

Certifiable Risk-Based Engineering Design Optimization

Anirban Chaudhuri*

Massachusetts Institute of Technology, Cambridge, MA, 02139, USA

Boris Kramer[†]

University of California San Diego, CA, 92093, USA

Matthew Norton[‡] Johannes O. Royset[§]

Naval Postgraduate School, Monterey, CA, 93943, USA

Karen E. Willcox[¶]

University of Texas at Austin, Austin, TX, 78712, USA

Abstract

Reliable, risk-averse design of complex engineering systems with optimized performance requires dealing with uncertainties. A conventional approach is to add safety margins to a design that was obtained from deterministic optimization. Safer engineering designs require appropriate cost and constraint function definitions that capture the *risk* associated with unwanted system behavior in the presence of uncertainties. The paper proposes two notions of certifiability. The first is based on accounting for the magnitude of failure to ensure data-informed conservativeness. The second is the ability to provide optimization convergence guarantees by preserving convexity. Satisfying these notions leads to *certifiable* risk-based design optimization (CRiBDO). In the context of CRiBDO, risk measures based on superquantile (a.k.a. conditional value-at-risk) and buffered probability of failure are analyzed. CRiBDO is contrasted with reliability-based design optimization (RBDO), where uncertainties are accounted for via the probability of failure, through a structural and a thermal design problem. A reformulation of the short column structural design problem leading to a convex CRiBDO problem is presented. The CRiBDO formulations capture more information about the problem to assign the appropriate conservativeness, exhibit superior optimization convergence by preserving properties of underlying functions, and alleviate the adverse effects of choosing hard failure thresholds required in RBDO.

1 Introduction

The design of complex engineering systems requires quantifying and accounting for risk in the presence of uncertainties. This is not only vital to ensure safety of designs but also to safeguard against costly design alterations late in the design cycle. The traditional approach is to add safety margins to compensate for uncertainties *after* a deterministic optimization is performed. This produces a sense of security, but is at best an imprecise recognition of risk and results in overly conservative designs that can limit performance. Properly accounting for risk *during* the design optimization of those systems could allow for more efficient designs. For example, payload increases for spacecraft and aircraft could be possible *without sacrificing safety*. The financial community has long recognized the superiority of specific risk measures in portfolio optimization (most importantly the conditional-value-at-risk (CVaR) pioneered by Rockafellar and Uryasev [1]), see [1, 2, 3]. In the financial context, it is understood that exposure to tail risk—rather rare events—can lead to catastrophic outcomes for companies, and adding too many “safety factors” (insurance, hedging) reduces

*Research Scientist, Department of Aeronautics and Astronautics, anirbanc@mit.edu.

[†]Assistant Professor, Department of Mechanical and Aerospace Engineering, bmkramer@ucsd.edu.

[‡]Assistant Professor, Department of Operations Research, mdnorto@gmail.com.

[§]Professor, Department of Operations Research, joroyset@nps.edu.

[¶]Director, Oden Institute for Computational Engineering and Sciences, kwillcox@oden.utexas.edu

profit. Analogously, in the engineering context, the problem is to find safe engineering designs without unnecessarily limiting performance and limiting the effects of the heuristic guesswork of choosing thresholds.

In general, there are two main issues when formulating a design optimization under uncertainty problem: (1) what to optimize and (2) how to optimize. The first issue involves deciding the design criterion, which in the context of decision theory could boil down to what type of utility function to use. What is a meaningful way of making design decisions under uncertainty? One would like to have a framework that can reflect stakeholders' preferences, but at the same time is relatively simple and can be explained to the public, to a governor, to a CEO, etc. The answer for what to optimize directly influences how you optimize. If the "what to optimize" was chosen poorly, the second issue becomes much more challenging. Design optimization of a real-world system is difficult, even in a deterministic setting, so it is essential to manage complexity as we formulate the design-under-uncertainty problem. Thus, any design criterion that preserves convexity and other desirable mathematical properties of the underlying functions is preferable as it simplifies the subsequent optimization.

This motivates us to incorporate specific mathematical measures of risk, either as a design constraint or cost function, into the design optimization formulation. To this end, we focus on two particular risk measures that have potentially superior properties: (i) superquantile/CVaR [4, 5], and (ii) buffered probability of failure (bPoF) [6]. Three immediate benefits of using these risk measures arise. First, both risk measures recognize extreme (tail) events which automatically enhances resilience. Second, they preserve convexity of underlying functions so that specialized and provably convergent optimizers can be employed. This drastically improves optimization performance. Third, superquantile and bPoF are conservative risk measures that add a buffer zone to the limiting threshold by taking into account the magnitude of failure. This can be handled by adding safety factors to the threshold; however, it has been shown before that probabilistic approaches lead to safer designs with optimized performance compared to the safety factor approach [7, 8, 9]. Superquantile/CVaR has been recently used in specific formulations in civil [10, 11], naval [12, 13] and aerospace [14, 15] engineering, as well as general PDE-constrained optimization [16, 17, 18, 19]. The bPoF risk measure has been shown to possess beneficial properties when used in optimization [6, 20, 21, 22], yet has been seldom used in engineering to-date [23, 24, 25, 26]. We contrast these above risk-based engineering design methods with the most common approach to address parametric uncertainties in engineering design, namely reliability-based design optimization (RBDO) [27, 28] which uses the probability of failure (PoF) as a design constraint. We discuss the specific advantages of using these ways of measuring risk in the design optimization cycle and their effect on the final design under uncertainty.

In this paper, we define two certifiability conditions for risk-based design optimization that can certify designs against near-failure and catastrophic failure events, and guarantee convergence to the global optimum based on preservation of convexity by the risk measures. We call the optimization formulations using risk measures satisfying any of the certifiability conditions as **Certifiable Risk-Based Design Optimization (CRiBDO)**. Risk measures satisfying both certifiability conditions lead to strongly certifiable risk-based design. We analyze superquantile and bPoF, which are examples of risk measures satisfying the certifiability conditions. We discuss how the nature of probabilistic conservativeness introduced through superquantile and bPoF makes practical sense since it is data-informed and based on the magnitude of failure. The data-informed probabilistic conservativeness of superquantiles and bPoF circumvents the guesswork associated with setting safety factors (especially, for the conceptual design phase) and transcends the limitations of setting hard thresholds for limit state functions used in PoF. This helps us move away from *being conservative blindly to being conservative to the level dictated by the data*. We compare the different risk-based design optimization formulations using a structural and a thermal design problem. For the structural design of a short column problem, we show a convex reformulation of the objective and limit state functions that leads to a convex CRiBDO formulation.

The remainder of this paper is organized as follows. We summarize the widely-used RBDO formulation in Section 2. The different risk-based optimization problem formulations along with the risk measures used in this work are described in Section 3. Section 4 explains the features of different risk-based optimization formulations through numerical experiments on the short column problem with a convex reformulation. Section 5 explores the different risk-based optimization formulations for the thermal design of a cooling fin problem with non-convex limit state. Section 6 presents the concluding remarks.

2 Reliability-based Design Optimization

In this section, we review the RBDO formulation, which uses PoF to quantify uncertainties. Let the quantity of interest of an engineering system be computed from the model $f : \mathcal{D} \times \Omega \mapsto \mathbb{R}$ as $f(\mathbf{d}, Z)$, where the inputs to the system are the n_d design variables $\mathbf{d} \in \mathcal{D} \subseteq \mathbb{R}^{n_d}$ and the n_z random variables Z with the probability distribution π . The realizations of the random variables Z are denoted by $\mathbf{z} \in \Omega \subseteq \mathbb{R}^{n_z}$. The space of design variables is denoted by \mathcal{D} and the space of random samples is denoted by Ω . The failure of the system is described by a limit state function $g : \mathcal{D} \times \Omega \mapsto \mathbb{R}$ and a critical threshold $t \in \mathbb{R}$, where, without loss of generality, $g(\mathbf{d}, \mathbf{z}) > t$ defines failure of the system. For a system under uncertainty, $g(\mathbf{d}, Z)$ is also a random variable given a particular design \mathbf{d} . The limit state function in most engineering applications requires the solution of a system of equations (such as ordinary differential equations or partial differential equations).

The most common RBDO formulation involves the use of a PoF constraint as

$$\begin{aligned} \min_{\mathbf{d} \in \mathcal{D}} \quad & \mathbb{E}[f(\mathbf{d}, Z)] \\ \text{subject to} \quad & p_t(g(\mathbf{d}, Z)) \leq 1 - \alpha_T, \end{aligned} \quad (1)$$

where $\alpha_T \in [0, 1]$ is the target reliability (i.e., $1 - \alpha_T$ is the target PoF) and the PoF is defined via the limit state function g and the failure threshold t as $p_t(g(\mathbf{d}, Z)) := \mathbb{P}[g(\mathbf{d}, Z) > t]$. The RBDO problem (1) designs a system with optimal mean characteristics, in terms of $f(\mathbf{d}, Z)$, such that it maintains a reliability of at least α . Note, however, that PoF has no information about the magnitude of the failure event as it is merely a measure of the set $\{g(\mathbf{d}, Z) > t\}$, see Figure 1.




Figure 1: Illustration for PoF indicated by the area of the shaded region.

For our upcoming discussion, it is helpful to point out that a constraint on the PoF is equivalent to a constraint on the α -quantile. The α -quantile, also known as the value-at-risk at level α , is defined in terms of the inverse cumulative distribution function of the limit state function $F_{g(\mathbf{d}, Z)}^{-1}$ as

$$Q_\alpha[g(\mathbf{d}, Z)] := F_{g(\mathbf{d}, Z)}^{-1}(\alpha). \quad (2)$$

PoF and Q_α are natural counterparts that are measures of the tail of the distribution of $g(\mathbf{d}, Z)$. When the largest $100(1 - \alpha)\%$ outcomes are the ones of interest (i.e., failed cases), the quantile is a measure of minimum value within the set of these tail events. When one knows that outcomes larger than a given threshold t are of interest, PoF provides a measure of the frequency of these “large” events. This equivalence of PoF and Q_α risk constraints is illustrated in Figure 2. In the context of our optimization problem, using the same value of t and α_T , (1) can be written equivalently as

$$\begin{aligned} \min_{\mathbf{d} \in \mathcal{D}} \quad & \mathbb{E}[f(\mathbf{d}, Z)] \\ \text{subject to} \quad & Q_{\alpha_T}[g(\mathbf{d}, Z)] \leq t. \end{aligned} \quad (3)$$

The most elementary method (although, inefficient) for estimating PoF is Monte Carlo (MC) simulation when dealing with nonlinear limit state functions. The MC estimate of the PoF for a given design \mathbf{d} is

$$\hat{p}_t(g(\mathbf{d}, Z)) = \frac{1}{m} \sum_{i=1}^m \mathbb{I}_{g(\mathbf{d})}(\mathbf{z}_i), \quad (4)$$




Figure 2: Illustration of equivalence of PoF (shown by the shaded region) and Q_α showing that the two quantities converge at the constraint threshold when the reliability constraint is active.

where $\mathbf{z}_1, \dots, \mathbf{z}_m$ are m samples distributed according to π , $\mathcal{G}(\mathbf{d}) = \{\mathbf{z} \mid g(\mathbf{d}, \mathbf{z}) > t\}$ is the failure set, and $\mathbb{I}_{\mathcal{G}(\mathbf{d})} : \Omega \rightarrow \{0, 1\}$ is the indicator function defined as

$$\mathbb{I}_{\mathcal{G}(\mathbf{d})}(\mathbf{z}) = \begin{cases} 1, & \text{if } \mathbf{z} \in \mathcal{G}(\mathbf{d}) \\ 0, & \text{else.} \end{cases} \quad (5)$$

The MC estimator is unbiased with the variance being $p_t(1 - p_t)/m$. The PoF estimation requires sampling from the tails of the distribution, which can often make MC estimators expensive. A wealth of literature exists for methods that have been developed to deal with the computational complexity of PoF estimation and the RBDO problem. First, reliability index methods (e.g., FORM, SORM, etc. [29, 30]) geometrically approximate the limit state function to reduce the computational effort of PoF estimation. However, when the limit state function is nonlinear, the reliability index method could lead to inaccuracies in the estimate. Second, MC variance reduction techniques such as importance sampling [31, 32, 33], adaptive importance sampling [34, 35, 36, 37, 38, 39], and multifidelity approaches [40, 41, 42] offer computational advantages. While the decay rate of the MC estimator cannot be improved upon, the variance of the MC estimator can be reduced, which offers computational advantages in that fewer (suitably chosen) MC samples are needed to obtain accurate PoF estimates. Third, adaptive data-driven surrogates for the limit state failure boundary identification can improve computational efficiency for the RBDO problem [43, 44, 45]. Fourth, bi-fidelity RBDO methods [46, 47] and recent multifidelity/multi-information-source methods for the PoF estimate [48, 49, 44] and the RBDO problem [50, 51] have led to significant computational savings.

