_0$, let

$$\boldsymbol{W} = (W^{[1]}, W^{[2]}, \ldots, W^{[q]}) \text{ with each } W^{[i]} = H_i(J^{[i]}) \sim \mathcal{U}[0, 1]. \tag{66}$$

In general, $W^{[1]}, W^{[2]}, \ldots, W^{[q]}$ are dependent, and their joint CDF is the copula K in (64) *induced* by H_0. Our experiments use two different copulas. Models 1 and 3 employ a Gaussian copula [55, pp. 190–191], in which H_0 is a multivariate normal CDF. Models 2 and 4 use a Student-t copula [56], where H_0 is a multivariate Student-t CDF.

To define a Gaussian copula, let H_0 be the $\mathcal{N}_q(\boldsymbol{0}, \boldsymbol{\Sigma})$ normal CDF with mean vector $\boldsymbol{0}$ and covariance matrix $\boldsymbol{\Sigma} = (\Sigma^{[i,j]}: i, j = 1, 2, \ldots, q)$, where each $\Sigma^{[i,i]} = 1$. Thus, each marginal is $H_i = \Phi$, the standard (i.e., mean 0 and variance 1) univariate normal CDF, and $\boldsymbol{\Sigma}$ is also the (Pearson linear) correlation matrix. Now let $\boldsymbol{Z} = (Z^{[1]}, Z^{[2]}, \ldots, Z^{[q]}) \sim \mathcal{N}_q(\boldsymbol{0}, \boldsymbol{\Sigma})$.

We can generate the (column) vector $\boldsymbol{Z}$ by multiplying a $q \times q$ matrix $\boldsymbol{\Gamma}$ satisfying $\boldsymbol{\Sigma} = \boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\mathrm{T}}$ (e.g., $\boldsymbol{\Gamma}$ is a Cholesky factor of $\boldsymbol{\Sigma}$) with a q-vector $\boldsymbol{D}$ of i.i.d. standard normals:

$$\boldsymbol{Z} = \boldsymbol{\Gamma}\boldsymbol{D}. \tag{67}$$

Then let $\boldsymbol{W}$ be as in (66) with $\boldsymbol{J} = \boldsymbol{Z}$ and each $H_i = \Phi$. The Gaussian copula K with input (Pearson linear) correlation matrix $\boldsymbol{\Sigma}$ is finally given by (64), and $\boldsymbol{W} \sim K$.

For a Student-t copula, let χ_ν^2 be a univariate chi-squared random variable with $\nu \geq 1$ degrees of freedom, independent of the $\boldsymbol{Z} \sim \mathcal{N}_q(\boldsymbol{0}, \boldsymbol{\Sigma})$ in (67). We then obtain a multivariate Student-t random vector with ν d.f. as

$$\boldsymbol{J} = \boldsymbol{Z}\sqrt{\nu/\chi_\nu^2}\,. \tag{68}$$

Let $\boldsymbol{W}$ be as in (66) with each $H_i = \Upsilon_\nu$, the univariate Student-t CDF with ν d.f., and then (64) defines the Student-t copula K with ν d.f. and input dispersion matrix $\boldsymbol{\Sigma}$, which is the (Pearson linear) correlation matrix of the normal $\boldsymbol{Z}$ used to build $\boldsymbol{J}$ in (68). If $\nu > 2$, then the covariance matrix of $\boldsymbol{J}$ is $\mathrm{Cov}[\boldsymbol{J}] = (\nu/(\nu - 2))\boldsymbol{\Sigma}$, and the Student-$t$ vector $\boldsymbol{J}$ in (68) also has (Pearson linear) correlation matrix $\boldsymbol{\Sigma}$. But when $\nu \leq 2$, each $J^{[i]}$ has infinite variance, so the Pearson correlation of $\boldsymbol{J}$ is undefined.

A key difference between the Gaussian and Student-t copulas lies in their tail dependence, as we now explain for dimension q=2. Suppose that $\boldsymbol{Y} = (Y^{[1]}, Y^{[2]})$ is a bivariate output random vector with copula K and marginal CDFs F_1 and F_2. Now define $\upsilon = \lim_{p \to 1} P(Y^{[2]} > F_2^{-1}(p) | Y^{[1]} > F_1^{-1}(p))$ as the limiting conditional probability that the second entry of $\boldsymbol{Y}$ exceeds its p-quantile given that the first is larger than its p-quantile. If the marginals F_1 and F_2 are continuous, then (65) implies that the value of υ depends only on the

copula K and not on F_1 and F_2. A Gaussian copula has $\upsilon = 0$, whereas $\upsilon > 0$ for a Student-t copula. Hence, a Gaussian copula has asymptotic independence in the tails, but a Student-t copula tends to produce joint extreme events. For more details, see Section 3.2 of [56].

As Section 7.1 already specified marginal CDFs for the loads and capacities, our next task is to pick a copula K so that the generated (load or capacity) vector $\boldsymbol{Y} = (Y^{[1]}, Y^{[2]}\ldots, Y^{[q]})$ has a desired dependence structure. For example, we may want a particular tail dependence υ for $\boldsymbol{Y}$. We may also want $\boldsymbol{Y}$ to have a target (Pearson linear) correlation matrix $\boldsymbol{\rho} = (\rho^{[i,j]}: i, j = 1, 2, \ldots, q)$, where $\rho^{[i,j]} = \mathrm{Cov}[Y^{[i]}, Y^{[j]}]/(\mathrm{Var}[Y^{[i]}]\mathrm{Var}[Y^{[j]}])^{1/2}$, which depends on both K and the marginals. If we have selected a copula family (e.g., Gaussian), it still remains to specify its parameters (e.g., the input correlation matrix $\boldsymbol{\Sigma}$ of a Gaussian copula) to achieve the target $\boldsymbol{\rho}$, assuming it is indeed attainable (Theorem 5.25 of [55]). There usually is no closed-form formula for $\boldsymbol{\rho}$ in terms of the copula parameters, and we can instead use a sampling-based search procedure, e.g., [57], to identify the copula parameters to obtain the target $\boldsymbol{\rho}$.

