SIAM/ASA J. Uncertainty Quantification
Vol. 8, No. 2, pp. 668–692

Adaptive Reduced-Order Model Construction for Conditional Value-at-Risk Estimation*

Matthias Heinkenschloss†, Boris Kramer‡, and Timur Takhtaganov§

Abstract. This paper shows how to systematically and efficiently improve a reduced-order model (ROM) to obtain a better ROM-based estimate of the Conditional Value-at-Risk (CVaR) of a computationally expensive quantity of interest (QoI). Efficiency is gained by exploiting the structure of CVaR, which implies that a ROM used for CVaR estimation only needs to be accurate in a small region of the parameter space, called the ϵ-risk region. Hence, any full-order model (FOM) queries needed to improve the ROM can be restricted to this small region of the parameter space, thereby substantially reducing the computational cost of ROM construction. However, an example is presented which shows that simply constructing a new ROM that has a smaller error with the FOM is in general not sufficient to yield a better CVaR estimate. Instead a combination of previous ROMs is proposed that achieves a guaranteed improvement, as well as ϵ-risk regions that converge monotonically to the FOM risk region with decreasing ROM error. Error estimates for the ROM-based CVaR estimates are presented. The gains in efficiency obtained by improving a ROM only in the small ϵ-risk region over a traditional greedy procedure on the entire parameter space are illustrated numerically.

Key words. reduced-order models, risk measures, Conditional Value-at-Risk, estimation, sampling, uncertainty quantification

AMS subject classifications. 35R60, 62H12, 65G99, 65Y20

DOI. 10.1137/19M1257433

1. Introduction.

In this paper we develop an approach to systematically and efficiently improve a reduced-order model (ROM) to obtain a better ROM-based estimate of the Conditional Value-at-Risk (CVaR) of a computationally expensive quantity of interest (QoI). This paper builds on our recent work [3], where we analyzed uses of ROMs to substantially decrease the computational cost of sampling-based estimation of CVaR. Our previous paper used the approximation properties of a ROM, but the ROMs could have been computed separately. This paper integrates the ROM generation into the estimation process. Efficiency is gained by exploiting the structure of CVaR, which implies that a ROM used for CVaR estimation only needs to be accurate in a small region of the parameter space. Hence, any expensive

*Received by the editors April 22, 2019; accepted for publication (in revised form) February 19, 2020; published electronically May 5, 2020.
https://doi.org/10.1137/19M1257433
Funding: The work of the first author was supported in part by NSF grants DMS-1522798, DMS-1819144, and CCF-1816219. The work of the second author was supported in part by Air Force Center of Excellence award FA9550-17-1-0195.

†Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005-1892 (heinken@rice.edu).
‡Department of Mechanical and Aerospace Engineering, University of California San Diego, San Diego, CA 92093-0411 (bmkramer@ucsd.edu).
§Center for Computational Sciences and Engineering, Lawrence Berkeley National Laboratory, MS 50A-3111, Berkeley, CA 94720 (timur.takhtaganov@gmail.com).

full-order model (FOM) queries needed to improve a given ROM can be restricted to this small region of the parameter space, thereby substantially reducing the computational cost of ROM construction. CVaR and related risk measures have been used to quantify risk in a variety of applications ranging from portfolio optimization [18, 8, 11], to engineering design [16, 23, 21, 19], to partial differential equation (PDE)–constrained optimization [7, 25]. While in special cases the CVaR for some random variables with known distributions can be computed analytically [12], for most science and engineering applications the distribution of the QoI is not known analytically. Instead, this distribution depends on the distribution of the random variables entering the system and on the dependence of the system state (often the solution of a PDE) on these random variables. In this situation CVaR must be estimated by sampling the QoI, and each sample requires a computationally expensive solution of the FOM system of equations. The ROM approach proposed in this paper provides sequences of CVaR estimates with guaranteed error bounds, and decreasing errors with substantially reduced total number of expensive FOM evaluations.