Although significant research has been devoted to PoF and RBDO, PoF as a risk measure does not factor in how catastrophic is the failure and thus, lacks resiliency. In other words, PoF neglects the magnitude of failure of the system and instead encodes a hard threshold via a binary function evaluation. We describe below this drawback of PoF.

Remark 1 (Limitations of hard-thresholding) *To motivate the upcoming use of risk measures, we take a closer look at the limit state function g and its use to characterize failure events. In the standard setting, a failure event is characterized by a realization of Z for some fixed design \mathbf{d} that leads to $g(\mathbf{d}, \mathbf{z}) > t$. However, this hard-threshold characterization of system failure potentially ignores important information quantified by the magnitude of $g(\mathbf{d}, \mathbf{z})$ and PoF fails to promote resilience, i.e., no distinction between bad and very bad. Let us consider a structure with $g(\mathbf{d}, \mathbf{z})$ being the load and the threshold t being the allowable strength. There may be a large difference between the event $g(\mathbf{d}, \mathbf{z}) = t + .01\text{kN}$ and $g(\mathbf{d}, \mathbf{z}) = t + 100\text{kN}$, the latter characterizing a catastrophic system failure. This is not captured when considering system failure only as a binary decision with a hard threshold. Similarly, one could also consider events $g(\mathbf{d}, \mathbf{z}) = t - .01\text{kN}$ and $g(\mathbf{d}, \mathbf{z}) = t - 100\text{kN}$. A hard-threshold assessment deems both of these events as non-failure events, even though $g(\mathbf{d}, \mathbf{z}) = t - .01\text{kN}$ is clearly a near-failure event compared to $g(\mathbf{d}, \mathbf{z}) = t - 100\text{kN}$. A hard-threshold characterization of failure would potentially overlook these important near-failure events and consider them as safe realizations of g . In reality, failure events do not usually occur using a hard-threshold rule. Even if they do, determination of the true threshold will also involve uncertainty, blending statistical estimation, expert knowledge, and system models. Therefore, the choice of threshold should be involved in any discussion of measures of failure risk and we analyze later in Remark 6, the advantage of the data-informed thresholding property of certain risk measures as compared to hard-thresholding. In addition, encoding magnitude of failure*

can help distinguish between designs with same PoF (see example 1 in Ref. [6]). As we show in the next section, superquantile and bPoF do not have this deficiency.

In the engineering community, PoF has been the preferred choice. Using PoF and RBDO offers some specific advantages starting with the simplicity of the risk measure and the natural intuition behind formulating the optimization problems, which is a major reason leading to the rich literature on this topic as noted before. Another advantage of PoF is the invariance to nonlinear reformulation for the limit state function. For example, let z_1 be a random load and z_2 be a random strength of a structure. Then the PoF would be the same regardless if the limit state function is defined as $z_1 - z_2$ or $z_1/z_2 - 1$. Since α -quantile leads to an equivalent formulation as PoF, both PoF and α -quantile formulations have this invariance for continuous distributions. However, there are several potential issues when using PoF as the risk measure for design optimization under uncertainty as noted below.

Remark 2 (Optimization considerations) *While there are several advantages of using PoF and RBDO, there are several potential drawbacks. First, PoF is not necessarily a convex function w.r.t. design variables \mathbf{d} even when the underlying limit state function is convex w.r.t. \mathbf{d} . Thus, we cannot formulate a convex optimization problem even when underlying functions f and g are convex w.r.t. \mathbf{d} . This is important because convexity guarantees convergence of standard and efficient algorithms to a globally optimal design under minimal assumptions since every local optimum is a global optimum in that case. Second, the computation of PoF gradients can be ill-conditioned, so traditional gradient-based optimizers that require accurate gradient evaluations tend to face challenges. While PoF is differentiable for the specific case when \mathbf{d} only contains parameters of the distribution of Z , such as mean and standard deviation, PoF is in general not a differentiable function. Consequently, PoF gradients may not exist and when using approximate methods, such as finite difference, the accuracy of the PoF gradients could be poor. Some of these drawbacks can be addressed by using other methods for estimating the PoF gradients, but they have been developed under potentially restrictive assumptions [52, 53, 54], which might not be easily verifiable for practical problems. Third, PoF can suffer from sensitivity to the failure threshold due to it being a discontinuous function w.r.t. threshold t . Since the choice of failure threshold could be uncertain, one would ideally prefer to have a measure of risk that is less sensitive to small changes in t . We further expand on this issue in Remark 4.*

3 Certifiable Risk-Based Design Optimization

Design optimization with a special class of risk measures can provide certifiable designs and algorithms. We first present two notions of certifiability in risk-based optimization in Section 3. 3.1. We then discuss two specific risk measures, superquantile in Section 3. 3.2 and bPoF in Section 3. 3.3, that satisfy these notions of certifiability.

3.1 Certifiability in risk-based design optimization

Risk in an engineering context can be quantified in several ways and the choice of risk measure, and its use as a cost or constraint, influences the design. We focus on a class of risk measures that can satisfy the following two *certifiability conditions*:

1. *Data-informed conservativeness*: Risk measures that take the magnitude of failure into account to decide the level of conservativeness required can certify the designs against near-failure and catastrophic failure events leading to increased resilience. The obtained designs can overcome the limitations of hard thresholding and are certifiably risk-averse against a continuous range of failure modes. In typical engineering problems, the limit state function distributions are not known and the information about the magnitude of failure is encoded through the generated data, thus making the conservativeness data-informed.
2. *Optimization convergence and efficiency*: Risk measures that preserve the convexity of underlying limit state functions (and/or cost functions) lead to convex risk-based optimization formulations. The resulting optimization problem is better behaved than a non-convex problem and can be solved more efficiently. Thus, one can find the design that is certifiably optimal in comparison with all alternate

designs at reduced computational cost. In general, the risk measure preserves the convexity of the limit state function, such that the complexity of the optimization under uncertainty problem remains similar to the complexity of the deterministic optimization problem using the limit state function.

We denote the risk-based design optimization formulations that use risk measures satisfying any of the two certifiability conditions as **Certifiable Risk-Based Design Optimization** (CRiBDO). Note that designs obtained through RBDO do not satisfy either of the above conditions since using PoF as the risk measure cannot guard against near-threshold or catastrophic failure events, see Remark 1, and cannot certify the design to be a global optimum, see Remark 2. Accounting for the magnitude of failure is critical to ensure appropriate conservativeness in CRiBDO designs and additionally, preservation of convexity is useful for optimizer efficiency and convergence guarantees to a globally optimal design. The optimization formulations satisfying both the conditions lead to *strongly certifiable* risk-based designs. In general engineering applications, the convexity condition is difficult to satisfy but encapsulates an ideal situation, highlighting the importance of research in creating (piece-wise) convex approximations for physical problems. In Sections 3. 3.2 and 3. 3.3, we discuss the properties of two particular risk measures, superquantile and bPoF, that lead to certifiable risk-based designs and have the potential to be strongly certifiable when underlying functions are convex. Although we focus on these two particular risk measures in this work, other measures of risk could also be used to produce certifiable risk-based designs, see [55, 56, 10].

3.2 Superquantile-based design optimization

This section describes the concept of superquantiles and associated risk-averse optimization problem formulations. Superquantiles emphasize tail events, and from an engineering perspective it is important to manage such tail risks.

3.2.1 Risk measure: superquantile

Intuitively, superquantiles can be understood as a tail expectation, or an average over a portion of worst-case outcomes. Given a fixed design \mathbf{d} and a distribution of potential outcomes $g(\mathbf{d}, Z)$, the superquantile at level $\alpha \in [0, 1]$ is the expected value of the largest $100(1 - \alpha)\%$ realizations of $g(\mathbf{d}, Z)$. In the literature, several other terms, such as CVaR and expected shortfall, have been used interchangeably with superquantile. We prefer the term superquantile because of its inherent connection with the long existing statistical quantity of quantiles and it being application agnostic.

The definition of α -superquantile is based on the α -quantile $Q_\alpha [g(\mathbf{d}, Z)]$ from Equation (2). The α -superquantile \bar{Q}_α can be defined as

$$\bar{Q}_\alpha [g(\mathbf{d}, Z)] := Q_\alpha [g(\mathbf{d}, Z)] + \frac{1}{1 - \alpha} \mathbb{E} \left[[g(\mathbf{d}, Z) - Q_\alpha [g(\mathbf{d}, Z)]]^+ \right], \quad (6)$$

where \mathbf{d} is the given design and $[c]^+ := \max\{0, c\}$. The expectation in the second part of the right hand side of Equation (6) can be interpreted as the expectation of the tail of the distribution exceeding the α -quantile. The α -superquantile can be seen as the sum of the α -quantile and a non-negative term and thus, $\bar{Q}_\alpha [g(\mathbf{d}, Z)]$ is a quantity higher (as indicated by “super”) than $Q_\alpha [g(\mathbf{d}, Z)]$. It follows from the definition that $\bar{Q}_\alpha [g(\mathbf{d}, Z)] \geq Q_\alpha [g(\mathbf{d}, Z)]$. When the cumulative distribution of $g(\mathbf{d}, Z)$ is continuous for any \mathbf{d} , we can also view $\bar{Q}_\alpha [g(\mathbf{d}, Z)]$ as the conditional expectation of $g(\mathbf{d}, Z)$ with the condition that $g(\mathbf{d}, Z)$ is not less than $Q_\alpha [g(\mathbf{d}, Z)]$, i.e., $\bar{Q}_\alpha [g(\mathbf{d}, Z)] = \mathbb{E} [g(\mathbf{d}, Z) \mid g(\mathbf{d}, Z) \geq Q_\alpha [g(\mathbf{d}, Z)]]$ [5]. We also note that by definition [4]

$$\begin{aligned} \text{for } \alpha = 0, \quad \bar{Q}_0 [g(\mathbf{d}, Z)] &= \mathbb{E} [g(\mathbf{d}, Z)], \text{ and} \\ \text{for } \alpha = 1, \quad \bar{Q}_1 [g(\mathbf{d}, Z)] &= \text{ess sup } g(\mathbf{d}, Z), \end{aligned} \quad (7)$$

where $\text{ess sup } g(\mathbf{d}, Z)$ is the essential supremum, i.e., the lowest value that $g(\mathbf{d}, Z)$ doesn't exceed with probability 1.

Figure 3 illustrates the \bar{Q}_α risk measure for two differently shaped, generic distributions of the limit state function. The figure shows that the magnitude of $\bar{Q}_\alpha - Q_\alpha$ (or the induced conservativeness) changes with the underlying distribution. Algorithm 1 describes standard MC sampling for approximating \bar{Q}_α . The second term on the right hand side in Equation (8) is a MC estimate of the expectation in Equation (6).




Figure 3: Illustration for \bar{Q}_α on two generic distributions: expectation of the worst-case $1 - \alpha$ outcomes shown in blue is $\bar{Q}_\alpha [g(\mathbf{d}, Z)]$.

Algorithm 1 Sampling-based estimation of Q_α and \bar{Q}_α .

Input: m i.i.d. samples $\mathbf{z}_1, \dots, \mathbf{z}_m$ of random variable Z , design variable \mathbf{d} , risk level $\alpha \in (0, 1)$, limit state function $g(\mathbf{d}, Z)$.

Output: Sample approximations $\hat{Q}_\alpha [g(\mathbf{d}, Z)]$, $\hat{\bar{Q}}_\alpha [g(\mathbf{d}, Z)]$.