Another possible dependence measure to specify for $\boldsymbol{Y}$ is *rank correlation*, which has different variants [55, Section 5.2.2]. The (population) Spearman rank correlation of $Y^{[i]} \sim F_i$ and $Y^{[j]} \sim F_j$ is defined as $\rho_S^{[i,j]} = \mathrm{Cov}[F_i(Y^{[i]}), F_j(Y^{[j]})]/(\mathrm{Var}[F_i(Y^{[i]})]\mathrm{Var}[F_j(Y^{[j]})])^{1/2}$. The Kendall (tau) rank correlation of $Y^{[i]}$ and $Y^{[j]}$ is $\rho_\tau^{[i,j]} = E[\mathrm{sgn}((Y^{[i]} - \widetilde{Y}^{[i]})(Y^{[j]} - \widetilde{Y}^{[j]}))]$, where $(\widetilde{Y}^{[i]}, \widetilde{Y}^{[j]})$ is an independent copy of $(Y^{[i]}, Y^{[j]})$, and $\mathrm{sgn}(x) = 1$ (resp., -1 and 0) if $x > 0$ (resp., $x < 0$ and x=0). When the marginal CDFs F_i are continuous (which we will assume for the rest of the paragraph), both rank correlations of $\boldsymbol{Y}$ depend only on the copula and not on the F_i [55, Proposition 5.29], in contrast to the (Pearson linear) correlation $\rho^{[i,j]}$, which depends on both. For a Gaussian copula, the Spearman and Kendall rank correlations relate to the input Pearson correlation $\Sigma^{[i,j]}$ of the Gaussian copula by $\rho_S^{[i,j]} = (6/\pi)\arcsin(\Sigma^{[i,j]}/2)$ and $\rho_\tau^{[i,j]} = (2/\pi)\arcsin(\Sigma^{[i,j]})$; see [55, Theorem 5.36]. Thus, if we want to generate $\boldsymbol{Y}$ with a Gaussian copula to have a specified Spearman or Kendall rank correlation, we choose the input correlation matrix $\boldsymbol{\Sigma}$ of the normal $\boldsymbol{Z}$ in (67) to have entries $\Sigma^{[i,j]} = 2\sin(\pi\rho_S^{[i,j]}/6)$ or $\Sigma^{[i,j]} = \sin(\pi\rho_\tau^{[i,j]}/2)$, respectively; then use (66) with $\boldsymbol{J} = \boldsymbol{Z}$ and $H_i = \Phi$, and build $\boldsymbol{Y}$ via (65). For a Gaussian copula, the difference between the Spearman rank correlation and the input Pearson correlation is quite small [55, p. 216]: $|\rho_S^{[i,j]} - \Sigma^{[i,j]}| \le 0.0181$. This fact is implicitly exploited by [58], which essentially employs a Gaussian copula to generate a sample with a specified sample Spearman rank correlation. As [55, Example 5.54] notes, if we instead want to generate vector $\boldsymbol{Y}$ with a Student-t copula, we can similarly calibrate the copula's input dispersion matrix $\boldsymbol{\Sigma}$ so that $\boldsymbol{Y}$ has a target Kendall rank correlation, as follows. For a Student-t copula, the Kendall rank correlation is again $\rho_\tau^{[i,j]} = (2/\pi)\arcsin(\Sigma^{[i,j]})$. Hence, to achieve a target Kendall value $\rho_\tau^{[i,j]}$, assign the copula's dispersion matrix $\boldsymbol{\Sigma}$ to have entries $\Sigma^{[i,j]} = \sin(\pi\rho_\tau^{[i,j]}/2)$; apply (66) with $H_i = Y_\nu$, and use (65) to obtain $\boldsymbol{Y}$. The Spearman rank correlation for a Student-t copula is not known in closed form.

In our experiments, we chose the input (correlation or dispersion) matrix $\boldsymbol{\Sigma}$ of the Gaussian or Student-t copula so that the generated output random vector $\boldsymbol{Y}$ in (65) has an appropriate (Pearson) correlation matrix $\boldsymbol{\rho}$. When $\boldsymbol{Y}$ is the load vector, which has size q=3, we calibrated $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_L = (\Sigma_L^{[i,j]}: i, j = 1, 2, 3)$ so as to match an estimated (Pearson) correlation matrix $\boldsymbol{\rho}_L = (\rho_L^{[i,j]}: i, j = 1, 2, 3)$ with $\rho_L^{[1,2]} = 0.85$, $\rho_L^{[1,3]} = 0.87$, and $\rho_L^{[2,3]} = 0.83$, which [10] computed from a sample of the (PCT, CWO, MLO) load vectors produced from multiple runs of a nuclear-specific computer code. Applying the search algorithm of [57] with the MLEs of the parameters of the specified marginal load CDFs (see Section 7.1), we found $\Sigma_L^{[1,2]} = 0.92$, $\Sigma_L^{[1,3]} = 0.96$, and $\Sigma_L^{[2,3]} = 0.86$ for the Gaussian copula to yield the target load correlations $\boldsymbol{\rho}_L$ (to two decimal places of accuracy). We used these input correlations for the loads across every scenario for Models 1 and 3. The Student-t copula in Models 2 and 4 have $\nu = 3$ d.f. in (68), and we identified the input dispersions as $\Sigma_L^{[1,2]} = 0.92$, $\Sigma_L^{[1,3]} = 0.96$, and $\Sigma_L^{[2,3]} = 0.85$, which are used across every scenario for the loads. For the input matrix $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_C$ of the capacity (Gaussian and Student-t) copulas, we let $\Sigma_C^{[1,2]} = \Sigma_C^{[1,3]} = \Sigma_C^{[2,3]} = 0.85$ for all models and scenarios.

7.3. Generating loads and capacities

We now combine the marginals chosen in Section 7.1 with copulas from Section 7.2 to specify the functions $\mathfrak{w}'_{(s)}$, $\mathfrak{l}_{(s)}$, and $\mathfrak{c}_{(s)}$ generating the load and capacity vectors for scenario s in Assumption A4 and (58). In general, to evaluate a d-dimensional integral, Monte Carlo methods become attractive alternatives to numerical quadrature when d is high, so we artificially increase the number $d_s = d_{s,L} + d_{s,C}$ of input variables of function $\mathfrak{w}'_{(s)}$ in Assumption A4. We first explain how to generate an observation of the load vector for scenario s using a Gaussian copula with the input correlation matrix $\boldsymbol{\Sigma}_L$ specified in Section 7.2. We define the load function $\mathfrak{l}_{(s)}$ in (58) to have $d_{s,L} = d_{s,L}^{[1]} + d_{s,L}^{[2]} + d_{s,L}^{[3]}$ i.i.d. $\mathcal{U}[0, 1]$ inputs $U_1, U_2, \ldots, U_{d_{s,L}}$, with each $d_{s,L}^{[k]} = 10$, and let $\boldsymbol{D}_L = (D_L^{[1]}, D_L^{[2]}, D_L^{[3]})$ with $D_L^{[1]} = \sum_{j=1}^{d_{s,L}^{[1]}} \Phi^{-1}(U_j)/\sqrt{d_{s,L}^{[1]}}$, $D_L^{[2]} = \sum_{j=d_{s,L}^{[1]}+1}^{d_{s,L}^{[1]}+d_{s,L}^{[2]}} \Phi^{-1}(U_j)/\sqrt{d_{s,L}^{[2]}}$, and $D_L^{[3]} = \sum_{j=d_{s,L}^{[1]}+d_{s,L}^{[2]}+1}^{d_{s,L}^{[1]}+d_{s,L}^{[2]}+d_{s,L}^{[3]}} \Phi^{-1}(U_j)/\sqrt{d_{s,L}^{[3]}}$. Thus, the entries of $\boldsymbol{D}_L$ are i.i.d. standard normal. Then, as in (67), define the vector $\boldsymbol{Z}_L = (Z_L^{[1]}, Z_L^{[2]}, Z_L^{[3]})$ as $\boldsymbol{Z}_L = \boldsymbol{\Gamma}_L\boldsymbol{D}_L$, where $\boldsymbol{\Gamma}_L$ is a Cholesky factor of $\boldsymbol{\Sigma}_L$. Letting $\boldsymbol{J} = \boldsymbol{Z}_L$, we obtain $\boldsymbol{W} = \boldsymbol{W}_L = (W_L^{[1]}, W_L^{[2]}, W_L^{[3]})$ as in (66), with each $H_i = \Phi$. By (65), the load vector

$$\boldsymbol{L}_{(s)} = (L_{(s)}^{[1]}, L_{(s)}^{[2]}, L_{(s)}^{[3]}) = ((F_{(s)}^{[1]})^{-1}(W_L^{[1]}), (F_{(s)}^{[2]})^{-1}(W_L^{[2]}), (F_{(s)}^{[3]})^{-1}(W_L^{[3]})) \sim F_{(s)} \tag{69}$$

has a Gaussian copula with input correlation matrix $\boldsymbol{\Sigma}_L$, and each $\boldsymbol{L}_{(s)}^{[k]}$ has marginal CDF $F_{(s)}^{[k]}$ given in Section 7.1.