Estimating the CVaR of a QoI requires sampling in the tail of the distribution of the QoI, and these samples lie in a small region, called the risk region, of the parameter space. Unfortunately, as indicated earlier, this risk region is not known analytically, but must be estimated from samples of the QoI. In [3] we have shown how to use a ROM for which an error estimate is available to construct a so-called ϵ-risk region that contains the true risk region of the original computationally expensive FOM QoI, and an estimate of the CVaR of the FOM QoI that only requires ROM evaluations. The error between the CVaR of the FOM QoI and this ROM-based CVaR estimate depends only on the ROM error in the ϵ-risk region. Therefore, we need to improve the ROM only in the ϵ-risk region. This is typically achieved by evaluating the FOM. Since these FOM queries are now restricted to the small ϵ-risk region and not the entire parameter space, our tailored process of improving the ROM is computationally substantially more efficient than traditional approaches. However, we present a simple example which shows that simply constructing a new ROM that has a smaller error with the FOM is in general not sufficient to yield a better CVaR estimate. Instead we propose a combination of the previously used ROM with the new ROM that achieves a guaranteed improvement in the CVaR estimate of the FOM QoI. We present error estimates for our ROM-based CVaR estimates, and we numerically demonstrate the gains in efficiency that can be obtained by improving a ROM only in the small ϵ-risk region over a traditional greedy procedure on the entire parameter space.

ROMs play a role in multifidelity methods for uncertainty quantification and optimization; see, e.g., the survey [13]. However, this survey focuses on the risk neutral expected value estimation. The use of ROMs for CVaR estimation and risk averse optimization is more recent and more limited. As we have already stated in [3], "Proper orthogonal decomposition based ROMs have recently been used in [21] to minimize CVaR_β for an aircraft noise problem modeled by the Helmholtz equation. However, they do not adaptively refine the ROMs, nor analyze the impact of ROMs on the CVaR_β estimation error. The design of an ultrahigh-speed hydrofoil by using CVaR_β optimization is considered by Royset et al. [19]. They propose building surrogates of the CVaR of their QoI and model these surrogates as random variables 'due to unknown error in the surrogate relative to the actual value' of the CVaR of their QoI. This randomness in the CVaR surrogate is then incorporated into the design process by

applying CVaR again, but with a different quantile level to the surrogate. Ultimately, they use a surrogate for the QoI that combines high-fidelity and low-fidelity QoI evaluations into a polynomial fit model. Our work does not require additional stochastic treatment of model error and focuses on the efficient and accurate sampling of CVaR using ROMs of the QoI that satisfy the original governing equations."

Zahr, Carlberg, and Kouri [22] extend the adaptive sparse-grid trust-region method of Kouri et al. [6] to include ROMs in optimization under uncertainty. The algorithm allows differentiable risk measures, such as a smoothed CVaR, but the numerical example in [22] considers risk neutral optimization using the expected value. While sparse grids can be very efficient for the integration of QoIs that are smooth in the random variables, numerical results [20, sec. 3.2.4] indicate that they may not be much more efficient than plain Monte Carlo sampling when applied to CVaR and other risk measures. Thus improving the efficiency of Monte Carlo sampling by integrating ROMs, CVaR structure, and Monte Carlo sampling as proposed in this paper seems beneficial for risk averse optimization.

Chen and Quarteroni [1] integrate ROMs into the evaluation of failure probabilities. An adaptive approach [1, Alg. 3] refines the ROM by a greedy method based on a criterion that tends to place snapshots near the boundary of the failure region in parameter space. However, no error estimates or improvement guarantees are given. The approach introduced in this paper could be integrated into [1, Alg. 3].

The paper by Zou, Kouri, and Aquino [26], which is an extension of [24], is closest to our paper in spirit. They compute estimates of general risk measures including CVaR based on a ROM and on error estimates that take into account the structure of the risk measure. However, their analysis is tied to their ROM approach, which uses a piecewise linear approximation over a Voronoi tessellation of the parameter space. To improve their ROM the Voronoi tessellation is refined as necessary. Their error estimates, which are tailored to the structure of the risk measure, tend to refine Voronoi tessellation primarily in subregions of the parameter space roughly corresponding to what we referred to earlier as the risk region. In contrast, our basic analysis is based on a generic ROM for which an error estimate is available, and we propose a combination of ROMs that leads to a guaranteed improvement of the ROM-based CVaR estimate. We then tailor our general framework to a class of widely used projection-based ROMs; see, e.g., [2], [4], or [15].

This paper is organized as follows. Section 2 introduces the problem formulation and reviews results from [3] that are needed for the integration of ROM construction. Section 3 presents our new adaptive ROM strategy for CVaR computation and gives a complete algorithm. Section 4 discusses practical aspects of the algorithm implementation as well as construction and error estimation for projection-based ROMs. In section 5 we present numerical results to support our theoretical findings and show the computational savings of our proposed adaptive ROM approach.