- 1: Evaluate limit state function at the samples to get $g(\mathbf{d}, \mathbf{z}_1), \dots, g(\mathbf{d}, \mathbf{z}_m)$.
- 2: Sort values of limit state function in descending order and relabel the samples so that

$$g(\mathbf{d}, \mathbf{z}_1) > g(\mathbf{d}, \mathbf{z}_2) > \dots > g(\mathbf{d}, \mathbf{z}_m).$$

- 3: Find the index $k_\alpha = \lceil m(1 - \alpha) \rceil$ to estimate $\hat{Q}_\alpha [g(\mathbf{d}, Z)] \leftarrow g(\mathbf{d}, \mathbf{z}_{k_\alpha})$.
- 4: Estimate

$$\hat{\bar{Q}}_\alpha [g(\mathbf{d}, Z)] = \hat{Q}_\alpha [g(\mathbf{d}, Z)] + \frac{1}{m(1 - \alpha)} \sum_{j=1}^m [g(\mathbf{d}, \mathbf{z}_j) - \hat{Q}_\alpha [g(\mathbf{d}, Z)]]^+. \quad (8)$$

3.2.2 Optimization problem: superquantiles as constraint

As noted before, the PoF constraint of the RBDO problem in (1) can be viewed as a Q_α constraint (as seen in (3)). The PoF constraint (and thus the Q_α constraint) does not consider the magnitude of the failure events, but only whether they are larger than the failure threshold. This could be a potential drawback for engineering applications. On the other hand, a \bar{Q}_α constraint considers the magnitude of the failure events by specifically constraining the expected value of the largest $100(1 - \alpha)\%$ realizations of $g(\mathbf{d}, Z)$. Additionally, depending upon the actual construction of $g(\mathbf{d}, \mathbf{z})$ and the accuracy of the sampling procedure, the \bar{Q}_α constraint may have numerical advantages over the Q_α constraint when it comes to optimization as discussed later. In particular, we have the optimization problem formulation

$$\begin{aligned} \min_{\mathbf{d} \in \mathcal{D}} \quad & \mathbb{E} [f(\mathbf{d}, Z)] \\ \text{subject to} \quad & \bar{Q}_{\alpha_T} [g(\mathbf{d}, Z)] \leq t, \end{aligned} \quad (9)$$

where α_T is the desired reliability level given the limit state failure threshold t . The \bar{Q}_α -based formulation typically leads to a more conservative design than when PoF is used. This can be observed by noting that $\bar{Q}_{\alpha_T} [g(\mathbf{d}, Z)] \leq t \implies Q_{\alpha_T} [g(\mathbf{d}, Z)] \leq t \iff p_t(g(\mathbf{d}, Z)) \leq 1 - \alpha_T$. Therefore, if the design satisfies the \bar{Q}_{α_T} constraint, then the design will also satisfy the related PoF constraint. Additionally, since the \bar{Q}_{α_T} constraint ensures that the average of the $(1 - \alpha_T)$ tail is no larger than t , it is likely that the probability of exceeding t (PoF) is strictly smaller than $1 - \alpha_T$ and is thus a conservative design for target reliability of α_T . Intuitively, this conservatism comes from the fact that \bar{Q}_{α_T} considers the magnitude of the worst failure events.

The formulation with \bar{Q}_{α_T} as the constraint is useful when the designer is unsure about the failure

boundary location for the problem but requires a certain level of reliability from the design. For example, consider the case where the failure is defined as maximum stress of a structure not exceeding a certain value. However, the designers cannot agree on the cut-off value for stress but can agree on the desired level of reliability they want. One can use this formulation to design a structure with a given reliability $(1 - \alpha_T)$ while constraining a conservative estimate of the cut-off value (\bar{Q}_{α_T}) on the stress.

Remark 3 (Convexity in \bar{Q}_α -based optimization) *It can be shown that \bar{Q}_α can be written in the form of an optimization problem [5] as*

$$\bar{Q}_\alpha [g(\mathbf{d}, Z)] = \min_{\gamma \in \mathbb{R}} \gamma + \frac{1}{1 - \alpha} \mathbb{E} [[g(\mathbf{d}, Z) - \gamma]^+], \quad (10)$$

where \mathbf{d} is the given design, γ is an auxiliary variable, and $[c]^+ := \max\{0, c\}$. At the optimum, $\gamma^* = Q_\alpha [g(\mathbf{d}, Z)]$. Using Equation (10), the formulation (9) can be reduced to an optimization problem involving only expectations as given by

$$\begin{aligned} & \min_{\gamma \in \mathbb{R}, \mathbf{d} \in \mathcal{D}} \mathbb{E} [f(\mathbf{d}, Z)] \\ \text{subject to} \quad & \gamma + \frac{1}{1 - \alpha_T} \mathbb{E} [[g(\mathbf{d}, Z) - \gamma]^+] \leq t. \end{aligned} \quad (11)$$

The formulation (11) is a convex optimization problem when $g(\mathbf{d}, Z)$ and $f(\mathbf{d}, Z)$ are convex in \mathbf{d} since $[\cdot]^+$ is a convex function and preserves the convexity of the limit state function. Another advantage of (11), as outlined in Ref. [5], is that the nonlinear part of the constraint, $\mathbb{E} [[g(\mathbf{d}, Z) - \gamma]^+]$, can be reformulated as a set of convex (linear) constraints if $g(\mathbf{d}, Z)$ is convex (linear) in \mathbf{d} and has a discrete (or empirical) distribution with the distribution of Z being independent of \mathbf{d} ¹. Specifically, consider a MC estimate where $\mathbf{z}_i, i = 1, \dots, m$ are m samples from probability distribution π . Then, using auxiliary variables $b_i, i = 1, \dots, m$ to define $\mathbf{b} = \{b_1, \dots, b_m\}$, we can reformulate (11) as

$$\begin{aligned} & \min_{\gamma \in \mathbb{R}, \mathbf{b} \in \mathbb{R}^m, \mathbf{d} \in \mathcal{D}} \mathbb{E} [f(\mathbf{d}, Z)] \\ \text{subject to} \quad & \gamma + \frac{1}{m(1 - \alpha_T)} \sum_{i=1}^m b_i \leq t, \\ & g(\mathbf{d}, z_i) - \gamma \leq b_i, i = 1, \dots, m, \\ & b_i \geq 0, i = 1, \dots, m. \end{aligned} \quad (12)$$

The formulation (12) is a linear program when $g(\mathbf{d}, Z)$ and $f(\mathbf{d}, Z)$ are linear in \mathbf{d} .

As noted in Remark 3, the formulations in (11) and (12) are convex (or linear) only when the underlying functions $g(\mathbf{d}, Z)$ and $f(\mathbf{d}, Z)$ are convex (or linear) in \mathbf{d} . However, the advantages and possibility of such formulations indicates that one can achieve significant gains by investing in convex (or linear) approximations for the underlying functions.

3.2.3 Optimization problem: superquantiles as objective

The α -superquantile \bar{Q}_α naturally arises as a replacement for Q_α in the constraint, but it can also be used as the objective function in the optimization problem formulation. For example, in PDE-constrained optimization, superquantiles have been used in the objective function [16, 17]. The optimization formulation is

$$\begin{aligned} & \min_{\mathbf{d} \in \mathcal{D}} \bar{Q}_{\alpha_T} [g(\mathbf{d}, Z)] \\ \text{subject to} \quad & \bar{Q}_{\beta_T} [f(\mathbf{d}, Z)] \leq C_T, \end{aligned} \quad (13)$$

¹In the case where π depends upon \mathbf{d} , one can perform optimization by using sampling-based estimators for the gradient of \bar{Q}_α [57, 58].

where α_T and β_T are the desired risk levels for g and f respectively, and C_T is a threshold on the quantity of interest f . This is a useful formulation when it is easier to define a threshold on the quantity of interest than deciding a risk level for the limit state function. For example, if the quantity of interest is the cost of manufacturing a rocket engine, one can specify a budget constraint and use the above formulation. The solution of this optimization formulation would result in the safest rocket engine design such that the expected budget does not exceed the given budget.

3.2.4 Discussion on superquantile-based optimization

From an optimization perspective, an important feature of \overline{Q}_α is that it preserves convexity of the function it is applied to, i.e., the limit state function or cost function. \overline{Q}_α -based formulations can lead to well-behaved convex optimization problems that allows one to provide convergence guarantees as described in Remark 3. The reformulation offers a major advantage, since an optimization algorithm can work directly on the limit state function without passing through an indicator function. This preserves the convexity and other mathematical properties of the limit state function. \overline{Q}_α also takes the magnitude of failure into account, which makes it more informative and resilient compared to PoF and builds in data-informed conservativeness.

As noted in [59], \overline{Q}_α estimators are less stable than estimators of Q_α since rare, large magnitude tail samples can have large effect on the sample estimate. This is more prevalent when the distribution of the random quantity is fat-tailed. Thus, there is a need for more research to develop efficient algorithms for \overline{Q}_α estimation. Despite offering convexity, a drawback of \overline{Q}_α is that it is non-smooth, and a direct \overline{Q}_α -based optimization would require either non-smooth optimization methods, for example variable-metric algorithms [60], or gradient-free methods. Note that smoothed approximations exist [16, 23], which significantly improve optimization performance. In addition, the formulation (12) offers a smooth alternative.

As noted in Remark 3, \overline{Q}_α -based formulations can be further reduced to a linear program. The formulation in (12) increases the dimensionality of the optimization problem from n_d+1 to n_d+m+1 , where m is the number of MC samples, which poses an issue when the number of MC samples is large. However, formulation (12) has mostly linear constraints and can also be completely converted into a linear program by using a linear approximation for $g(\mathbf{d}, z_i)$ (following similar ideas as reliability index methods described in Section 1). There are extremely efficient methods for finding solutions to linear programs even for high-dimensional problems.

3.3 bPoF-based design optimization

Buffered probability of failure was first introduced by Rockafellar and Royset [6] as an alternative to PoF. This section describes bPoF and the associated optimization problem formulations. When used as constraints, bPoF and superquantile lead to equivalent optimization formulations but bPoF provides an alternative interpretation of the \overline{Q}_α constraint that is, arguably, more natural for applications dealing with constraints in terms of failure probability instead of constraints involving quantiles. When considered as an objective function, bPoF and superquantile lead to different optimal design solutions.

3.3.1 Risk measure: bPoF

The bPoF is an alternate measure of reliability which adds a buffer to the traditional PoF. The definition of bPoF at a given design \mathbf{d} is based on the superquantile as given by

$$\overline{p}_t(g(\mathbf{d}, Z)) := \begin{cases} \{1 - \alpha \mid \overline{Q}_\alpha[g(\mathbf{d}, Z)] = t\}, & \text{if } \overline{Q}_0[g(\mathbf{d}, Z)] < t < \overline{Q}_1[g(\mathbf{d}, Z)] \\ 0, & \text{if } t \geq \overline{Q}_1[g(\mathbf{d}, Z)] \\ 1, & \text{otherwise.} \end{cases} \quad (14)$$

The domains of the threshold t in Equation (14) can be interpreted in more intuitive terms using Equation (7) for $\overline{Q}_0[g(\mathbf{d}, Z)]$ and $\overline{Q}_1[g(\mathbf{d}, Z)]$. The relationship between superquantiles and bPoF in the first condition in Equation (14) can also be viewed in the same way as that connecting α -quantile and PoF by recalling that

$$Q_\alpha[g(\mathbf{d}, Z)] \leq t \iff p_t(g(\mathbf{d}, Z)) \leq 1 - \alpha \text{ and here, } \overline{Q}_\alpha[g(\mathbf{d}, Z)] \leq t \iff \overline{p}_t(g(\mathbf{d}, Z)) \leq 1 - \alpha. \quad (15)$$

To make the concept of *buffer* concrete, we further analyze the case in the first condition in Equation (14) when $t \in (\overline{Q}_0[g(\mathbf{d}, Z)], \overline{Q}_1[g(\mathbf{d}, Z)])$ and $g(\mathbf{d}, Z)$ is a continuous random variable, which leads to $\bar{p}_t(g(\mathbf{d}, Z)) = \{1 - \alpha \mid \overline{Q}_\alpha[g(\mathbf{d}, Z)] = t\}$. Using the definition of quantiles from Equation (2) and its connection with superquantiles (see Equation (6) and Figure 3), we can see that $1 - \alpha = \mathbb{P}[g(\mathbf{d}, Z) \geq Q_\alpha[g(\mathbf{d}, Z)]]$. This leads to another definition of bPoF in terms of probability of exceeding a quantile given the condition on α as

$$\bar{p}_t(g(\mathbf{d}, Z)) = \mathbb{P}[g(\mathbf{d}, Z) \geq Q_\alpha[g(\mathbf{d}, Z)]] = 1 - \alpha, \quad \text{where } \alpha \text{ is such that } \overline{Q}_\alpha[g(\mathbf{d}, Z)] = t. \quad (16)$$

We know that superquantiles are conservative as compared to quantiles (Section 3.3.2.3.2.1), which leads to $Q_\alpha \leq t$ since $\overline{Q}_\alpha = t$. Thus, Equation (16) can be split as a sum of PoF and the probability of *near-failure* as

$$\bar{p}_t(g(\mathbf{d}, Z)) = \mathbb{P}[g(\mathbf{d}, Z) > t] + \mathbb{P}[g(\mathbf{d}, Z) \in [Q_\alpha[g(\mathbf{d}, Z)], t]] = p_t(g(\mathbf{d}, Z)) + \mathbb{P}[g(\mathbf{d}, Z) \in [\lambda, t]], \quad (17)$$

where $\lambda = Q_\alpha[g(\mathbf{d}, Z)]$. The value of λ is affected by the condition on α through superquantiles (see Equation (16)) and takes into account the frequency and magnitude of failure. Thus, the near-failure region $[\lambda, t]$ is determined by the frequency and magnitude of tail events around t and can be intuitively seen as the *buffer* on top of the PoF. An illustration of the bPoF risk measure is shown in Figure 4. Algorithm 2 describes standard MC sampling for estimating bPoF. Note that all the quantities discussed in this work (PoF, superquantiles, and bPoF) can be viewed as expectations. Estimating them via Monte Carlo simulation therefore yields estimates whose error decreases with the rate $1/\sqrt{\text{number of samples}}$. All the estimates suffer from an increasing constant associated with the estimator variance as one moves further out in the tail, i.e., larger threshold or larger α . The computational effort can be reduced for any of the risk measures by using Monte Carlo variance reduction strategies.