We apply similar ideas to define the function $\mathfrak{c}_{(s)}$ in (58) to get an observation of the capacity vector for scenario s when employing the Gaussian copula with input correlation matrix $\boldsymbol{\Sigma}_C$ specified in Section 7.2. We use $d_{s,C} = d_{s,C}^{[1]} + d_{s,C}^{[2]} + d_{s,C}^{[3]}$ i.i.d. $\mathcal{U}[0, 1]$ inputs $U_{d_{s,L}+1}, \ldots, U_{d_{s,L}+d_{s,C}}$, with each $d_{s,C}^{[k]} = 1$, to obtain normal vector $\boldsymbol{Z}_C = (Z_C^{[1]}, Z_C^{[2]}, Z_C^{[3]})$ as $\boldsymbol{Z}_C = \boldsymbol{\Gamma}_C\boldsymbol{D}_C$, where $\boldsymbol{\Gamma}_C\boldsymbol{\Gamma}_C^\top = \boldsymbol{\Sigma}_C$ and $\boldsymbol{D}_C = (\Phi^{-1}(U_{d_{s,L}+1}), \Phi^{-1}(U_{d_{s,L}+2}), \Phi^{-1}(U_{d_{s,L}+3}))$. Letting $\boldsymbol{J} = \boldsymbol{Z}_C$, we obtain $\boldsymbol{W} = \boldsymbol{W}_C = (W_C^{[1]}, W_C^{[2]}, W_C^{[3]})$ as in (66), with each $H_i = \Phi$. By (65), the capacity vector

$$\boldsymbol{C}_{(s)} = (C^{[1]}, C^{[2]}, C^{[3]}) = ((G^{[1]})^{-1}(W_C^{[1]}), (G^{[2]})^{-1}(W_C^{[2]}), (G^{[3]})^{-1}(W_C^{[3]})) \sim G \tag{70}$$

has Gaussian copula with input correlation matrix $\boldsymbol{\Sigma}_C$, and each $C^{[k]} \sim G^{[k]}$ from Section 7.1. As $(U_1, \ldots, U_{d_{s,L}})$ and $(U_{d_{s,L}+1}, \ldots, U_{d_{s,L}+d_{s,C}})$ are independent as in Assumption A4, $\boldsymbol{L}_{(s)}$ and $\boldsymbol{C}_{(s)}$ are also (Proposition 1). Note that $\mathfrak{w}'_{(s)}$ in A4 uses $d_s = d_{s,L} + d_{s,C} = 33$ i.i.d. uniforms to generate $\boldsymbol{X}_{(s)} = (\boldsymbol{L}_{(s)}, \boldsymbol{C}_{(s)})$ when employing Gaussian copulas.

When instead assuming Student-t copulas, we make the following changes to the load and capacity functions $\mathfrak{l}_{(s)}$ and $\mathfrak{c}_{(s)}$. We now let $d_{s,L} = d_{s,L}^{[1]} + d_{s,L}^{[2]} + d_{s,L}^{[3]} + 1$ and $d_{s,C} = d_{s,C}^{[1]} + d_{s,C}^{[2]} + d_{s,C}^{[3]} + 1$, with each $d_{s,L}^{[k]} = 10$ and $d_{s,C}^{[k]} = 1$ as before. We set $\chi_{\nu,L}^2 = \Psi_\nu^{-1}(U_{d_{s,L}})$ and $\chi_{\nu,C}^2 = \Psi_\nu^{-1}(U_{d_{s,L}+d_{s,C}})$, where Ψ_ν is the CDF of a univariate chi-squared random variable with $\nu = 3$ d.f. Then let $\boldsymbol{J} = \boldsymbol{Z}_L\sqrt{\nu/\chi_{\nu,L}^2}$ with $\boldsymbol{Z}_L$ as before, and apply (66) with each $H_i = Y_\nu$ to get $\boldsymbol{W} = \boldsymbol{W}_L$. Finally, (69) yields a load vector having a Student-t copula with $\nu = 3$ d.f. and the input dispersion matrix $\boldsymbol{\Sigma}_L$ specified in Section 7.2. For the capacities, we instead use $\boldsymbol{J} = \boldsymbol{Z}_C\sqrt{\nu/\chi_{\nu,C}^2}$ (with $\boldsymbol{Z}_C$ as before) in (66) to get $\boldsymbol{W} = \boldsymbol{W}_C$ and apply (70). Thus, $\mathfrak{w}'_{(s)}$ in A4 takes $d_s = d_{s,L} + d_{s,C} = 35$ uniforms to produce $\boldsymbol{L}_{(s)}$ and $\boldsymbol{C}_{(s)}$ with Student-t cields a load vector having a Studentopulas.

Our experiments use the series-system failure function $\mathfrak{d}'$ from (7) in Assumption A5, so (59) has $a^{[0]} = 0$ and each $a^{[k_1,k_2,\ldots,k_p]} = (-1)^{p-1}$. To ensure the tractability of the response function (63) when CMC is

applied, Assumption A6 requires that we can compute the marginal CDF $G^{[k_1,k_2,\ldots,k_p]}$ in (61) for each subcollection $C^{[k1]}$, $C^{[k2]}$, …, $C^{[k_p]}$ of $1 \le p \le 3$ criteria's capacities, which we next show holds for both copulas. When p=1, we have a univariate CDF $G^{[k]}$, which is trivial to evaluate as each $G^{[k]}$ is triangular. For p=2 with the PCT and CWO capacities $C^{[1]}$ and $C^{[2]}$, note that (70) implies that their marginal joint CDF satisfies

$$\begin{aligned} G^{[1,2]}(x_1, x_2) &= P(C^{[1]} \le x_1,\ C^{[2]} \le x_2) = P((G^{[1]})^{-1}(W_C^{[1]}) \le x_1,\ (G^{[2]})^{-1}(W_C^{[2]}) \le x_2) \\ &= P(W_C^{[1]} \le G^{[1]}(x_1),\ W_C^{[2]} \le G^{[2]}(x_2)) \\ &= P(H_1(J^{[1]}) \le G^{[1]}(x_1),\ H_2(J^{[2]}) \le G^{[2]}(x_2)) \\ &= P(J^{[1]} \le H_1^{-1}(G^{[1]}(x_1)),\ J^{[2]} \le H_2^{-1}(G^{[2]}(x_2))) \end{aligned} \tag{71}$$

by (66). When the capacities are based on a Gaussian (resp., Student-t) copula, the vector $\boldsymbol{J}$ used in (66) and (71) is multivariate normal (resp., multivariate Student-t) and each H_i is a standard univariate normal (resp., univariate Student-t) CDF, so (71) is a multivariate normal (resp., multivariate Student-t) CDF evaluation, which can be handled numerically. Similarly, we can numerically compute the other marginal joint CDFs $G^{[1,3]}$, $G^{[2,3]}$, and $G^{[1,2,3]}$ in Assumption A6. These computations are somewhat involved, but in an actual RISMC analysis, the computer code $\mathfrak{l}_{(s)}$ in (58) generating loads typically takes substantially more time, making the overhead of CMC negligible.

7.4. Discussion of Monte Carlo results

We ran Monte Carlo experiments with our four models, whose differences are summarized in Table 1. Because of the analytical tractability of our models (which would not be the case for a real NPP analysis), we are able to numerically evaluate various multi-dimensional integrals to determine the true failure probability θ in (3) and (16). The last column of Table 1 shows that for our models, the value of θ does not vary much for the different copulas, but we will see below that the performances of the Monte Carlo methods are affected, sometimes substantially, by the change in the dependence structure. (Section 7.5 will give results for another model demonstrating significant differences in θ as we change the copula.)