2. Problem formulation and background.

This section introduces the basic problem setting and notation and reviews some results on CVaR. Specifically, in subsection 2.1 we define the state equation and the QoI. Subsection 2.2 defines the CVaR and its corresponding risk region, and subsection 2.3 briefly reviews the sampling-based computation of CVaR.

2.1. The state equation and quantity of interest. Given a random variable $\boldsymbol{\xi}$ with values $\xi \in \Xi \subset \mathbb{R}^M$ and with density ρ, we are interested in the efficient approximation of risk measures of the random variable

$$\xi \mapsto s(y(\xi)), \tag{2.1}$$

where $s : \mathbb{R}^N \mapsto \mathbb{R}$ is a quantity of interest (QoI) which depends on $y : \Xi \mapsto \mathbb{R}^N$, which is implicitly defined as the solution of the state equation

$$F(y(\xi), \xi) = 0 \qquad \text{for almost all } \xi \in \Xi, \tag{2.2}$$

with $F : \mathbb{R}^N \times \Xi \mapsto \mathbb{R}^N$. For now we assume that (2.2) has a unique solution $y(\xi)$ for almost all $\xi \in \Xi$. Later we will verify this assumption for the specific applications we consider.

For many results in this paper, the specific structure (2.1), (2.2) of the QoI is not important. Therefore, we define

$$X = s(y(\cdot)). \tag{2.3}$$

We assume that $X \in L^1_\rho(\Xi)$. The expected value of a random variable X is $\mathbb{E}[X] = \int_\Xi X(\xi)\rho(\xi)d\xi$.

2.2. Conditional Value-at-Risk. We review basic properties of the Conditional Value-at-Risk (CVaR) at level β, denoted as CVaR_β, that are required within this paper. The CVaR_β is based on the Value-at-Risk (VaR_β). For a given level $\beta \in (0,1)$ the $\mathrm{VaR}_\beta[X]$ is the β-quantile of the random variable X,

$$\mathrm{VaR}_\beta[X] = \min_{t\in\mathbb{R}} \Big\{ \Pr\big[\, \{\xi \in \Xi \,:\, X(\xi) \le t\} \,\big] \ge \beta \Big\}. \tag{2.4}$$

We often use the shorthand notation $\{X \le t\} = \{\xi \in \Xi \,:\, X(\xi) \le t\}$ and the indicator function

$$\mathbb{I}_S(\xi) = \begin{cases} 1 & \text{if } \xi \in S, \\ 0 & \text{else.} \end{cases}$$

Different equivalent definitions of CVaR_β exist. The following definition is due to Rockafellar and Uryasev [17, 18]. The CVaR_β at level $\beta \in (0,1)$ is

$$\mathrm{CVaR}_\beta[X] = \mathrm{VaR}_\beta[X] + \frac{1}{1-\beta}\,\mathbb{E}\left[(X - \mathrm{VaR}_\beta[X])_+\right]. \tag{2.5}$$

The representation (2.5) of $\mathrm{CVaR}_\beta[X]$ motivates the following definition.

Definition 2.1. *The risk region corresponding to* $\mathrm{CVaR}_\beta[X]$ *is given by*

$$\mathbb{G}_\beta[X] := \{\xi \in \Xi \,:\, X(\xi) \ge \mathrm{VaR}_\beta[X]\}. \tag{2.6}$$

As mentioned before, $\mathrm{VaR}_\beta[X]$ and $\mathrm{CVaR}_\beta[X]$ depend only on the values of X that lie in the upper tail of the cumulative distribution function (c.d.f.). In particular, for any set $\widehat{\mathbb{G}}$ with

$$\mathbb{G}_\beta[X] \subset \widehat{\mathbb{G}} \subset \Xi \tag{2.7}$$

we can write the VaR_β in (2.4) as

$$\mathrm{VaR}_\beta[X] = \min_{t\in\mathbb{R}} \left\{ \Pr\left[\left\{ \xi \in \widehat{\mathbb{G}} \ : \ X(\xi) \le t \right\} \right] \ge \beta \right\} \tag{2.8}$$

and the CVaR_β (2.5) as

$$\mathrm{CVaR}_\beta[X] = \mathrm{VaR}_\beta[X] + \frac{1}{1-\beta} \int_{\widehat{\mathbb{G}}} \left(X(\xi) - \mathrm{VaR}_\beta[X]\right)_+ \rho(\xi) d\xi. \tag{2.9}$$

These representations show that we only need values of X in a subdomain $\widehat{\mathbb{G}}$ of the parameter space that includes the risk region. In section 3 we will use ROMs to compute approximations $\widehat{\mathbb{G}}$ of the risk region with the property (2.7), and for parameters $\xi \in \widehat{\mathbb{G}}$ we will approximate the FOM QoI X by the ROM approximation. However, before we introduce ROMs, we briefly discuss sampling-based estimation of CVaR_β, upon which practical ROM-based CVaR_β estimators are based.