Figure 4: Illustration for bPoF: for a given threshold t , PoF equals the area in *red* while bPoF equals the combined area in *red and blue*.

In general, we can see that for any design \mathbf{d} ,

$$\bar{p}_t(g(\mathbf{d}, Z)) \geq p_t(g(\mathbf{d}, Z)). \quad (18)$$

Through Equation (17), we can see that the conservatism of bPoF comes from the data-dependent mechanism that selects the conservative threshold $\lambda \leq t$, which acts to establish a *buffer* zone. If realizations of $g(\mathbf{d}, Z)$ beyond t are very large (potentially catastrophic failures), λ will need to be smaller (making bPoF bigger) to drive the expectation beyond λ to t . Thus, the larger bPoF serves to account for not only the frequency of failure events, but also their magnitude. The bPoF also accounts for the frequency of *near-failure* events that have magnitude below, but very close to t . If there are a large number of near-failure events, bPoF will take this into account, since it will be included in the λ -tail which must have average equal to t . Thus, the bPoF is a conservative estimate of the PoF for any design \mathbf{d} and carries more information about failure than PoF since it takes into consideration the magnitude of failure. It has been shown that for exponential distribution of the limit state function, the bPoF is $e \approx 2.718$ times the PoF [61]. However, the degree of conservativeness of bPoF w.r.t. PoF is dependent on the distribution of $g(\mathbf{d}, Z)$, which is typically not known in closed-form in engineering applications.

Algorithm 2 Sampling-based estimation of bPoF.

Input: m i.i.d. samples $\mathbf{z}_1, \dots, \mathbf{z}_m$ of random variable Z , design variable \mathbf{d} , failure threshold t , and limit state function $g(\mathbf{d}, Z)$.

Output: Sample approximation $\widehat{p}_t(g(\mathbf{d}, Z))$.

- 1: Evaluate limit state function at the samples to get $g(\mathbf{d}, \mathbf{z}_1), \dots, g(\mathbf{d}, \mathbf{z}_m)$.
- 2: Sort values of limit state function in descending order and relabel the samples so that

$$g(\mathbf{d}, \mathbf{z}_1) > g(\mathbf{d}, \mathbf{z}_2) > \dots > g(\mathbf{d}, \mathbf{z}_m).$$

- 3: $c = g(\mathbf{d}, \mathbf{z}_1)$ ▷ Initialize superquantile estimate
 - 4: $k = 1$
 - 5: **while** $c \geq t$ **do** ▷ Check if superquantile estimate equals threshold
 - 6: $k \leftarrow k + 1$
 - 7: $c = \frac{1}{k} \sum_{i=1}^k g(\mathbf{d}, \mathbf{z}_i)$ ▷ Update superquantile estimate
 - 8: **end while**
 - 9: Estimate bPoF as $\widehat{p}_t(g(\mathbf{d}, Z)) \approx \frac{k-1}{m}$ ▷ Estimate bPoF as $1 - \alpha$ when $c \approx t$
-

Remark 4 (Continuity of bPoF w.r.t. threshold) *In practice, thresholds are sometimes set by regulatory commissions, informed by industry standards; see chapter 18 in Ref. [62] for a discussion on code calibration. As discussed before, the data-informed conservativeness of bPoF reduces the adverse effects of poorly chosen thresholds by building a buffer around the threshold t . Another issue with poorly set thresholds is that the values could change as one learns more about the system. In such cases, continuity of the risk measure w.r.t. the threshold becomes important. bPoF is continuous w.r.t. the threshold but PoF is not. Consequently, if an engineer makes small changes to the threshold t , then it can have significant effects on the resulting design when PoF is used in the optimization formulation. On the other hand, small changes in t will only have small effect on the bPoF-based optimal design due to bPoF being continuous w.r.t. t . The following example illustrates the continuity of bPoF vs PoF w.r.t. the threshold. Let X be a random variable with finite distribution probability mass function given by*

$$\mathbb{P}(X = x) = \begin{cases} 0.8, & \text{if } x = -1 \\ 0.1, & \text{if } x = 0 \\ 0.1, & \text{if } x = 1, \end{cases}$$

which is visualized in Figure 5(a). For this simple distribution, one can derive the PoF and bPoF analytically for any given threshold t . The PoF values for different values of t are

$$p_t = \begin{cases} 1, & \text{if } t < -1 \\ 0.2, & \text{if } t \in [-1, 0) \\ 0.1, & \text{if } t \in [0, 1) \\ 0, & \text{if } t \geq 1, \end{cases}$$

which is clearly not continuous in t . The bPoF values for different values of t are

$$\bar{p}_t = \begin{cases} 1, & \text{if } t < -0.7 \\ 0.3/(t+1), & \text{if } t \in [-0.7, 0.5) \\ 0.1/t, & \text{if } t \in [0.5, 1) \\ 0, & \text{if } t \geq 1, \end{cases}$$

which is continuous in t on the interval $(-\infty, \bar{Q}_1) = (-\infty, 1)$. The PoF and bPoF values as a function of the threshold t are plotted in Figure 5(b) showing the continuity of bPoF in t . In a similar way, superquantiles \bar{Q}_α are continuous in α but quantiles Q_α are not.

3.3.2 Optimization problem: bPoF as constraint

One of the advantages of bPoF, which provides data-informed conservativeness, is the intuitive relatability to the widely used PoF. This helps in easy transition from PoF-based formulations to bPoF-based formulations.




Figure 5: Illustrating (a) the probability mass function of X and (b) the continuity of bPoF w.r.t. changing threshold values as compared to discontinuous nature of PoF.

Consider the optimization problem (1) with the PoF constraint replaced by the bPoF constraint,

$$\begin{aligned} \min_{\mathbf{d} \in \mathcal{D}} \quad & \mathbb{E}[f(\mathbf{d}, Z)] \\ \text{subject to} \quad & \bar{p}_t(g(\mathbf{d}, Z)) \leq 1 - \alpha_T. \end{aligned} \quad (19)$$

Just as a PoF constraint is equivalent to a \bar{Q}_α constraint, it can be shown that the bPoF constraint formulation described above is equivalent to a \bar{Q}_α constraint (see Equation (15)). It can be observed that the bPoF-based formulation (19) is equivalent to the \bar{Q}_α -based optimization formulation (9) by noting that the bPoF constraint being active implies that the \bar{Q}_α constraint is also active, i.e., $\bar{p}_t(g(\mathbf{d}, Z)) = 1 - \alpha_T \implies \bar{Q}_{\alpha_T}[g(\mathbf{d}, Z)] = t$. However, formulation (19) is useful when considered in the context of interpretability w.r.t. the originally intended PoF reliability constraint along with the data-informed conservative buffer provided by bPoF. In engineering applications, the exact failure threshold is often uncertain and chosen by a subject matter expert. Thus, it is beneficial that bPoF can provide a reliability constraint that is robust to uncertain or inexact choices of failure threshold.

Remark 5 (Convexity in bPoF-based optimization) *It can be shown that bPoF² can be written in the form of a convex optimization problem, similar to \bar{Q}_α , as [63, 20]*

$$\bar{p}_t(g(\mathbf{d}, Z)) = \min_{\lambda < t} \frac{\mathbb{E}[g(\mathbf{d}, Z) - \lambda]^+]}{t - \lambda}, \quad (20)$$

where \mathbf{d} is the given design, $[c]^+ = \max\{0, c\}$, and λ is an auxiliary variable. Note that the optimal λ^* from Equation (20) is the threshold from Equation (16) that provides $\mathbb{E}[g(\mathbf{d}, Z) | g(\mathbf{d}, Z) \geq \lambda^*] = t$.

Using Equation (20), formulation (19) can be reduced to an optimization problem involving only expectations as given by

$$\begin{aligned} \min_{\lambda < t, \mathbf{d} \in \mathcal{D}} \quad & \mathbb{E}[f(\mathbf{d}, Z)] \\ \text{subject to} \quad & \frac{\mathbb{E}[g(\mathbf{d}, Z) - \lambda]^+]}{t - \lambda} \leq 1 - \alpha_T. \end{aligned} \quad (21)$$

Note that (21) can be reformulated by a simple rearrangement of the constraint to become equivalent to the \bar{Q}_α constrained problem given by (11). Thus, it is a convex problem when $g(\mathbf{d}, Z)$ and $f(\mathbf{d}, Z)$ are convex. The same linearization trick can also be performed as in (12).

²assuming $\bar{Q}_0[g(\mathbf{d}, Z)] < t < \bar{Q}_1[g(\mathbf{d}, Z)]$ and $g(\mathbf{d}, Z)$ is integrable

3.3.3 Optimization problem: bPoF as objective

A bPoF objective provides us with an optimization problem focused on optimal reliability subject to satisfaction of other design metrics. While the use of bPoF as a constraint is equivalent to a \overline{Q}_α constraint, the same can not be said about the case in which bPoF and \overline{Q}_α are used as an objective function. Consider the PoF minimization problem

$$\begin{aligned} \min_{\mathbf{d} \in \mathcal{D}} \quad & p_t(g(\mathbf{d}, Z)) \\ \text{subject to} \quad & \mathbb{E}[f(\mathbf{d}, Z)] \leq C_T. \end{aligned} \tag{22}$$

As mentioned in Remark 2, PoF is often nonconvex and discontinuous, making gradient calculations ill-posed or unstable. However, (22) is a desirable formulation if reliability is paramount. The formulation in (22) defines the situation where given our design performance specifications, characterized by $\mathbb{E}[f(\mathbf{d}, Z)] \leq C_T$, we desire the most reliable design achievable.

We can consider an alternative to the problem in (22) using a bPoF objective function as

$$\begin{aligned} \min_{\mathbf{d} \in \mathcal{D}} \quad & \overline{p}_t(g(\mathbf{d}, Z)) \\ \text{subject to} \quad & \mathbb{E}[f(\mathbf{d}, Z)] \leq C_T. \end{aligned} \tag{23}$$

Using Equation (20), the optimization problem in (23) can be rewritten in terms of expectations as

$$\begin{aligned} \min_{\lambda < t, \mathbf{d} \in \mathcal{D}} \quad & \frac{\mathbb{E} \left[[g(\mathbf{d}, Z) - \lambda]^+ \right]}{t - \lambda}, \\ \text{subject to} \quad & \mathbb{E}[f(\mathbf{d}, Z)] \leq C_T. \end{aligned} \tag{24}$$

This allows one to minimize bPoF, a conservative upper bound of PoF, subject to constraints on costs or other performance measures, while maintaining convexity and other mathematical properties of the limit state function.