Tables 2–4 present results from our Monte Carlo experiments with Models 1, 3, and 4, respectively. We omit the results for Model 2 as they are very similar to those for Model 1. Our experiments construct upper confidence bounds $B(n)$ satisfying (5) using four different combinations of simulation methods: SS (see Section 4), SS+CMC (developed in [25]), SS+rLHS (Section 5.2), and SS+CMC+rLHS (Section 6.2). Table 3 includes Model 3's results for only SS+rLHS and SS+CMC+rLHS, which also appear in [40], as the other methods' results do not differ significantly from those for Model 4. We varied the total sample size $n = 4^{\nu} \times 100$ for $\nu = 1, 2, 3, 4$. For each n and each Monte Carlo method, we ran 10^4 independent experiments. The stratification allocation has $\eta_s = 0.25$ for each scenario $s = 1, 2, 3, 4$. For rLHS, we used b=10 independent replicates of LHS samples within a scenario, so for each scenario s, each replicate has size $m_s = \eta_s n/b$. Methods that do not require replications to create an upper confidence bound are denoted with b=1. For the safety requirement (4), we set $\theta_0 = 0.05$ and $\gamma = 0.95$. As seen in Table 1, all four models have $\theta < \theta_0$, although our experiments do not use this knowledge.

In Tables 2–4, the column labeled "95% AHW" presents the average half-width across the 10^4 experiments, where the half-width is the difference between the 95% UCB and the point estimate of θ. For a UCB $B(n)$ and overall sample size n, the *coverage* is the probability $P(B(n) \ge \theta)$. As noted throughout the paper, each method's UCB satisfies (5), so the coverage converges to the nominal level $\gamma = 0.95$ as $n \to \infty$. But for a fixed n, the coverage may differ from γ. We estimate the coverage as the fraction of the 10^4 experiments in which $B(n) \ge \theta$, and one gauge of the convergence of a method is how far its coverage is from γ. The *probability of correct decision* (PCD) is estimated as the fraction of the 10^4 experiments that the decision rule in (6) correctly determined that $\theta < \theta_0$. The column "CV" gives the coefficient of variation of the point estimate of θ, computed as the

Table 1
Summary of model differences.

Model	PCT load CDF parameters	Copula	θ
1	Based on data from [10]	Gaussian	0.0254630
2	Based on data from [10]	Student-t with ν = 3 d.f.	0.0254639
3	Taken from [24]	Gaussian	0.0464784
4	Taken from [24]	Student-t with ν = 3 d.f.	0.0464774

Table 2
Results for Model 1 (θ = 0.0254630).

Method	n	b	95% AHW	Coverage	PCD	CV	VRF SS/x
SS	400	1	2.16E−02	0.900	0.578	5.51E−01	1.00
	1600	1	1.14E−02	0.908	0.929	2.74E−01	1.00
	6400	1	5.68E−03	0.928	1.000	1.38E−01	1.00
	25,600	1	2.82E−03	0.936	1.000	6.90E−02	1.00
SS+CMC	400	1	1.40E−02	0.846	0.785	3.51E−01	2.45
	1600	1	7.31E−03	0.904	0.998	1.78E−01	2.37
	6400	1	3.64E−03	0.931	1.000	8.78E−02	2.45
	25,600	1	1.82E−03	0.937	1.000	4.44E−02	2.41
SS+rLHS	400	10	2.26E−02	0.923	0.567	5.13E−01	1.15
	1600	10	1.14E−02	0.923	0.940	2.52E−01	1.18
	6400	10	5.69E−03	0.936	1.000	1.26E−01	1.20
	25,600	10	2.84E−03	0.944	1.000	6.27E−02	1.21
SS+CMC +rLHS	400	10	1.36E−02	0.871	0.820	3.09E−01	3.17
	1600	10	6.77E−03	0.928	0.999	1.48E−01	3.42
	6400	10	3.32E−03	0.936	1.000	7.38E−02	3.47
	25,600	10	1.64E−03	0.941	1.000	3.67E−02	3.55

Table 3
LHS results for Model 3 (θ = 0.0464784).

Method	n	b	95% AHW	Coverage	PCD	CV	VRF SS/x
SS+rLHS	400	10	7.40E−03	0.777	0.616	1.15E−01	1.39
	1600	10	4.43E−03	0.793	0.395	5.59E−02	1.48
	6400	10	2.20E−03	0.921	0.801	2.67E−02	1.67
	25,600	10	1.08E−03	0.936	1.000	1.30E−02	1.73
SS+CMC +rLHS	400	10	8.59E−04	0.933	1.000	1.06E−02	161.39
	1600	10	3.45E−04	0.942	1.000	4.15E−03	268.13
	6400	10	1.62E−04	0.940	1.000	2.00E−03	299.20
	25,600	10	7.91E−05	0.942	1.000	9.79E−04	307.71

Table 4
Results for Model 4 (θ = 0.0464744).

Method	n	b	95% AHW	Coverage	PCD	CV	VRF SS/x
SS	400	1	9.25E−03	0.864	0.345	1.38E−01	1.00
	1600	1	5.08E−03	0.879	0.394	6.89E−02	1.00
	6400	1	2.59E−03	0.919	0.695	3.45E−02	1.00
	25,600	1	1.31E−03	0.934	0.992	1.73E−02	1.00
SS+CMC	400	1	2.56E−03	0.946	0.725	3.39E−02	16.63
	1600	1	1.30E−03	0.954	0.997	1.66E−02	17.11
	6400	1	6.37E−04	0.947	1.000	8.44E−03	16.70
	25,600	1	3.21E−04	0.951	1.000	4.25E−03	16.68
SS+rLHS	400	10	7.82E−03	0.798	0.581	1.16E−01	1.41
	1600	10	4.55E−03	0.793	0.404	5.79E−02	1.42
	6400	10	2.34E−03	0.911	0.756	2.87E−02	1.45
	25,600	10	1.19E−03	0.934	0.996	1.45E−02	1.44
SS+CMC +rLHS	400	10	1.25E−03	0.947	0.998	1.48E−02	87.41
	1600	10	5.63E−04	0.947	1.000	6.79E−03	102.71
	6400	10	2.78E−04	0.949	1.000	3.34E−03	106.44
	25,600	10	1.36E−04	0.948	1.000	1.67E−03	108.48

sample standard deviation of the point estimate of θ across the 10^4 experiments divided by the numerically-computed value of θ in Table 1. For each particular method x, the last column presents the *variance-reduction factor* (VRF), which for a given overall sample size n is the ratio of the sample variances for the base case of SS and for method x. The VRF gives the factor by which the overall sample size for method x can be reduced compared to SS to achieve about the same AHW. For example, in Table 4, SS+CMC has VRF ≈ 16, and its AHW for n=400 (resp., 1600) is about the same as the SS AHW for n=6400 (resp., 25,600).

Tables 2–4 clearly show that each method's performance depends on the model. For example, for small n in Model 1 (and 2, not shown), the coverages for the methods with CMC are lower than those for the corresponding approaches without CMC, but the opposite is true for Models 3 and 4. Also, each method exhibits variations across models in the sample size n at which convergence appears to have been roughly achieved (as defined by coverage being within 5% of the nominal level $\gamma = 0.95$). For example, SS+rLHS appears to approximately converge by n=400 for Model 1; in contrast, Model 3 requires n=6400. However, by n=6400, the methods SS and SS+rLHS seem to have nearly converged for all models, whereas SS+CMC and SS+CMC+rLHS appear to need n=1600.

But even with the performance differences across models for each method, Tables 2–4 also illustrate some universal trends. For all models, each combined VRT has VRF > 1, so we see the effectiveness of combining methods. This agrees with the asymptotic theory established in Theorems 1 and 3 for SS+LHS and SS+CMC+LHS, respectively. In particular, the combination of SS+CMC+rLHS provides the most variance reduction for each model, with VRFs of about 300 and 100 for Models 3 (Table 3) and 4 (Table 4), respectively. The VRFs for SS+CMC+rLHS are more modest (about 3.5) for Models 1 (Table 2) and Table 2 (not shown). Thus, the effectiveness of the method depends on the marginal distributions (compare Models 1 and 3) and the copula (Models 3 and 4).