2.3. Sampling-based estimation of VaR_β and CVaR_β. Algorithm 2.1 below is used to obtain sampling-based estimates of $\mathrm{VaR}_\beta[X]$ and $\mathrm{CVaR}_\beta[X]$. The algorithm is standard; see, e.g., [18]. For additional information, see [3].

Algorithm 2.1 Sampling-based estimation of VaR_β and CVaR_β.

Input: Set $\Xi_m = \{\xi^{(1)}, \ldots, \xi^{(m)}\} \subset \Xi$ of finitely many parameters and corresponding probabilities $p^{(1)}, \ldots, p^{(m)}$, risk level $\beta \in (0,1)$, and random variable $X : \Xi \to \mathbb{R}$.

Output: Estimate $\widehat{\mathrm{VaR}}_\beta[X]$ and $\widehat{\mathrm{CVaR}}_\beta[X]$.

1: Evaluate X at the parameter samples: $X(\xi^{(1)}), \ldots, X(\xi^{(m)})$.

2: Sort values of X in descending order, relabel the samples so that

$$X(\xi^{(1)}) > X(\xi^{(2)}) > \cdots > X(\xi^{(m)}), \tag{2.10}$$

and reorder the probabilities accordingly (so that $p^{(j)}$ corresponds to $\xi^{(j)}$).

3: Compute an index k_β such that

$$\sum_{j=1}^{k_\beta - 1} p^{(j)} \le 1 - \beta < \sum_{j=1}^{k_\beta} p^{(j)}.$$

4: Set

$$\widehat{\mathrm{VaR}}_\beta[X] = X(\xi^{(k_\beta)}), \tag{2.11}$$

$$\widehat{\mathbb{G}}_\beta[X] = \left\{ \xi \in \Xi_m \ : \ X(\xi) \ge \widehat{\mathrm{VaR}}_\beta[X] \right\}, \tag{2.12}$$

$$\widehat{\mathrm{CVaR}}_\beta[X] = \frac{1}{1-\beta} \sum_{j=1}^{k_\beta - 1} p^{(j)} X(\xi^{(j)}) + \frac{1}{1-\beta} \left(1 - \beta - \sum_{j=1}^{k_\beta - 1} p^{(j)} \right) \widehat{\mathrm{VaR}}_\beta[X]. \tag{2.13}$$

We note that the second term on the right-hand side of (2.13) in Algorithm 2.1 is nonzero for the case $\sum_{j=1}^{k_\beta - 1} p^{(j)} \ne 1 - \beta$ and is based on the idea of splitting the probability atom

at $\mathrm{VaR}_\beta[X]$ (see [18]). An important observation is that the estimates (2.11) and (2.13) depend only on the parameters in the sample risk region $\widehat{\mathbb{G}}_\beta[X]$ (2.12) and their corresponding probabilities. Thus Algorithm 2.1 called with a parameter set Ξ_m and a parameter set $\widetilde{\Xi}$ such that $\widehat{\mathbb{G}}_\beta[X] \subset \widetilde{\Xi} \subset \Xi_m$ gives the same estimates $\widehat{\mathrm{VaR}}_\beta[X]$ and $\widehat{\mathrm{CVaR}}_\beta[X]$.

As discussed in [3, p. 1418], we can also compute confidence intervals using the asymptotic results in [5, sec. 2.1, 2.2]. Since we will use it in our computations, we note that the $100(1-\alpha)\%$ confidence interval (CI) for $\mathrm{CVaR}_\beta[X]$ is

$$\left[\widehat{\mathrm{CVaR}}_\beta[X] - z_\alpha \frac{\widehat{\kappa}_\beta}{\sqrt{m}}, \widehat{\mathrm{CVaR}}_\beta[X] + z_\alpha \frac{\widehat{\kappa}_\beta}{\sqrt{m}}\right], \tag{2.14}$$