3.3.4 Discussion on bPoF-based optimization

There are several advantages of using the bPoF-based optimization problem described by Equation (19) as compared to the RBDO problem. First, the bPoF-based optimization problem leads to a data-informed conservative design as compared to RBDO. This data-informed conservativeness of the bPoF-based optimal design is more desirable and resilient because it takes into account the magnitude of failure (or the tail of the distribution) that guards against more serious catastrophic failures as highlighted later in Remark 6. Second, the bPoF-based optimization problem preserves convexity if the underlying limit state function is convex as noted in Remark 5 (as compared to PoF, which does not preserve convexity). This leads to well-behaved convex optimization problems even for the risk-based formulation and pushes us to pay more attention to devising limit state functions that are convex or nearly convex. Additionally, if π is independent of \mathbf{d} , the same linear reformulation trick used with \overline{Q}_α constraints (see (12)) can be used to transform the objective in (24) into a linear function with additional linear constraints and auxiliary variables, offering similar advantages as noted in Section 3.3.2.3.2.4. Third, under certain conditions, it is possible to calculate (quasi)-gradients for bPoF [64]. Note that non-smoothness can also be avoided by using smoothed versions given by [65, 23].

4 Structural Design: Short Column Problem With Convex Reformulation

In this section, we use the short column structural design problem that has been widely used in the RBDO community [28, 43] to compare some of the properties of PoF- and bPoF-based optimization formulations. We show a convex reformulation of the objective and limit state functions leading to strongly certifiable risk-based designs for bPoF-based optimization.

4.1 Short column problem description

The problem consists of designing a short column with rectangular cross-section of dimensions w and h , subjected to uncertain loads (axial force F and bending moment M). The yield stress of the material, Y , is also considered to be uncertain. The random variables are $Z = [F, M, Y]^\top$ with a joint distribution π . Table 1 describes the random variables used in the short column design. The correlation coefficient between F and M is 0.5. The design variables, $\mathbf{d} = [w, h]^\top$, are the width and depth of the cross-section as shown in Table 2. The objective function is the cross-sectional area given by wh . Along with a failure threshold $t = 1$, the limit state function is defined as

$$g(\mathbf{d}, \mathbf{z}) = \frac{4M}{wh^2Y} + \frac{F^2}{w^2h^2Y^2}. \quad (25)$$

Table 1: Random variables used in the short column application.

Random variable	Units	Distribution	Mean	Standard deviation
F	kN	Normal	500	100
M	kNm	Normal	2000	400
Y	MPa	Log-normal	5	0.5

Table 2: Design variables used in the short column application.

Design variable	Lower bound (m)	Upper bound (m)
w	5	15
h	15	25

4.2 Optimization problem formulations

This section provides the optimization formulations based on PoF and bPoF for the short column structural design. We show a convex reformulation of the bPoF-based optimization short column problem to emphasize the specific advantage of bPoF risk measure making it strongly certifiable. For each case, we solve multiple optimization problems each with a different fixed value of $1 - \alpha_T$, where α_T is the desired reliability.

4.2.1 Short column RBDO

The RBDO problem is given by

$$\begin{aligned} \min_{w, h} \quad & wh \\ \text{subject to} \quad & p_t(g(\mathbf{d}, Z)) \leq 1 - \alpha_T, \\ & \ell_w \leq w \leq u_w, \\ & \ell_h \leq h \leq u_h, \end{aligned} \quad (26)$$

where $(\ell_w, \ell_h, u_w, u_h)$ denote the lower and upper bounds on w and h as defined in Table 2.

4.2.2 Short column bPoF-constrained CRiBDO with convex reformulation

The bPoF-based optimization problem for the short column design is

$$\begin{aligned} & \min_{\lambda < t, w, h} \quad wh \\ & \text{subject to} \quad \frac{\mathbb{E} \left[\left[\frac{4M}{wh^2Y} + \frac{F^2}{w^2h^2Y^2} - \lambda \right]^+ \right]}{t - \lambda} \leq 1 - \alpha_T, \\ & \quad \ell_w \leq w \leq u_w, \\ & \quad \ell_h \leq h \leq u_h, \end{aligned} \tag{27}$$

where we used (21). We first show that the optimization problem with a bPoF constraint can be formulated as a convex optimization problem in this case. This will permit us to take advantage of convex optimization solvers and offer guarantees for the optimization problem. To reformulate this as a more manageable convex optimization problem³, we rearrange the first constraint in (27) to achieve an equivalent form⁴,

$$\lambda + \frac{1}{(1 - \alpha_T)} \mathbb{E} \left[\left[\frac{4M}{wh^2Y} + \frac{F^2}{w^2h^2Y^2} - \lambda \right]^+ \right] \leq t. \tag{28}$$

Next, we note that both w and h are nonnegative and thus we can make the change of variable $w = e^{x_1}$, $h = e^{x_2}$ with $x_1, x_2 \in \mathbb{R}$. Then the limit state function becomes $g(x_1, x_2, \mathbf{z}) = \frac{4M}{Y} e^{-x_1 - 2x_2} + \frac{F^2}{Y^2} e^{-2x_1 - 2x_2}$ and objective function becomes $wh = e^{x_1 + x_2}$. These are both convex functions in the new design variables (x_1, x_2) . Thus, using the change of variable and Equation (28) we can reformulate (27) as

$$\begin{aligned} & \min_{\lambda < t, x_1, x_2} \quad e^{x_1 + x_2} \\ & \text{subject to} \quad \lambda + \frac{1}{(1 - \alpha_T)} \mathbb{E} \left[\left[\frac{4M}{Y} e^{-x_1 - 2x_2} + \frac{F^2}{Y^2} e^{-2x_1 - 2x_2} - \lambda \right]^+ \right] \leq t, \\ & \quad \ln \ell_w \leq x_1 \leq \ln u_w, \\ & \quad \ln \ell_h \leq x_2 \leq \ln u_h. \end{aligned} \tag{29}$$

Furthermore, since the distribution of the random variables are independent of the design variables, we can empirically estimate the expectation in the constraint for any design by using a fixed set of m samples $\{\mathbf{z}_1, \dots, \mathbf{z}_m\}$ that are sampled *a priori* from the distribution π . This gives us a convex sample-average-approximation optimization problem for the given set of m samples as

$$\begin{aligned} & \min_{\lambda < t, x_1, x_2} \quad e^{x_1 + x_2} \\ & \text{subject to} \quad \lambda + \frac{1}{m(1 - \alpha_T)} \sum_{i=1}^m \left[\frac{4M_i}{Y_i} e^{-x_1 - 2x_2} + \frac{F_i^2}{Y_i^2} e^{-2x_1 - 2x_2} - \lambda \right]^+ \leq t, \\ & \quad \ln \ell_w \leq x_1 \leq \ln u_w \\ & \quad \ln \ell_h \leq x_2 \leq \ln u_h. \end{aligned} \tag{30}$$

Note that the constraint is convex since $[\cdot]^+$ is a convex function and preserves the convexity of the limit state function. Thus, the formulation in (30) leads to a strongly certifiable formulation since it satisfies both conditions of certifiability.

4.3 Experimental comparison between RBDO and bPoF-based CRiBDO

We now compare the behavior of the RBDO formulation and the bPoF-constrained CRiBDO formulation. We solve the RBDO problem with various values of α_T using the gradient-free COBYLA optimizer [66].

³We were required to find a formulation that was recognized as convex by modeling language CVXPY with convex solver MOSEK

⁴Note that this is in the same form as a \bar{Q}_{α_T} constraint from (11).

We estimate PoF in the RBDO problem by iteratively adding samples until the MC error reaches below 1%. We solve the bPoF-constrained problem (Equation (30)) with various values of α_T , utilizing the convex optimization solver MOSEK with CVXPY [67] as our modeling interface to the solver.

We illustrate that, as hypothesized, PoF and bPoF are indeed natural counterparts with both measuring notions of failure risk. Each point in Figure 6(a) corresponds to an optimal design achieved by solving RBDO (“Using PoF”) or bPoF-constrained optimization for some α_T (“Using bPoF”). The x-axis provides the estimated value of PoF of the design, estimated using a separate evaluation sample of size 5×10^6 . The y-axis provides the cross-section area of the optimal design. Figure 6(b) is similar, except with the value of bPoF (estimated using 5×10^6 samples) as the x-axis. We see that RBDO and bPoF-based CRiBDO achieve nearly identical frontiers of cross-section area for the optimal designs with similar PoF or bPoF. The optimal designs similar to the ones obtained using RBDO can be obtained by setting higher values for $1 - \alpha_T$ in the bPoF constraint as compared to PoF constraint, and vice versa. This highlights the fact that bPoF is indeed a natural replacement for PoF, achieving similar goals and controlling failure probability. All the optimal designs for the short column problem have $h \approx 25$ as shown in Figure 6(d). However, an advantage of the bPoF-based CRiBDO formulation is that while both achieve similar design frontiers, we have a guarantee that the design given by the bPoF formulation is *globally optimal* for the given sample-average-approximation problem in Equation (30) since it is a convex problem. We have no such guarantee for the designs given by RBDO. Thus, even if similar designs are achieved by RBDO and bPoF-constrained optimization, we have additional guarantees about the quality of the bPoF-based design due to the underlying convexity.

Another advantage is the conservative nature of bPoF as compared to PoF. When formulated with identical levels of α_T in the constraint, the bPoF-based optimization achieves a more conservative design than RBDO. The conservative nature of bPoF for the same desired reliability α_T can be seen from Figure 7 and is directly reflected in more conservative optimal designs with larger cross-sectional areas as seen in Figure 6(c). This type of probabilistically derived conservativeness can be seen as desirable as highlighted in Remark 6 below.

Figure 7 compares different levels of desired $1 - \alpha_T$ versus the estimated PoF or bPoF for the optimal designs obtained through PoF- and bPoF-based optimization. For these plots, we use 5×10^6 samples⁵ to get accurate estimates of the PoF or bPoF at the optimum. Figure 7(a) shows the desired $1 - \alpha_T$ and the PoF/bPoF for the optimum design obtained using the RBDO problem. We can see that since the MC error for PoF estimate in the RBDO problem was always ensured to be below 1%, the desired PoF and the PoF at the optimum overlap. The figure also emphasizes the conservative property of bPoF for the same desired $1 - \alpha_T$.

A key observation is illustrated by Figure 7(b), which compares the results for optimal designs obtained using bPoF-based CRiBDO for different *a priori* sample sizes, i.e. the value of m in Equation (30). We make this comparison to analyze the effect of fixing the sample set for all optimization iterations before starting the optimization, which is required to obtain the convex optimization formulation shown in Equation (30). We can see that for lower sample sizes of 10^3 and 10^4 , the bPoF at the optimum and the desired bPoF do not overlap reflecting inaccurate MC estimates of bPoF. However, it should be noted that the bPoF formulation is still effective in controlling the PoF, even when sample size is small. In other words, even when small number of samples are used within the optimization, the nature of bPoF yields an optimal design with desirable conservativeness and thus, an acceptably low PoF. Additionally, even when formulated with a small sample size, the bPoF-based convex optimization problem is still considerably stable leading to good optimal designs. One of the primary drawbacks of RBDO is the potential fragility of the optimization, particularly when sample sizes are small, where the estimates of PoF and/or gradients (if a gradient-based solver is used) are unstable and produce poor or inconsistent optimization results. The bPoF formulation does not seem to suffer in the same way for the short column design as illustrated here.

Remark 6 (Desirable data-informed conservativeness) *We take a closer look at the conservativeness induced by the bPoF-based CRiBDO for the same desired reliability level α_T as compared to PoF-based optimization (as shown in Figures 6(c) and 7) and why this type of conservativeness would be desirable and resilient. There are other ways of introducing conservativeness, such as safety factors, basis values,*

⁵These samples are not used in the optimization, but only to estimate at the optimal design after the optimization is completed.

⁷The *a priori* sample sizes m used for the bPoF CRiBDO are indicated in the legend using bPoF- m .