For all four models, the VRF for SS+CMC+rLHS exceeds the product of the VRFs for SS+CMC and SS+rLHS, illustrating the synergistic effect of integrating LHS with CMC. As explained in Section 5.1, LHS can be particularly effective when the response is a nearly additive function of the input random variables. Without CMC, the response function (35) is an indicator, for which an additive approximation can be a poor fit, so LHS by itself may not reduce variance by much. But incorporating CMC leads to a smoother response function (47) through the integration performed for the conditional expectation, so its additive fit (54) can be much better; as Section 6.1 notes, LHS can then produce substantially more variance reduction, as Tables 2–4 show.

Another benefit of reduced variance becomes evident when examining the PCD. Many application domains critically rely on a high probability of correctly deciding if the safety requirement $\theta < \theta_0$ in (4) holds. But it can be difficult to achieve high PCD when θ is close to θ_0, as in Models 3 and 4 where $\theta \approx 0.0465$ and $\theta_0 = 0.05$. In general it is easier for the decision rule (6) to make a correct determination when the point estimator of θ has small statistical error. Because the error asymptotically decreases at rate $n^{-1/2}$ by virtue of the relevant CLT, the PCD converges to 1 as $n \to \infty$, as seen throughout Tables 2–4 for each method, so we can ensure high PCD by having a large sample size. But this may not always be feasible, e.g., when a simulation run is extremely time-consuming. Instead, we can also achieve higher PCD for a fixed n by applying VRTs that reduce the error. For example, Table 4 shows that for n=400, method SS has PCD = 0.345, whereas SS +CMC+rLHS has PCD = 0.998. Thus, the much smaller variance for SS +CMC+rLHS results in a substantially higher chance of making a correct decision for the same sample size.

7.5. Additional numerical results

While Table 1 displays little difference in the exact values of θ for the Gaussian and Student-t copulas, this is not always the case, as we will see later. But first, we explain why the two copulas lead to essentially the same values for θ for Models 1 and 2 and for Models 3 and 4. For each criterion k = 1, 2, 3, let $\theta^{[k]} = \sum_{s=1}^{s_0} \lambda_s P(L_{(s)}^{[k]} \geq C_{(s)}^{[k]})$ be the marginal failure probability for criterion k by itself. Models 1–4 all have that $\theta^{[2]}$ = 2.13E−15 and $\theta^{[3]}$ = 2.73E−6. For criterion k=1 (PCT), Models 1 and 2 (resp., 3 and 4) have $\theta^{[1]} = 0.0255$ (resp., $\theta^{[1]} = 0.0465$), which overwhelm $\theta^{[2]}$ and $\theta^{[3]}$. Thus, the overall failure probability θ is mainly determined by PCT, so the particular dependence structure among the criteria has little impact.

To demonstrate that differing copulas can yield substantial changes in θ, we altered the parameters of the load distributions. For each scenario s, the lognormal CDF parameters for criterion k=1 are instead $\mu_{(s)}^{[1]} = 7.33 + 0.1\,s$ and $\sigma_{(s)}^{[1]} = 0.01 + 0.01\,s$; the Weibull parameters for criterion k=2 (CWO) are $\alpha_{(s)}^{[2]} = 0.22 + 0.005\,s$ and $\beta_{(s)}^{[2]} = 0.90 + 0.02\,s$; and the Weibull parameters for criterion k=3 (MLO) are $\alpha_{(s)}^{[3]} = 3.1 + 0.3\,s$ and $\beta_{(s)}^{[3]} = 0.82 + 0.03\,s$. For the loads, we slightly changed the input correlations to $\Sigma_L^{[1,2]} = 0.93$, $\Sigma_L^{[1,3]} = 0.96$, and $\Sigma_L^{[2,3]} = 0.84$ for the Gaussian copula, and the input dispersions to $\Sigma_L^{[1,2]} = 0.94$, $\Sigma_L^{[1,3]} = 0.97$, and $\Sigma_L^{[2,3]} = 0.83$ for the Student-t copula, in both cases producing load target correlations $\rho_L^{[1,2]} = 0.83$, $\rho_L^{[1,3]} = 0.84$, and $\rho_L^{[2,3]} = 0.81$. The capacities' marginal distributions and copulas are the same as in Models 1–4. With these parameters we now have that $\theta^{[1]} = 0.0201$, $\theta^{[2]} = 0.0201$, and $\theta^{[3]} = 0.0208$, so no single criterion's marginal failure probability dominates the others, in contrast to what occurs for Models 1–4.

Table 5 presents the values of θ, computed numerically, for the Gaussian and Student-t copulas. Now we see a distinct difference, with θ decreasing by about 5% when changing the copula from Gaussian to Student-t *for system failure defined as for a series system*.

We can observe an even greater difference in θ between the copulas by changing the definition of a failure from that for a series system in (7) to instead be for a parallel system, as in (8). As seen in the bottom of Table 5, now θ increases by about 19% when changing the copula from Gaussian to Student-t. To explain the larger change when failure is defined by (8) instead of (7), note that the marginal failure probabilities $\theta^{[k]}$ are fairly small (≈0.02), so the tail dependence plays an important role in determining θ. As we discussed in the paragraph after (68) in Section 7.2, the two copulas give rise to very different dependence behaviors in the tails. Because the Student-t copula asymptotically has non-zero tail dependence, if one criterion is violated, then there is a reasonable chance that the others are as well, leading to higher θ when failure is defined for a parallel system. In contrast, the asymptotic tail independence of the Gaussian copula does not produce such joint behavior, so θ is notably smaller.

8. Concluding remarks

We combined variance-reduction techniques to estimate a failure

Table 5
Values of θ for another set of distribution parameters for the loads with different copulas and different definitions for failure.

Copula	Failure definition	θ
Gaussian	Series system, Eq. (7)	0.0493
Student-t	Series system, Eq. (7)	0.0470
Gaussian	Parallel system, Eq. (8)	6.83E−4
Student-t	Parallel system, Eq. (8)	8.11E−4

probability. We developed the methods in a general setting, where a failure is defined in terms of a basic random object, which may be a random vector or a more general random object, e.g., a stochastic process. We also specialized our approaches to the setting in which a failure is specified as a function of loads and capacities of q criteria (e.g., as in the RISMC problem for a nuclear PSA), and we exploited the added structure to obtain readily implementable techniques. The combination SS+CMC+LHS can be especially effective, substantially reducing variance compared to SS, SS+CMC, and SS+LHS, as shown in the numerical experiments in Section 7. To account for the statistical error of our Monte Carlo estimators, we further devised upper confidence bounds for θ, and we proved the asymptotic validity (as the total sample size $n \to \infty$) of the UCBs when applying SS+rLHS and SS+CMC+rLHS in Theorems 2 and 4, respectively.