where $z_\alpha = \Phi^{-1}(1-\alpha/2)$, Φ is the c.d.f. of the standard normal variable, and $\widehat{\kappa}_\beta = \widehat{\psi}_\beta/(1-\beta)$ with

$$(\widehat{\psi}_\beta)^2 = \frac{1}{m}\sum_{j=1}^{m} \mathbb{I}_{\widehat{\mathbb{G}}_\beta[X]}(\xi^{(j)}) \left(X(\xi^{(j)}) - \widehat{\mathrm{VaR}}_\beta[X]\right)^2 - \left(\frac{1}{m}\sum_{j=1}^{m} \mathbb{I}_{\widehat{\mathbb{G}}_\beta[X]}(\xi^{(j)}) \left(X(\xi^{(j)}) - \widehat{\mathrm{VaR}}_\beta[X]\right)\right)^2.$$

3. Adaptive surrogate-based CVaR_β approximation.

For our target application, FOM (2.2) is a large-scale system that arises from the discretization of a PDE. For given ξ the solution of (2.2) for $y(\xi)$ is expensive, and therefore sampling the QoI (2.1) for CVaR_β computations is expensive. In this section, we propose a method that combines adaptive ROM refinement with knowledge of the CVaR_β computation to generate efficient approximation of the CVaR_β of the QoI (2.1).

We review ROM-based CVaR_β computation in subsection 3.1. In subsection 3.2 we propose our new method that adaptively refines surrogate models to achieve monotonically converging risk regions. Subsection 3.3 then presents our complete algorithm for adaptive surrogate-based CVaR_β approximation.

3.1. Reduced-order models for CVaR_β computation.

A ROM of (2.2) is a model of small dimension, i.e.,

$$F_k(y_k(\xi), \xi) = 0 \qquad \text{for almost all } \xi \in \Xi, \tag{3.1}$$

with $F_k : \mathbb{R}^{N_k} \times \Xi \mapsto \mathbb{R}^{N_k}$, $N_k \ll N$, and an $s_k : \mathbb{R}^{N_k} \mapsto \mathbb{R}$ such that

$$\xi \mapsto s_k(y_k(\xi)) \tag{3.2}$$

is a good approximation of (2.1). We will provide a more detailed discussion of projection-based ROMs in subsection 4.1. For now, let $X_k : \Xi \to \mathbb{R}$, $k = 1, \ldots,$ denote an approximation of the QoI X. We refer to X_k as a model of X. At this point it is not important that the evaluation of X requires the solution of a computationally expensive system (2.2)–(2.1), nor is it important how the models X_k are computed. However, we assume that we have an estimate for the errors between X_k and X, namely

$$|X_k(\xi) - X(\xi)| \le \epsilon_k(\xi) \quad \text{for almost all } \xi \in \Xi, \quad k = 1, \ldots. \tag{3.3}$$

We next show how to construct estimates of the risk region that satisfy (2.7) from approximations X_k of X, and we derive approximations of $\mathrm{VaR}_\beta[X]$ and $\mathrm{CVaR}_\beta[X]$ based on X_k; for more information see our previous work in [3]. Recall the risk region of the QoI X from (2.6). The *ϵ-risk region* associated with X_k is defined as

$$\mathbb{G}_\beta^k = \Big\{\xi \,:\, X_k(\xi) + \epsilon_k(\xi) \geq \mathrm{VaR}_\beta[X_k - \epsilon_k]\Big\}. \tag{3.4}$$

Note that if the error ϵ_k is constant, then the translation equivariance of VaR_β implies $\mathrm{VaR}_\beta[X_k - \epsilon_k] = \mathrm{VaR}_\beta[X_k] - \epsilon_k$. Since

$$X_k(\xi) + \epsilon_k(\xi) \geq X(\xi) \geq X_k(\xi) - \epsilon_k(\xi),$$

the monotonicity of VaR_β gives

$$\mathrm{VaR}_\beta[X] \geq \mathrm{VaR}_\beta[X_k - \epsilon_k].$$

Hence $X_k(\xi) + \epsilon_k(\xi) \geq X(\xi) \geq \mathrm{VaR}_\beta[X] \geq \mathrm{VaR}_\beta[X_k - \epsilon_k]$ for almost all $\xi \in \mathbb{G}_\beta[X]$. Similarly, $X_k(\xi) + \epsilon_k(\xi) \geq X_k(\xi) \geq \mathrm{VaR}_\beta[X_k] \geq \mathrm{VaR}_\beta[X_k - \epsilon_k]$ for almost all $\xi \in \mathbb{G}_\beta[X_k]$. The previous inequalities imply

$$\mathbb{G}_\beta[X] \subset \mathbb{G}_\beta^k \quad \text{and} \quad \mathbb{G}_\beta[X_k] \subset \mathbb{G}_\beta^k. \tag{3.5}$$

Here and in the following we still use the set inclusion $S_1 \subset S_2$ if $\Pr\big[S_1 \setminus S_2\big] = 0$.