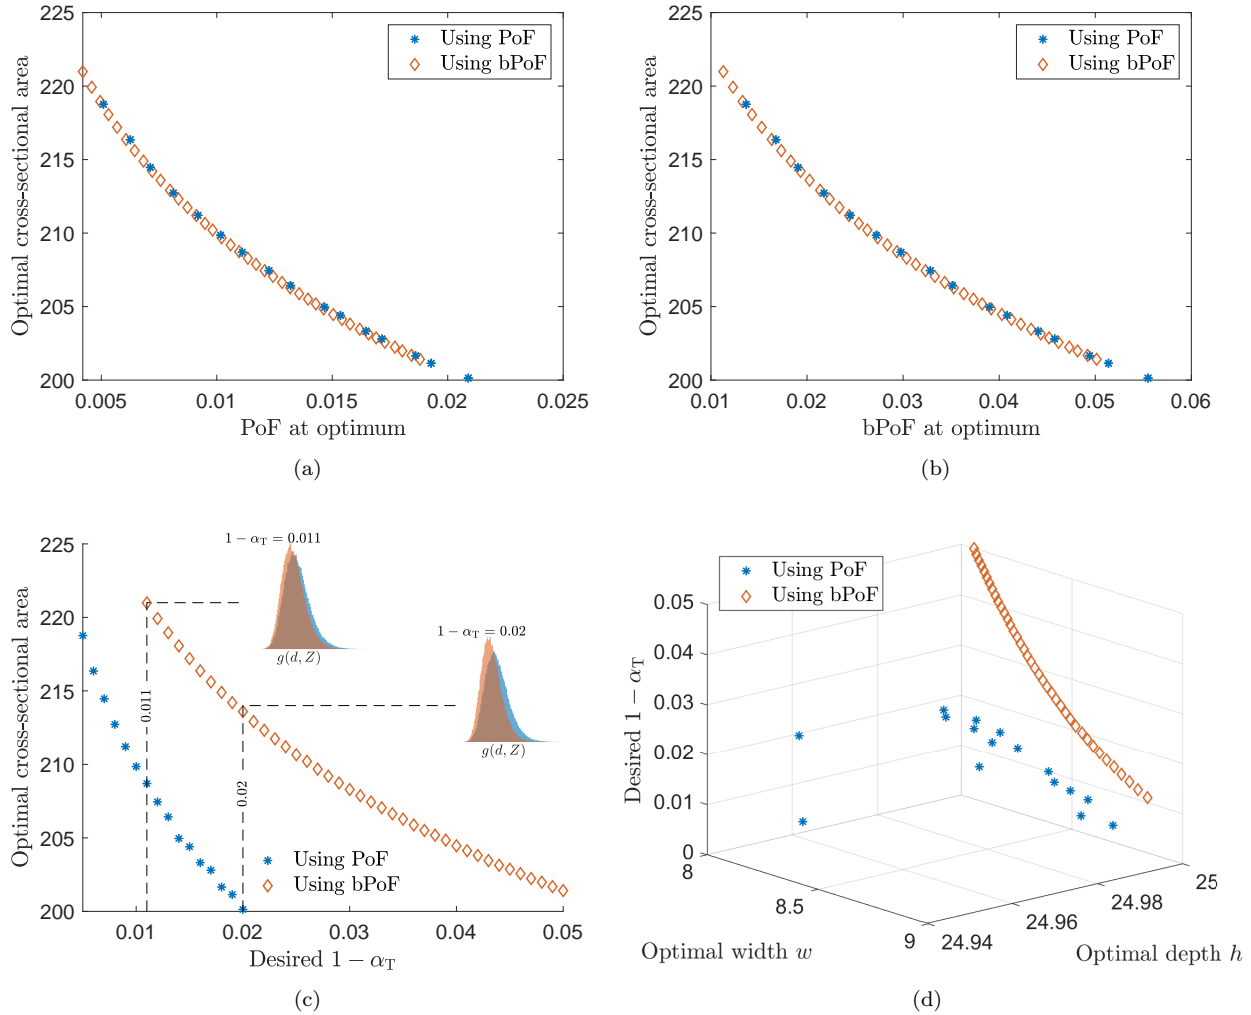


Figure 6: Optimal cross-sectional area and designs obtained for different levels of target $1 - \alpha_T$.

and stricter reliability levels. Typically, using safety factors leads to overly conservative designs. When inappropriate safety factor values are used, the deterministic optimization setup could also potentially lead to unreliable designs. This is because converting to deterministic optimization using just safety factors (or basis values) to account for the uncertainty in the system does not take into account the distribution of the limit state function and lacks sufficient information to make good decisions. These well-known issues with safety-factor- and basis-values-based deterministic optimization formulations have progressively led us to consider risk-based optimization under uncertainty. Another way to introduce conservativeness in risk-based optimization is by using lower values of $1 - \alpha_T$, which leads to stricter reliability constraints. However, this will just lead to overly reliable designs without any information about the distribution of the limit state function.

The bPof-based CRiBDO can be seen as a better way to induce conservativeness because it encodes more information about the underlying limit state function through the data on the magnitude of failure (as seen from Equations (16) and (18)). In Section 5, we highlight a similar observation on conservativeness for Q_α -based CRiBDO through the thermal design problem. These CRiBDO formulations lead to a probabilistic data-informed way of achieving a conservative design, which can be seen as more desirable in practice.

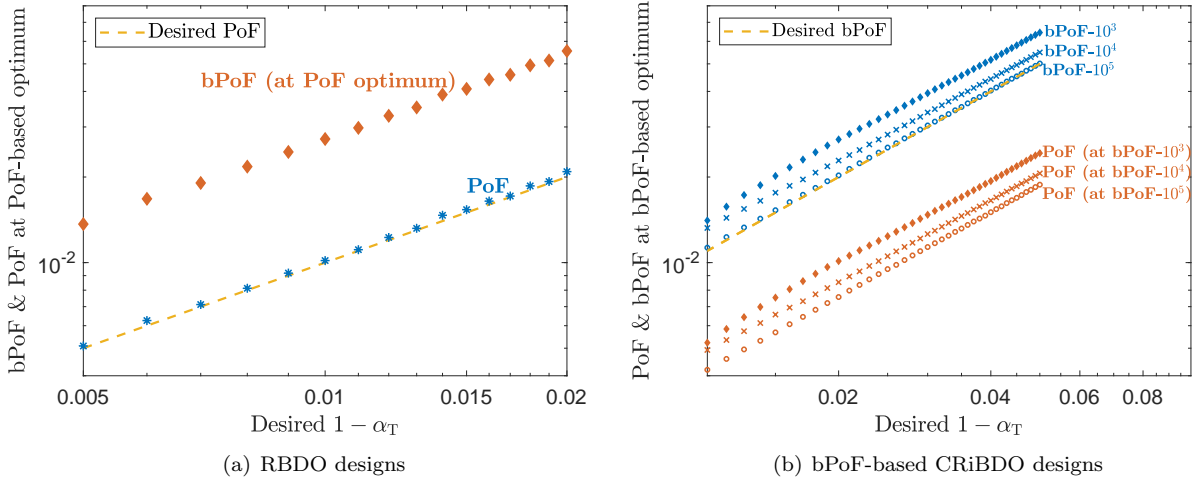


Figure 7: Comparing (a) bPoF estimates at the optimal designs obtained through RBDO, and (b) PoF estimates at the optimal designs obtained through bPoF-based CRiBDO using different *a priori* sample sizes⁷ for different desired $1 - \alpha_T$ in log-log scale.

5 Thermal Design: Cooling Fin Problem

In this section, we compare the properties of PoF- and \bar{Q}_α -based optimization formulations for the thermal design of a cooling fin problem. A cooling fin is used in many engineering applications to dissipate heat. This application also highlights the usefulness of CRiBDO even without the second certifiability condition of convexity. This problem considers a discretized partial differential equation, where closed-form solutions (such as in the previous short column) are not available, and only computer simulations are available.

5.1 Cooling fin model description

We consider a cooling fin with fixed geometry as shown in Figure 8, consisting of a vertical post with horizontal fins attached. We briefly review the problem here and refer to [68] for more details. The fin array consists of four horizontal sub-fins with width 2.5 and thickness 0.25, as well as a fin post with unit width and height of four. The thermal design is parametrized by the fin conductivities $k_i, i = 1, \dots, 4$ and the post conductivity k_0 , as well as the Biot number Bi , which is a non-dimensionalized heat transfer coefficient for thermal transfer from the fins to the surrounding air. The design variables, $\mathbf{d} = [k_1, k_2, k_3, k_4]$ are the thermal conductivities of the four fins as shown in Table 3. The post conductivity is $k_0 = 5$ and the Biot number is $Bi = 0.5$. We introduce manufacturing and operational uncertainties in all the parameters through the random variable $Z = [\xi_0, \xi_1, \xi_2, \xi_3, \xi_4, \xi_{Bi}]^\top$ with a joint distribution π given in Table 4. The random variables ξ_i model the additive uncertainty for the respective thermal conductivities $k_i, i = 0, \dots, 4$ and ξ_{Bi} models the additive uncertainty for the Biot number Bi . The system is governed by Poisson's equation in two spatial dimensions denoted by \mathbf{x} whose solution is the temperature field $y(\mathbf{x}, \mathbf{d}, Z)$. The PDE is semi-discretized with the finite element method and yields a system with 4,760 degrees of freedom.

The fin conducts heat away from the root Γ_{root} , so the lower the root temperature, the more effective the cooling fin. Thus, our objective function depends on the measure of the average temperature at the root, i.e.,

$$\mathcal{Y}(\mathbf{d}, Z) = \int_{\Gamma_{\text{root}}} y(\mathbf{x}, \mathbf{d}, Z) dx. \quad (31)$$

We also include a quantity proportional to the cost of the material based on the area and material thermal conductivity in the objective function as shown in Section 5.5.2. The limit state function for the cooling fin




Figure 8: Fin geometry and model parameters.

problem is based on the maximum temperature and is defined as

$$g(\mathbf{d}, Z) = \max_x y(\mathbf{x}, \mathbf{d}, Z). \quad (32)$$

We choose $t = 0.35$ as the constraint on the limit state function to define the maximum allowable temperature of the system.

Table 3: Design variables used in the cooling fin application.

Design variable	Lower bound	Upper bound
$k_i, i = 1, \dots, 4$	1	10

Table 4: Random variables used in the cooling fin application.

Random variable	Distribution	Mean μ	Standard deviation σ
$\xi_i, i = 0, \dots, 4$	truncated normal ($[\mu - 4\sigma, \mu + 2\sigma]$)	0	0.1
ξ_{Bi}		0	0.02

5.2 Optimization problem formulations

This section provides the optimization formulations based on PoF and \bar{Q}_α for the cooling fin thermal design.

5.2.1 Cooling fin RBDO

The RBDO problem is given by

$$\begin{aligned} \min_{\mathbf{d} \in \mathcal{D}} \quad & \mathcal{Y}(\mathbf{d}, \mu_Z) + \frac{\sum_{i=0}^4 A_i k_i}{5A_0 + \sum_{i=1}^4 10A_i} \\ \text{subject to} \quad & p_t(g(\mathbf{d}, Z)) \leq 1 - \alpha_T, \end{aligned} \quad (33)$$

where A_i denotes the area for the material with thermal conductivity of $k_i, i = 0, \dots, 4$ and $A_i k_i$ represents a quantity proportional to the cost of the material. Here, the fin post area is $A_0 = 4$ and sub-fin areas are $A_i = 1.25, i = 1, \dots, 4$. The cost part is normalized by the maximum proportionate cost. We use two different values of $1 - \alpha_T \in \{0.001, 0.05\}$.

5.2.2 Cooling fin \bar{Q}_α -constrained CRiBDO

The \bar{Q}_α -constrained CRiBDO formulation for the cooling fin design is

$$\begin{aligned} \min_{\mathbf{d} \in \mathcal{D}} \quad & \mathcal{Y}(\mathbf{d}, \mu_Z) + \frac{\sum_{i=0}^4 A_i k_i}{5A_0 + \sum_{i=1}^4 10A_i} \\ \text{subject to} \quad & \bar{Q}_{\alpha_T}[g(\mathbf{d}, Z)] \leq t, \end{aligned} \quad (34)$$

where we find the optimal designs for two different values of $1 - \alpha_T \in \{0.001, 0.05\}$. In this case, the underlying limit state function is not known to be convex making the CRiBDO formulation (34) certifiable in one condition, which is the data-informed conservativeness.

5.3 Experimental comparison between RBDO and \bar{Q}_α -constrained CRiBDO

We compare the optimal results obtained through RBDO and \bar{Q}_α -constrained CRiBDO formulations under the same α_T values. We solve the RBDO and CRiBDO problems using the gradient-free COBYLA optimizer. We estimate the PoF in each RBDO iteration by iteratively adding samples until the MC error reaches below 1% with the maximum number of samples capped at 10^5 . We estimate the \bar{Q}_α in each CRiBDO iteration by using 10^5 MC samples.