Theorem 1 (resp., 3) further established that SS+LHS (resp., SS +CMC+LHS) has smaller asymptotic variance than SS (resp., SS+LHS). We also provided analysis (Sections 5.1 and 6.1) giving insight into why combining CMC with LHS can be so powerful (e.g., Table 3 shows SS +CMC+rLHS reduces variance by about a factor of 300 compared to SS). It is known (e.g., Section 10.3 of [22]) that LHS can tremendously reduce variance when the response function whose mean is being estimated is well approximated by an additive function of the input random numbers. But when estimating a probability without applying CMC, the response is an indicator function, and the additive approximation to it can be a poor fit. Incorporating CMC smooths the response function through a conditional expectation (so the response is no longer necessarily binary-valued), which can lead to a better additive fit and consequently lower variance when combining CMC with LHS. In our development of SS+CMC+LHS in Section 6.3, where failure is defined in terms of loads and capacities of q criteria, Assumption A4 implied that the load vector is independent of the capacity vector (Proposition 1), which holds in many applications. (If it does not for a particular problem, SS+CMC+LHS can still be applied but it then requires working instead with the response function (47) which may be less tractable than (63).) Our methods can be applied to high-dimensional problems; e.g., the models in Section 7 used d_s=33 and 35 input uniforms for each scenario s. We should note, however, that the models in Section 7 are synthetic examples, although they are loosely based on data from actual nuclear computer codes [10,3,19]. While SS+CMC+LHS is guaranteed to outperform SS+LHS in terms of asymptotic variance, it would be interesting to see how much our techniques reduce variance when using actual nuclear computer codes. As noted in Section 1, LHS is most appropriate for estimating measures of central tendency, e.g., a mean, so adding LHS to SS+CMC may not be effective in sufficiently reducing variance when estimating rare-event probabilities, smaller than say 10^{-4}. In such cases, other Monte Carlo methods, such as importance sampling or splitting [14, Chapter VI], may be needed.

Developing appropriate stochastic models for multivariate capacities also deserves further investigation. The marginal distributions and pairwise (Pearson linear) correlations do not fully specify the entire dependence structure, and there can be infinitely many joint distributions fitting that limited information. On the other hand, replacing correlations with a copula does completely determine the joint distribution for given continuous marginals. In our numerical experiments, we assumed that each criterion's (PCT, CWO, and MLO) capacity has a marginal triangular distribution, with the joint dependence structure defined by a Gaussian or Student-t copula. We chose marginal triangular capacity distributions because [3,19] use this when considering only a single criterion, PCT. As our numerical results indicate, the value of the failure probability θ and the performance of our developed methodologies depend on the particular stochastic structure specified, so it is critical to build accurate models.

Acknowledgements

This work has been supported in part by the National Science Foundation under Grants No. CMMI-1200065, DMS-1331010, and CMMI-1537322. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Also, additional funding came from a NJIT Provost Undergraduate Summer Research Fellowship.

Appendix A. Proof of Theorem 2

Note that $P(B_{\mathrm{SS+rLHS},\eta,b}(n) \ge \theta) = P(\sqrt{b}\,[\widehat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n) - \theta]/S_b(n) \ge -\tau_{b-1,\gamma})$, so it suffices to show that

$$\frac{\sqrt{b}\left[\widehat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n) - \theta\right]}{S_b(n)} = \frac{\sqrt{bn}\left[\widehat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n) - \theta\right]}{\sqrt{n}\,S_b(n)} \Rightarrow T_{b-1} \tag{A.1}$$

as $n \to \infty$ for fixed $b \ge 2$, where we recall that T_{b-1} is a Student-t random variable with $b-1$ d.f. To accomplish this, we will start by analyzing $\widehat{\theta}_{(s)}^{\langle r\rangle}(n)$ in (39) and build up to $\widehat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n)$ and $S_b(n)$ through (40) and (41).

The estimator $\widehat{\theta}_{(s)}^{\langle r\rangle}(n)$ in (39) of $\theta_{(s)}$ from replication r has bounded response functions $\mathcal{I}_{(s),i}^{\langle r\rangle}$ defined in (38) and (25), so the estimator satisfies the LHS CLT of [49]:

$$\sqrt{n}\left[\widehat{\theta}_{(s)}^{\langle r\rangle}(n) - \theta_{(s)}\right] = \sqrt{\frac{bm_s}{\eta_s}}\left[\frac{1}{m_s}\sum_{i=1}^{m_s}\mathcal{I}_{(s),i}^{\langle r\rangle} - \theta_{(s)}\right] \Rightarrow N_{(s)}^{\langle r\rangle} \overset{\mathcal{D}}{=} \sqrt{\frac{b}{\eta_s}}\,\mathcal{N}(0,\,\sigma_{(s),\mathrm{SS+LHS}}^2) \overset{\mathcal{D}}{=} \mathcal{N}(0,\,\sigma_{(s),\mathrm{SS+LHS}}^2 b/\eta_s) \tag{A.2}$$

as $n \to \infty$ because $m_s = \eta_s n/b \to \infty$, where an expression for $\sigma_{(s),\mathrm{SS+LHS}}^2$ is given in (36). The mutual independence of $\widehat{\theta}_{(s)}^{\langle r\rangle}(n)$, $1 \le r \le b$, $1 \le s \le s_0$, ensures the joint convergence

$$\left(\sqrt{n}\left[\widehat{\theta}_{(s)}^{\langle r\rangle}(n) - \theta_{(s)}\right]\colon r = 1, 2, \ldots, b;\ s = 1, 2, \ldots, s_0\right) \Rightarrow \left(N_{(s)}^{\langle r\rangle}\colon r = 1, 2, \ldots, b;\ s = 1, 2, \ldots, s_0\right) \tag{A.3}$$

as $n \to \infty$ by Problem 29.2 of [43], where $N_{(s)}^{\langle r\rangle}, 1 \le r \le b, 1 \le s \le s_0$, are independent $\mathcal{N}(0,\,\sigma_{(s),\mathrm{SS+LHS}}^2 b/\eta_s)$ random variables. Thus, for each replication $r = 1, 2, \ldots, b$, we have that (16) and (40) imply $\widehat{\theta}^{\langle r\rangle}(n)$ satisfies

$$\sqrt{n}\left[\widehat{\theta}^{\langle r\rangle}(n) - \theta\right] = \sum_{s=1}^{s_0}\lambda_s\sqrt{n}\left[\widehat{\theta}_{(s)}^{\langle r\rangle}(n) - \theta_{(s)}\right] \Rightarrow \sum_{s=1}^{s_0}\lambda_s N_{(s)}^{\langle r\rangle} \equiv N^{\langle r\rangle} \overset{\mathcal{D}}{=} \mathcal{N}(0,\,\psi^2) \tag{A.4}$$

as $n \to \infty$ by (A.3) and the continuous-mapping theorem (e.g., Theorem 29.2 of [43]), where the asymptotic variance $\psi^2 = \sum_{s=1}^{s_0} \lambda_s^2 \sigma^2_{(s),\mathrm{SS+LHS}} b/\eta_s = b\sigma^2_{\mathrm{SS+LHS},\eta}$ by (31) and the independence of $N^{\langle r \rangle}_{(s)}$, $1 \le s \le s_0$. Moreover, $N^{\langle r \rangle}$, $1 \le r \le b$, are i.i.d. Hence, (41) and (A.4) ensure that for the numerator in the middle term of (A.1), we have

$$\sqrt{bn}\,[\hat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n) - \theta] = \frac{1}{\sqrt{b}} \sum_{r=1}^{b} \sqrt{n}\left[\hat{\theta}^{\langle r \rangle}(n) - \theta\right] \Rightarrow \frac{1}{\sqrt{b}} \sum_{r=1}^{b} N^{\langle r \rangle} \overset{\mathcal{D}}{=} \psi \mathcal{N}(0,\,1) \tag{A.5}$$

as $n \to \infty$ for fixed $b \ge 2$ because $\sum_{r=1}^{b} N^{\langle r \rangle} \sim \mathcal{N}(0,\, b\psi^2)$.