We have shown in [3, Thm. 3.3] that if (3.3) holds, then

$$\Big|\mathrm{CVaR}_\beta[X] - \mathrm{CVaR}_\beta[X_k]\Big| \leq \frac{1}{1-\beta} \int_{\mathbb{G}_\beta^k} |X(\xi) - X_k(\xi)| \rho(\xi) d\xi \tag{3.6}$$

and

$$\Big|\mathrm{CVaR}_\beta[X] - \mathrm{CVaR}_\beta[X_k]\Big| \leq \left(1 + \frac{1}{1-\beta}\right) \operatorname{ess\,sup}_{\xi \in \mathbb{G}_\beta^k} \epsilon_k(\xi). \tag{3.7}$$

We note that under continuity conditions on the c.d.f.'s of X and X_k, which often hold, the factor $1 + 1/(1-\beta)$ on the right-hand side of (3.7) can typically be replaced by 1; see [3, Thm. 3.3] for details. Moreover, the first inequality (3.6) appears in the proof of [3, Thm. 3.3].

We see from (3.6)–(3.7) that for the accurate estimation of $\mathrm{CVaR}_\beta[X]$ with a surrogate model, we need a model X_k that is accurate in the ϵ-risk region $\mathbb{G}_\beta^k$. Moreover, applying (2.8) and (2.9) with X and $\widehat{\mathbb{G}}$ replaced by X_k and $\mathbb{G}_\beta^k$ shows that we only need to evaluate X_k in the ϵ-risk region $\mathbb{G}_\beta^k$ to evaluate $\mathrm{CVaR}_\beta[X_k]$.

3.2. Improving CVaR_β computation with adaptive reduced-order models. What happens if $\mathrm{CVaR}_\beta[X_k]$ is not a good enough approximation of $\mathrm{CVaR}_\beta[X]$? In that case, we would like to generate a new model X_{k+1}, so that $\mathrm{CVaR}_\beta[X_{k+1}]$ is a better estimate of $\mathrm{CVaR}_\beta[X]$ than $\mathrm{CVaR}_\beta[X_k]$, or at least so that the upper bound (3.6) for the error is reduced. The upper bound (3.6) for the CVaR_β approximation error is nonincreasing if the ϵ-risk region is

nonexpanding, $\mathbb{G}_\beta^{k+1} \subset \mathbb{G}_\beta^k$, and the approximation error is nonincreasing, $\epsilon_{k+1}(\xi) \le \epsilon_k(\xi)$ for $\xi \in \mathbb{G}_\beta^{k+1}$, since then

$$\text{ess sup}_{\xi\in\mathbb{G}_\beta^{k+1}}\epsilon_{k+1}(\xi) \le \text{ess sup}_{\xi\in\mathbb{G}_\beta^{k+1}}\epsilon_k(\xi) \le \text{ess sup}_{\xi\in\mathbb{G}_\beta^{k}}\epsilon_k(\xi). \tag{3.8}$$

The CVaR_β approximation error is reduced if $\mathbb{G}_\beta^{k+1} \subset \mathbb{G}_\beta^k$, $\Pr\left[\mathbb{G}_\beta^k \setminus \mathbb{G}_\beta^{k+1}\right] > 0$, and $\epsilon_{k+1}(\xi) \le \epsilon_k(\xi) - \delta_k$ for $\xi \in \mathbb{G}_\beta^{k+1}$ and some $\delta_k > 0$.

In general, however, a model X_{k+1} with a smaller error $\epsilon_{k+1} < \epsilon_k$ a.e. in Ξ alone does not guarantee that $\mathbb{G}_\beta^{k+1} \subset \mathbb{G}_\beta^k$, as the following example shows.