Table 5 shows the optimal designs obtained from the different optimization formulations. We start the optimization with an initial design that is feasible for all the optimization formulations. The optimal designs obtained using \bar{Q}_α -constrained CRiBDO for a given α_T are more conservative than the RBDO designs. This highlights one of the major advantages of using \bar{Q}_α -constrained CRiBDO that certifies designs through the data-informed conservativeness. The conservative nature of the \bar{Q}_α -constrained CRiBDO can be clearly seen by comparing the limit state function distributions at the optimal designs as shown in Figure 9. As discussed in Remark 6, this conservativeness is desirable and required to prevent catastrophic failures. Figure 10 compares the specified hard thresholds with the \bar{Q}_α for the different optimal designs to further highlight the fact that \bar{Q}_α considering the magnitude of failure and not using hard thresholding leads to appropriately conservative designs. The data-informed nature of the conservativeness is a significant advantage since the magnitude of conservativeness induced automatically changes according to the data from the underlying limit state function distribution for a particular design, i.e., \bar{Q}_α is more conservative only when it is required as dictated by the underlying distribution. We explicitly show the data-informed nature of conservativeness in the next section.

Table 5: Optimal designs obtained from RBDO and \bar{Q}_α -constrained CRiBDO.

Design variable/ Output statistic	Initial design	RBDO		\bar{Q}_α -constrained CRiBDO	
		$1 - \alpha_T = 0.001$	$1 - \alpha_T = 0.05$	$1 - \alpha_T = 0.001$	$1 - \alpha_T = 0.05$
k_1	5.5	4.154	3.3271	4.5175	3.6468
k_2	5.5	4.1132	3.2015	3.965	3.1974
k_3	5.5	1	1.0053	1	1
k_4	5.5	1	1	1	1
Objective function	1.0141	0.7649	0.7094	0.781	0.725
PoF	0	0.001	0.0486	4.3×10^{-4}	0.0195
\bar{Q}_{α_T}	0.3376 ($1 - \alpha_T = 0.001$) 0.3301 ($1 - \alpha_T = 0.05$)	0.3524	0.3526	0.35	0.35

5.4 Effect of input uncertainties on data-informed conservativeness

In this section, we demonstrate the data-informed nature of the conservativeness induced by \bar{Q}_α as described in Remark 6. The magnitude of conservativeness is naturally adjusted for different limit state function distributions. Since, \bar{Q}_α and Q_α are natural counterparts, we quantify the magnitude of conservativeness by the percentage difference $(\bar{Q}_\alpha - Q_\alpha)/Q_\alpha\%$. We fix the design and the α value for the comparison in this

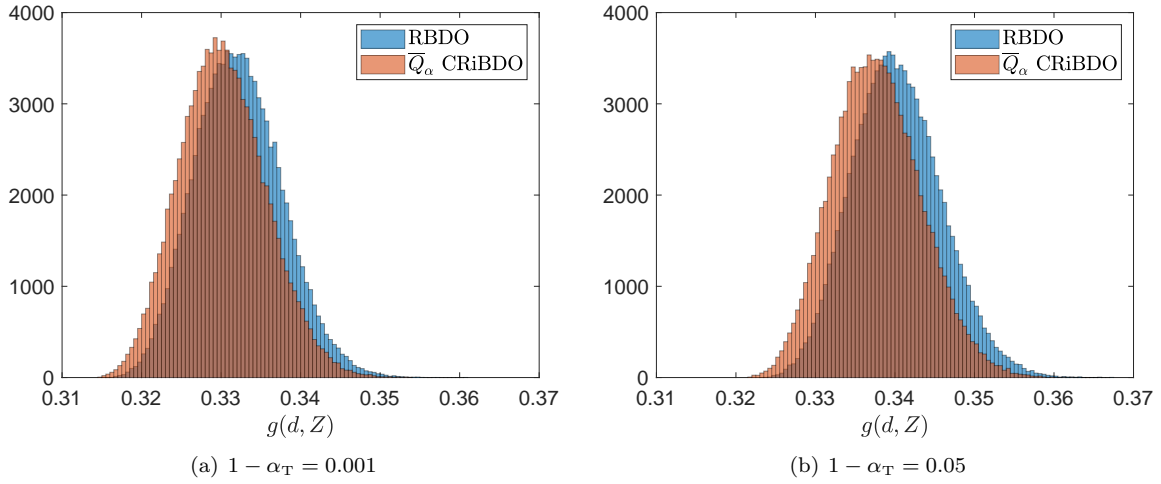


Figure 9: Histograms comparing limit state function distributions for optimal designs obtained through different optimization formulations.




Figure 10: Comparing specified thresholds and \bar{Q}_α for optimal designs obtained through different optimization formulations.

section. We use the optimal design obtained by RBDO with $1 - \alpha_T = 0.05$ given in Table 5 as the fixed design

$\mathbf{d} = [3.3271, 3.2015, 1.0053, 1]$. Different limit state function distributions are generated by changing the input uncertainties through modifying the truncation range for the distribution given in Table 4. Figure 11 shows histograms of the different limit state function distributions for the fixed design. Three different truncation ranges for the input random variables are indicated in the Figure 11 legend. The distribution of the limit state function is obtained by first sampling the input random variables Z from their given distribution and then running the simulation to obtain realizations of $g(\mathbf{d}, Z)$. All the results are generated using 10^5 MC samples.

Figure 12 shows the different levels of conservativeness of \bar{Q}_α when compared to Q_α for different limit state function distributions. We can see that \bar{Q}_α is always conservative when compared to Q_α . Furthermore, it can be seen that the magnitude of conservativeness depends on the distribution, which exemplifies the data-informed nature of the induced conservativeness, i.e., \bar{Q}_α is as conservative as required by the underlying distribution. Specifically, the magnitude of conservativeness for superquantile is higher for the fatter-tailed distribution as seen in Figure 12 (a). For fat-tailed distributions, i.e., distributions with significant tail risk, superquantile provides additional conservativeness by nature of being a tail-integral.




Figure 11: Histograms for different limit state function distributions generated through modifying input uncertainty truncation range for a fixed design $\mathbf{d} = [3.3271, 3.2015, 1.0053, 1]$.

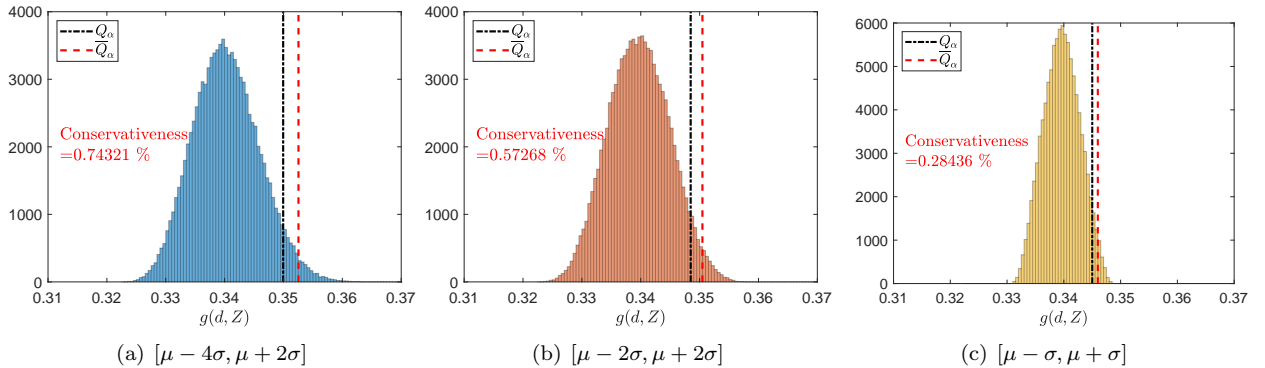


Figure 12: Conservativeness induced by \bar{Q}_α compared to Q_α for different limit state function distributions generated through modifying input uncertainty truncation range for a fixed design $\mathbf{d} = [3.3271, 3.2015, 1.0053, 1]$ and $1 - \alpha = 0.05$.

6 Concluding Remarks

In this work, we propose two certifiability conditions that lead to certifiable risk-based design optimization (CRiBDO): (a) data-informed conservativeness: the resulting designs should be certifiably risk-averse against near-failure and catastrophic failure events, and (b) optimization convergence and efficiency: the resulting designs should be certifiably optimal in comparison with all alternate designs at reduced computational cost for the optimization. The risk measures satisfying either of the certifiability conditions are classified under CRiBDO while satisfying both conditions makes the resulting optimal designs strongly certifiable. We compare and contrast the existing RBDO formulation based on probability of failure (PoF) with risk-based optimization formulations using the buffered probability of failure (bPoF) and the superquantile (a.k.a. conditional value-at-risk) risk measures. We show that RBDO does not satisfy either of the certifiability conditions while superquantiles and bPoF lead to CRiBDO formulations. An additional advantage of bPoF is the intuitive relation to PoF, which allows widely used PoF-based optimization formulations to be easily transitioned to bPoF-based optimization formulations.

Both bPoF and superquantile risk measures introduce data-informed conservativeness by encoding extra information about the limit state function distribution in the form of the magnitude of failure. This highlights a way to get desirable conservativeness in engineering risk-based optimization by switching to a different optimization formulation and getting certifiably risk-averse designs. Superquantiles and bPoF provide alternate measures of risk that avoid hard-threshold characterizations of failure events used in PoF, thus, relaxing the guesswork associated with picking limit state function thresholds. We highlight the data-informed nature of the conservativeness through numerical experiments on bPoF-based CRiBDO for short column structural design problem and superquantile-based CRiBDO for cooling fin thermal design problem.

When the underlying limit state and objective functions are convex w.r.t. the design variables, using bPoF and superquantiles can lead to strongly certifiable designs because these risk measures can preserve convexity and lead to certifiably optimal designs. We show this property for bPoF-based CRiBDO through convex reformulation of the limit state and the objective function for the short column design problem. The bPoF-based CRiBDO results in a convex optimization problem that can be solved using convex optimizers and provide global convergence guarantees. Although, such convex reformulations are not always possible, using convex approximations for the limit state functions is one way of addressing non-convex problems.

Estimation of bPoF and superquantiles can be expensive, and recent results for Monte Carlo variance reduction techniques along with approximation techniques have been shown to be effective ways of reducing the computational effort for estimating superquantiles and bPoF [12, 13, 25, 16, 17, 19, 18, 14, 15]. More such research efforts will lead to flexibility in switching to appropriate and advantageous CRiBDO formulations for designing safe engineering systems.

Acknowledgement

This work has been supported in part by the Air Force Office of Scientific Research (AFOSR) MURI on managing multiple information sources of multi-physics systems award numbers FA9550-15-1-0038 and FA9550-18-1-0023, and Air Force Center of Excellence on Multi-Fidelity Modeling of Rocket Combustor Dynamics award FA9550-17-1-0195. The fourth author acknowledges the support from Office of Naval Research under MIPR N0001420WX00519.

References

- [1] Rockafellar, R. T. and Uryasev, S., “Conditional value-at-risk for general loss distributions,” *Journal of Banking & Finance*, Vol. 26, No. 7, 2002, pp. 1443–1471.
- [2] Krokmal, P., Palmquist, J., and Uryasev, S., “Portfolio optimization with conditional value-at-risk objective and constraints,” *Journal of Risk*, Vol. 4, No. 2, 2002, pp. 11–27.
- [3] Mansini, R., Ogryczak, W., and Speranza, M. G., “Conditional value at risk and related linear programming models for portfolio optimization,” *Annals of Operations Research*, Vol. 152, 2007, pp. 227–256.