Now we consider the square of the denominator of the middle term in (A.1):

$$nS_b^2(n) = \frac{1}{b-1} \sum_{r=1}^{b} \left[\sqrt{n}\left(\hat{\theta}^{\langle r \rangle}(n) - \frac{1}{b}\sum_{r'=1}^{b} \hat{\theta}^{\langle r' \rangle}(n)\right)\right]^2 = \frac{1}{b-1} \sum_{r=1}^{b} \left[\sqrt{n}\left[\hat{\theta}^{\langle r \rangle}(n) - \theta\right] - \frac{1}{b}\sum_{r'=1}^{b} \sqrt{n}\left[\hat{\theta}^{\langle r' \rangle}(n) - \theta\right]\right]^2$$
$$\Rightarrow \frac{1}{b-1} \sum_{r=1}^{b} \left[N^{\langle r \rangle} - \frac{1}{b}\sum_{r'=1}^{b} N^{\langle r' \rangle}\right]^2 \overset{\mathcal{D}}{=} \psi^2 \frac{\chi^2_{b-1}}{b-1} \tag{A.6}$$

as $n \to \infty$ by (A.3), (A.4), and the continuous-mapping theorem, where χ^2_{b-1} is a chi-squared random variable with $b-1$ d.f., which is independent of $\sum_{r=1}^{b} N^{\langle r \rangle}$ in (A.5). Putting (A.5) and (A.6) into the middle term of (A.1) and using (A.3) yields $\sqrt{b}\,[\hat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n) - \theta]/S_b(n) \Rightarrow \psi\mathcal{N}(0,\,1)/\sqrt{\psi^2\chi^2_{b-1}/(b-1)} = \mathcal{N}(0,\,1)/\sqrt{\chi^2_{b-1}/(b-1)} \overset{\mathcal{D}}{=} T_{b-1}$ as $n \to \infty$ by the continuous-mapping theorem and the independence of the limit's numerator and denominator. Thus, (A.1) holds, so the portmanteau theorem ensures $P(B_{\mathrm{SS+rLHS},\eta,b}(n) \ge \theta) = P(\sqrt{b}\,[\hat{\theta}_{\mathrm{SS+rLHS},\eta,b}(n) - \theta]/S_b(n) \ge -\tau_{b-1,\gamma}) \to P(T_{b-1} \ge -\tau_{b-1,\gamma}) = \gamma$ as $n \to \infty$ by the continuity and symmetry of the Student-t distribution, completing the proof.

Appendix B. Proof of Theorem 3

As in Assumption A3, let $\boldsymbol{U}' = (U_1, U_2, \ldots, U_{d'_s})$ be a subvector of vector $\boldsymbol{U} = (U_1, U_2, \ldots, U_{d_s})$ of i.i.d. uniforms. It can be shown that the conditions of Theorem 5ii of [23] hold under our Assumptions A2 and A3. Thus, as established in Eq. (77) of the proof of Theorem 5ii of [23], the asymptotic variance $\sigma^2_{(s),\mathrm{SCL}}$ in (52) for scenario s satisfies

$$\sigma^2_{(s),\mathrm{SCL}} = \mathrm{Var}[E[\epsilon_{(s)}(\boldsymbol{U})|\,\boldsymbol{M}_{(s)}(\boldsymbol{U}')]] = \mathrm{Var}[\epsilon_{(s)}(\boldsymbol{U})] - E[\mathrm{Var}[\epsilon_{(s)}(\boldsymbol{U})|\,\boldsymbol{M}_{(s)}(\boldsymbol{U}')]] = \sigma^2_{(s),\mathrm{SS+LHS}} - E[\mathrm{Var}[\epsilon_{(s)}(\boldsymbol{U})|\,\boldsymbol{M}_{(s)}(\boldsymbol{U}')]] \tag{B.1}$$

by a variance decomposition and (36). Hence, putting (B.1) into (53) yields

$$\sigma^2_{\mathrm{SCL},\eta} = \sum_{s=1}^{s_0} \frac{\lambda_s^2}{\eta_s}\left(\sigma^2_{(s),\mathrm{SS+LHS}} - E[\mathrm{Var}[\epsilon_{(s)}(\boldsymbol{U})|\,\boldsymbol{M}_{(s)}(\boldsymbol{U}')]]\right) = \sigma^2_{\mathrm{SS+LHS},\eta} - \sum_{s=1}^{s_0} \frac{\lambda_s^2}{\eta_s} E[\mathrm{Var}[\epsilon_{(s)}(\boldsymbol{U})|\,\boldsymbol{M}_{(s)}(\boldsymbol{U}')]]$$

by (31), establishing the equality in (55). The inequality in (55) holds as the conditional variances in (55) are nonnegative, implying that $\sigma^2_{\mathrm{SCL},\eta} \le \sigma^2_{\mathrm{SS+LHS},\eta}$.

Appendix C. Proof of Theorem 4

We can prove Theorem 4 employing the same basic arguments applied to show Theorem 2. The only difference is that we now use the estimator (57) of $\theta_{(s)}$ from each replication r instead of the SS+rLHS estimator without CMC in (39). By (47), the response function $\mathcal{J}_{(s)}$ in (57) is bounded, so the analogue of the CLT in (A.2) also holds for each r, and the rest of the proof of Theorem 2 goes through.

References

[1] Nuclear Energy Agency Committee on the Safety of Nuclear Installations. Safety margins action plan (SMAP) – final report. Tech Rep NEA/CSNI/R(2007)9. Organization for Economic Cooperation and Development; 2007.

[2] Smith C, Rabiti C, Martineau R. Risk informed safety margin characterization (RISMC) pathway technical program plan. Tech Rep INL/EXT-11-22977. Idaho National Laboratory; 2012.

[3] Dube DA, Sherry RR, Gabor JR, Hess SM. Application of risk informed safety margin characterization to extended power uprate analysis. Reliab Eng Syst Saf 2014;129:19–28.

[4] U.S. Nuclear Regulatory Commission. Acceptance criteria for emergency core cooling systems for light-water nuclear power reactors. Title 10, Code of Federal Regulations Section 50.46 (10CFR50.46). U.S. Nuclear Regulatory Commission, Washington, DC; 2010.

[5] U.S. Nuclear Regulatory Commission. Applying statistics. U.S. Nuclear Regulatory Commission Report NUREG-1475, Rev 1. U.S. Nuclear Regulatory Commission, Washington, DC; 2011.

[6] U.S. Nuclear Regulatory Commission. Safety evaluation by the Office of the Nuclear Reactor Regulation Topical Report WCAP-16009-P. Revision 0, realistic large break LOCA evaluation methodology using automated statistical treatment of uncertainty method (ASTRUM). Tech Rep. U.S. Nuclear Regulatory Commission; 2004.

[7] Hess S, Gaertner J, Gabor J, Shanley L, Enos-Sylla L, Prasad S, Hollingsworth S. Framework for risk-informed safety margin characterization. Tech Rep 1019206. Electric Power Research Institute, Palo Alto, California; 2009.

[8] U.S. Nuclear Regulatory Commission. Best-estimate calculations of emergency core cooling performance. Nuclear Regulatory Commission Regulatory Guide 1.157. U. S. Nuclear Regulatory Commission, Washington, DC; 1989.

[9] Guba A, Makai M, Pál L. Statistical aspects of best estimate method–I. Reliab Eng Syst Saf 2003;80:217–32.

[10] Nutt WT, Wallis GB. Evaluation of nuclear safety from the outputs of computer codes in the presence of uncertainties. Reliab Eng Syst Saf 2004;83:57–77.

[11] Wilks SS. Determination of sample sizes for setting tolerance limits. Ann Math Stat 1941;12:91–6.

[12] Grabaskas D, Nakayama MK, Denning R, Aldemir T. Advantages of variance reduction techniques in establishing confidence intervals for quantiles. Reliab Eng Syst Saf 2016;149:187–203.

[13] Glasserman P. Monte Carlo methods in financial engineering. New York: Springer; 2004.

[14] Asmussen S, Glynn P. Stochastic simulation: algorithms and analysis. New York: Springer; 2007.