- [4] Rockafellar, R. T. and Royset, J. O., “Superquantiles and their applications to risk, random variables, and regression,” *Theory Driven by Influential Applications*, Informs, 2013, pp. 151–167.
- [5] Rockafellar, R. T. and Uryasev, S., “Optimization of conditional value-at-risk,” *Journal of Risk*, Vol. 2, 2000, pp. 21–42.
- [6] Rockafellar, R. T. and Royset, J. O., “On buffered failure probability in design and optimization of structures,” *Reliability Engineering & System Safety*, Vol. 95, No. 5, 2010, pp. 499–510.
- [7] Roland, H. E. and Moriarty, B., *System safety engineering and management*, John Wiley & Sons, 1990.
- [8] Möller, N. and Hansson, S. O., “Principles of engineering safety: Risk and uncertainty reduction,” *Reliability Engineering & System Safety*, Vol. 93, No. 6, 2008, pp. 798–805.
- [9] Suzuki, Y. and Haftka, R. T., “Analytical benchmark example for risk allocation in structural optimization,” *Structural and Multidisciplinary Optimization*, Vol. 50, No. 1, 2014, pp. 1–7.
- [10] Rockafellar, R. T. and Royset, J. O., “Engineering decisions under risk averseness,” *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, Vol. 1, No. 2, 2015, pp. 04015003.
- [11] Zhang, W., Rahimian, H., and Bayraksan, G., “Decomposition algorithms for risk-averse multistage stochastic programs with application to water allocation under uncertainty,” *INFORMS Journal on Computing*, Vol. 28, No. 3, 2016, pp. 385–404.
- [12] Royset, J. O., Bonfiglio, L., Vernengo, G., and Brizzolara, S., “Risk-adaptive set-based design and applications to shaping a hydrofoil,” *Journal of Mechanical Design*, Vol. 139, No. 10, 2017, pp. 101403.
- [13] Bonfiglio, L. and Royset, J. O., “Multidisciplinary Risk-Adaptive Set-Based Design of Supercavitating Hydrofoils,” *AIAA Journal*, Vol. 57, No. 8, 2019, pp. 3360–3378.
- [14] Yang, H. and Gunzburger, M., “Algorithms and analyses for stochastic optimization for turbofan noise reduction using parallel reduced-order modeling,” *Comput. Methods Appl. Mech. Engrg.*, Vol. 319, 2017, pp. 217–239.
- [15] Chaudhuri, A., Peherstorfer, B., and Willcox, K., “Multifidelity Cross-Entropy Estimation of Conditional Value-at-Risk for Risk-Averse Design Optimization,” *AIAA Scitech 2020 Forum*, 2020, p. 2129.
- [16] Kouri, D. P. and Surowiec, T. M., “Risk-averse PDE-constrained optimization using the conditional value-at-risk,” *SIAM Journal on Optimization*, Vol. 26, No. 1, 2016, pp. 365–396.
- [17] Zou, Z., Kouri, D. P., and Aquino, W., “A locally adapted reduced basis method for solving risk-averse PDE-constrained optimization problems,” *2018 AIAA Non-Deterministic Approaches Conference, Kissimmee, Florida, AIAA SciTech Forum, (AIAA 2018-2174)*, 2018.
- [18] Heinkenschloss, M., Kramer, B., Takhtaganov, T., and Willcox, K., “Conditional-value-at-risk estimation via reduced-order models,” *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 6, No. 4, 2018, pp. 1395–1423.
- [19] Heinkenschloss, M., Kramer, B., and Takhtaganov, T., “Adaptive reduced-order model construction for conditional value-at-risk estimation,” *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 8, No. 2, 2020, pp. 668–692.
- [20] Mafusalov, A. and Uryasev, S., “Buffered probability of exceedance: mathematical properties and optimization,” *SIAM Journal on Optimization*, Vol. 28, No. 2, 2018, pp. 1077–1103.
- [21] Norton, M., Khokhlov, V., and Uryasev, S., “Calculating CVaR and bPOE for common probability distributions with application to portfolio optimization and density estimation,” *Annals of Operations Research*, 2019, pp. 1–35.

- [22] Rockafellar, R. T. and Uryasev, S., “Minimizing buffered probability of exceedance by progressive hedging,” *Mathematical Programming*, 2020, pp. 1–20.
- [23] Basova, H. G., Rockafellar, R. T., and Royset, J. O., “A computational study of the buffered failure probability in reliability-based design optimization,” *Proceedings of International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP), Zurich, Switzerland*, Citeseer, 2011.
- [24] Minguez, R., Castillo, E., and Lara, J., “Iterative scenario reduction technique to solve reliability based optimization problems using the buffered failure probability,” *Proceedings of ICOSSAR*, 2013.
- [25] Harajli, M. M., Rockafellar, R. T., and Royset, J. O., “Importance sampling in the evaluation and optimization of buffered failure probability,” *Proceedings of International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP), Vancouver, Canada, July 12-15*, 2015.
- [26] Royset, J. O., Günay, S., and Mosalam, K. M., “Risk-Adaptive Learning of Seismic Response using Multi-Fidelity Analysis,” *Proceedings of the International Conference on Applied Statistics and Probability in Civil Engineering (ICASP), Seoul, Korea*, 2019.
- [27] Yao, W., Chen, X., Luo, W., van Tooren, M., and Guo, J., “Review of uncertainty-based multidisciplinary design optimization methods for aerospace vehicles,” *Progress in Aerospace Sciences*, Vol. 47, No. 6, 2011, pp. 450–479.
- [28] Aoues, Y. and Chateaufneuf, A., “Benchmark study of numerical methods for reliability-based design optimization,” *Structural and Multidisciplinary Optimization*, Vol. 41, No. 2, 2010, pp. 277–294.
- [29] Sobieszczanski-Sobieski, J., Morris, A., and van Tooren, M., *Multidisciplinary design optimization supported by knowledge based engineering*, John Wiley & Sons, 2015.
- [30] Du, X. and Chen, W., “A most probable point-based method for efficient uncertainty analysis,” *Journal of Design and Manufacturing Automation*, Vol. 4, No. 1, 2001, pp. 47–66.
- [31] Rubinstein, R. Y. and Kroese, D. P., *Simulation and the Monte Carlo method*, Vol. 10, John Wiley & Sons, 2016.
- [32] Melchers, R., “Importance sampling in structural systems,” *Structural Safety*, Vol. 6, No. 1, 1989, pp. 3–10.
- [33] Owen, A. B., *Monte Carlo theory, methods and examples*, 2013.
- [34] Au, S.-K. and Beck, J. L., “A new adaptive importance sampling scheme for reliability calculations,” *Structural Safety*, Vol. 21, No. 2, 1999, pp. 135–158.
- [35] Dey, A. and Mahadevan, S., “Ductile structural system reliability analysis using adaptive importance sampling,” *Structural Safety*, Vol. 20, No. 2, 1998, pp. 137–154.
- [36] De Boer, P.-T., Kroese, D. P., Mannor, S., and Rubinstein, R. Y., “A tutorial on the cross-entropy method,” *Annals of Operations Research*, Vol. 134, No. 1, 2005, pp. 19–67.
- [37] Papaioannou, I., Betz, W., Zwirgmaier, K., and Straub, D., “MCMC algorithms for subset simulation,” *Probabilistic Engineering Mechanics*, Vol. 41, 2015, pp. 89–103.
- [38] Depina, I., Papaioannou, I., Straub, D., and Eiksund, G., “Coupling the cross-entropy with the line sampling method for risk-based design optimization,” *Structural and Multidisciplinary Optimization*, Vol. 55, No. 5, 2017, pp. 1589–1612.
- [39] Kurtz, N. and Song, J., “Cross-entropy-based adaptive importance sampling using Gaussian mixture,” *Structural Safety*, Vol. 42, 2013, pp. 35–44.
- [40] Li, J., Li, J., and Xiu, D., “An efficient surrogate-based method for computing rare failure probability,” *Journal of Computational Physics*, Vol. 230, No. 24, 2011, pp. 8683–8697.

- [41] Peherstorfer, B., Kramer, B., and Willcox, K., “Combining multiple surrogate models to accelerate failure probability estimation with expensive high-fidelity models,” *Journal of Computational Physics*, Vol. 341, 2017, pp. 61–75.
- [42] Peherstorfer, B., Kramer, B., and Willcox, K., “Multifidelity preconditioning of the cross-entropy method for rare event simulation and failure probability estimation,” *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 6, No. 2, 2018, pp. 737–761.
- [43] Bichon, B. J., Eldred, M. S., Mahadevan, S., and McFarland, J. M., “Efficient global surrogate modeling for reliability-based design optimization,” *Journal of Mechanical Design*, Vol. 135, No. 1, 2013, pp. 011009.
- [44] Chaudhuri, A., Marques, A. N., and Willcox, K. E., “mfEGRA: Multifidelity Efficient Global Reliability Analysis through Active Learning for Failure Boundary Location,” *Structural and Multidisciplinary Optimization*, 2020.
- [45] Moustapha, M. and Sudret, B., “Surrogate-assisted reliability-based design optimization: a survey and a unified modular framework,” *Structural and Multidisciplinary Optimization*, Vol. 60, No. 5, 2019, pp. 1–20.
- [46] Gano, S. E., Renaud, J. E., Agarwal, H., and Tovar, A., “Reliability-based design using variable-fidelity optimization,” *Structures and Infrastructure Engineering*, Vol. 2, No. 3-4, 2006, pp. 247–260.
- [47] Li, X., Qiu, H., Jiang, Z., Gao, L., and Shao, X., “A VF-SLP framework using least squares hybrid scaling for RBDO,” *Structural and Multidisciplinary Optimization*, Vol. 55, No. 5, 2017, pp. 1629–1640.
- [48] Marques, A., Lam, R., and Willcox, K., “Contour location via entropy reduction leveraging multiple information sources,” *Advances in Neural Information Processing Systems*, 2018, pp. 5217–5227.
- [49] Kramer, B., Marques, A., Peherstorfer, B., Villa, U., and Willcox, K., “Multifidelity probability estimation via fusion of estimators,” *Journal of Computational Physics*, Vol. 392, 2019, pp. 385–402.
- [50] Chaudhuri, A., Marques, A. N., Lam, R., and Willcox, K. E., “Reusing Information for Multifidelity Active Learning in Reliability-Based Design Optimization,” *AIAA Scitech 2019 Forum*, 2019, p. 1222.
- [51] Chaudhuri, A., Kramer, B., and Willcox, K. E., “Information Reuse for Importance Sampling in Reliability-Based Design Optimization,” *Reliability Engineering & System Safety*, 2020, pp. 106853.
- [52] Uryasev, S., “Derivatives of probability functions and some applications,” *Annals of Operations Research*, Vol. 56, No. 1, 1995, pp. 287–311.
- [53] Royset, J. and Polak, E., “Extensions of stochastic optimization results to problems with system failure probability functions,” *Journal of Optimization Theory and Applications*, Vol. 133, No. 1, 2007, pp. 1–18.
- [54] Tretiakov, G. L., “Star-shaped approximation approach for stochastic programming problems with probability function,” *Optimization*, Vol. 47, No. 3-4, 2000, pp. 303–317.
- [55] Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D., “Coherent measures of risk,” *Mathematical Finance*, Vol. 9, No. 3, 1999, pp. 203–228.
- [56] Rockafellar, R. T. and Uryasev, S., “The fundamental risk quadrangle in risk management, optimization and statistical estimation,” *Surveys in Operations Research and Management Science*, Vol. 18, No. 1-2, 2013, pp. 33–53.
- [57] Lan, G. and Zhou, Z., “Algorithms for stochastic optimization with expectation constraints,” *arXiv preprint arXiv:1604.03887*, 2016.
- [58] Tamar, A., Glassner, Y., and Mannor, S., “Optimizing the CVaR via sampling,” *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015, pp. 2993–2999.

- [59] Yamai, Y., Yoshiba, T., et al., “Comparative analyses of expected shortfall and value-at-risk: their estimation error, decomposition, and optimization,” *Monetary and Economic Studies*, Vol. 20, No. 1, 2002, pp. 87–121.
- [60] Uryasev, S., “New variable-metric algorithms for nondifferentiable optimization problems,” *Journal of Optimization Theory and Applications*, Vol. 71, No. 2, 1991, pp. 359–388.
- [61] Pertaia, G., Prokhorov, A., and Uryasev, S., “A new approach to credit ratings,” *Journal of Banking & Finance*, 2021, pp. 106097.
- [62] Ditlevsen, O. and Madsen, H. O., *Structural reliability methods*, Vol. 178, Wiley New York, 1996.
- [63] Norton, M. and Uryasev, S., “Maximization of AUC and buffered AUC in binary classification,” *Mathematical Programming*, Vol. 174, No. 1-2, 2018, pp. 575–612.
- [64] Zhang, T., Uryasev, S., and Guan, Y., “Derivatives and subderivatives of buffered probability of exceedance,” *Operations Research Letters*, Vol. 47, No. 2, 2019, pp. 130–132.
- [65] Kouri, D., “Higher-moment buffered probability,” *Optimization Letters*, Vol. 13, No. 6, 2019, pp. 1223–1237.
- [66] Powell, M. J., “A direct search optimization method that models the objective and constraint functions by linear interpolation,” *Advances in optimization and numerical analysis*, Springer, 1994, pp. 51–67.
- [67] Diamond, S. and Boyd, S., “CVXPY: A Python-embedded modeling language for convex optimization,” *Journal of Machine Learning Research*, Vol. 17, No. 83, 2016, pp. 1–5.
- [68] Prud’homme, C., Rovas, D. V., Veroy, K., Machiels, L., Maday, Y., Patera, A. T., and Turinici, G., “Reliable Real-Time Solution of Parametrized Partial Differential Equations: Reduced-Basis Output Bound Methods,” *Journal of Fluids Engineering*, Vol. 124, No. 1, 2002, pp. 70–80.