[15] Chu F, Nakayama MK. Confidence intervals for quantiles when applying variance-reduction techniques. ACM Trans Model Comput Simul; 36; 2012 Article 7 (25 pages plus 12–page online–only appendix).

[16] Nakayama MK. Asymptotically valid confidence intervals for quantiles and values-at-risk when applying Latin hypercube sampling. Int J Adv Syst Meas 2011;4:86–94.

[17] Szilard RH, Youngblood R, Frepoli C, Yurko JP, Swindlehrst G, Zhang H, et al. Industry application emergency core cooling system cladding acceptance criteria problem statement. Tech Rep INL/EXT-15-35073. DOE Office of Nuclear Energy, Idaho National Laboratory; 2015.

[18] Papazoglou IA. Mathematical foundations of event trees. Reliab Eng Syst Saf 1998;61:169–83.

[19] Sherry RR, Gabor JR, Hess SM. Pilot application of risk informed safety margin characterization to a total loss of feedwater event. Reliab Eng Syst Saf

2013;117:65–72.
[20] Melchers RE. Structural reliability analysis and prediction, 2nd edition. Chichester, UK: John Wiley & Sons; 1999.
[21] Iman RL. Statistical methods for including uncertainties associated with the geologic isolation of radioactive waste which allow for a comparison with licensing criteria. In: Kocher DC, editor. Proceedings of the symposium on uncertainties associated with the regulation of the geologic disposal of high-level radioactive waste. US Nuclear Regulatory Commission, Directorate of Technical Information and Document Control, Washington, DC; 1981. p. 145–157.
[22] Owen AB, Monte Carlo theory, methods and examples, Draft; 2013 (In preparation).
[23] Avramidis AN, Wilson JR. Integrated variance reduction strategies for simulation. Oper Res 1996;44:327–46.
[24] Nakayama MK. Estimating a failure probability using a combination of variance-reduction tecniques. In: Yilmaz L, Chan WKV, Moon I, Roeder TMK, Macal C, Rossetti MD, editors. Proceedings of the IEEE winter simulation conference; 2015. p. 621–32.
[25] Nakayama MK, Grabaskas D. Conditional Monte Carlo for efficient simulation for risk-informed safety margin characterization of nuclear power plants; 2016 (In preparation).
[26] Ayyub BM, Chia C-Y. Generalized conditional expectation for structural reliability assessment. Struct Saf 1992;11:131–46.
[27] Grossi P, Kunreuther H, editors. Catastrophe modeling: a new approach to managing risk. Springer, New York; 2005.
[28] Basel Committee on Banking Supervision. Basel II: international convergence of capital measurement and capital standards: a revised framework. Tech Rep. Bank for International Settlements, Basel, Switzerland ⟨http://www.bis.org/publ/bcbs107.htm⟩; 2004.
[29] Committee on the evaluation of quantification of margins and uncertainties methodology for assessing and certifying the reliability of the nuclear stockpile. Division on Engineering and Physical Sciences. Evaluation of quantification of margins and uncertainties methodology for assessing and certifying the reliability of the nuclear stockpile. Tech Rep. National Research Council of the National Academies, Washington, DC; 2008.
[30] Helton JC. Quantification of margins and uncertainties: conceptual and computational basis. Reliab Eng Syst Saf 2011;96:976–1013.
[31] Helton JC, Pilch M, Sallaberry C. Probability of loss of assured safety in systems with multiple time-dependent failure modes: representations with aleatory and epistemic uncertainty. Reliab Eng Syst Saf 2014;124:171–200.
[32] Agency NE, editor. Methods for safety assessment of geological disposal facilities for radioactive waste: outcomes of the NEA MeSA initiative. No. NEA No. 6923. Organisation for Economic Co-operation and Development, Issy-les-Moulineaux, France; 2012.
[33] Helton JC, Hansen CW, Sallaberry CJ. Conceptual structure and computational organization of the 2008 performance assessment for the proposed high-level radioactive waste repository at Yucca Mountain, Nevada. Reliab Eng Syst Saf; 122; 2014. p. 223–48.
[34] Helton JC, Anderson DR, Basabilvazo G, Jow H-N, Marietta MG. Conceptual structure of the 1996 performance assessment for the waste isolation pilot plant. Reliab Eng Syst Saf; 69; 2000. p. 151–65.
[35] Naess A, Leira BJ, Batsevych O. System reliability by enhanced Monte Carlo. Struct Saf 2009;31:349–55.
[36] Bucher C. Asymptotic sampling for high-dimensional reliability analysis. Probab Eng Mech 2009;24:504–10.
[37] Gaspar B, Soares CG. Hull girder reliability using a Monte Carlo based simulation method. Probab Eng Mech 2013;31:65–75.
[38] Au S-K, Beck JL. Estimation of small failure probabilities by subset simulation. Probab Eng Mech 2001;16:263–77.
[39] Bjerager P. Probability integration by directional simulation. J Eng Mech 1988;114(8):1285–302.
[40] Alban A, Darji H, Imamura A, Nakayama MK. Variance reduction for estimating a failure probability with multiple criteria. In: Roeder TMK, Frazier PI, Szechtman R, Zhou E, Huschka T, Chick SE, editors. Proceedings of the winter simulation conference; 2016. p. 302–13.
[41] Billingsley P. Convergence of probability measures, 2nd edition. New York: John Wiley and Sons; 1999.
[42] Nelsen RB. An introduction to copulas. New York: Springer; 1999.
[43] Billingsley P. Probability and measure, 3rd edition. New York: John Wiley and Sons; 1995.
[44] Lehmann EL. Elements of large-sample theory. New York: Springer; 1999.
[45] Orechwa Y. Comments on evaluation of nuclear safety from the outputs of computer codes in the presence of uncertainties by WT Nutt and GB Wallis. Reliab Eng Syst Saf; 87; 2005. p. 133–5.
[46] Heidtmann KD. Improved method of inclusion-exclusion applied to k-out-of-n systems. IEEE Trans Reliab 1982;R-31(1):36–40.
[47] Pál L, Makai M. Remarks on the statistical aspects of the safety analysis. Tech Rep. arxiv, 1336v1 [physics.data-an]; 6 Jun 2013.
[48] McKay MD, Conover WJ, Beckman RJ. A comparison of three methods for selecting input variables in the analysis of output from a computer code. Technometrics 1979;21:239–45.
[49] Owen AB. A central limit theorem for Latin hypercube sampling. J R Stat Soc B 1992;54:541–51.
[50] Loh W-L. On Latin hypercube sampling. Ann Stat 1996;24:2058–80.
[51] Helton JC, Davis FJ. Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. Reliab Eng Syst Saf 2003;81:23–69.
[52] Stein M. Large sample properties of simulations using Latin hypercube sampling. Technometrics 1987;32:367, correction 32:367.
[53] Goyal A, Shahabuddin P, Heidelberger P, Nicola V, Glynn PW. A unified framework for simulating Markovian models of highly dependable systems. IEEE Trans Comput 1992;C-41:36–51.
[54] Fu MC, Hu JQ. Conditional Monte Carlo: gradient estimation and optimization applications. New York: Kluwer Academic; 1997.
[55] McNeil AJ, Frey R, Embrechts P. Quantitative risk management: concepts, techniques, tools. Princeton, NJ: Princeton University Press; 2005.
[56] Demarta S, McNeil AJ. The t copula and related copulas. Int Stat Rev 2005;73(1):111–29.
[57] Cario MC, Nelson BL. Modeling and generating random vectors with arbitrary marginal distributions and correlation matrix. Tech Rep. Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL; 1997.
[58] Iman RL, Conover WJ. A distribution-free approach to inducing rank correlation among input variables. Commun Stat: Simul Comput 1982;B11:311–34.