Example 3.1. Let $X \ge 0$ be a nonnegative random variable, and consider the surrogate model $X_k = X + \frac{1}{k}(-1)^k X$ with error $\epsilon_k(\xi) = |X(\xi) - X_k(\xi)| = \frac{1}{k}X$. For $k = 1, \ldots$ the ϵ-risk regions are

$$\begin{aligned}
\mathbb{G}_\beta^{2k-1} &= \{\xi : X_{2k-1} + \epsilon_{2k-1} \ge \text{VaR}_\beta\left[X_{2k-1} - \epsilon_{2k-1}\right]\} \\
&= \left\{\xi : X(\xi) \ge \text{VaR}_\beta\left[X - \frac{2}{2k-1}X\right]\right\} = \left\{\xi : X(\xi) \ge \frac{2k-3}{2k-1}\text{VaR}_\beta[X]\right\}, \\
\mathbb{G}_\beta^{2k} &= \{\xi : X_{2k} + \epsilon_{2k} \ge \text{VaR}_\beta[X_{2k} - \epsilon_{2k}]\} \\
&= \left\{\xi : X(\xi) + \frac{1}{k}X(\xi) \ge \text{VaR}_\beta[X]\right\} = \left\{\xi : X(\xi) \ge \frac{k}{k+1}\text{VaR}_\beta[X]\right\}.
\end{aligned}$$

We have the inclusions

$$\mathbb{G}_\beta^{2k} \subset \mathbb{G}_\beta^{2k-1},$$

since $(2k-3)/(2k-1) < k/(k+1)$, but

$$\mathbb{G}_\beta^{2k} \subset \mathbb{G}_\beta^{2k+1},$$

since $(2(k+1)-3)/(2(k+1)-1) < k/(k+1)$. Thus, there is no monotonicity (in the sense of inclusion) of the ϵ-risk regions. Note that the ϵ-risk regions are based on the models X_k. While the models X_k become more accurate, the lack of monotonicity of the ϵ-risk regions is due to the fact that here the ϵ_k neighborhoods around the X_k are alternatingly below or above the true X.

When does the use of a new model X_{k+1} improve the approximation of $\text{CVaR}_\beta[X]$? A sufficient condition for improvement is the monotonicity condition

$$X_k(\xi) + \epsilon_k(\xi) \ge X_{k+1}(\xi) + \epsilon_{k+1}(\xi) \ge X(\xi) \ge X_{k+1}(\xi) - \epsilon_{k+1}(\xi) \ge X_k(\xi) - \epsilon_k(\xi) \quad \text{a.e. in } \Xi. \tag{3.9}$$

In fact, monotonicity of VaR_β gives $\text{VaR}_\beta[X] \ge \text{VaR}_\beta[X_{k+1} - \epsilon_{k+1}] \ge \text{VaR}_\beta[X_k - \epsilon_k]$. These inequalities and (3.9) yield

$$\begin{aligned}
&X_k(\xi) + \epsilon_k(\xi) \ge X_{k+1}(\xi) + \epsilon_{k+1}(\xi) \ge X(\xi) \ge \text{VaR}_\beta[X] \\
&\ge \text{VaR}_\beta[X_{k+1} - \epsilon_{k+1}] \ge \text{VaR}_\beta[X_k - \epsilon_k] \qquad \text{a.e. in } \mathbb{G}_\beta[X]
\end{aligned}$$

and

$$X_k(\xi) + \epsilon_k(\xi) \geq X_{k+1}(\xi) + \epsilon_{k+1}(\xi) \geq \mathrm{VaR}_\beta[X_{k+1} - \epsilon_{k+1}] \geq \mathrm{VaR}_\beta[X_k - \epsilon_k] \quad \text{a.e. in } \mathbb{G}_\beta^k,$$

which imply

$$\mathbb{G}_\beta[X] \subset \mathbb{G}_\beta^{k+1} \subset \mathbb{G}_\beta^k. \tag{3.10}$$

Unfortunately, models X_k, $k = 1, \ldots,$ typically do not satisfy the monotonicity relations (3.9), as the simple Example 3.1 shows. However, we can combine the models X_k, $k = 1, \ldots,$ into models $\widetilde{X}_k$, $k = 1, \ldots,$ that satisfy (3.9). We define these new models $\widetilde{X}_k$ in the next lemma.

Lemma 3.2. *If the models X_k and error functions ϵ_k satisfy (3.3), $k = 1, \ldots,$ then the models $\widetilde{X}_k$ and corresponding error functions $\widetilde{\epsilon}_k$ defined by $\widetilde{X}_1 = X_1$, $\widetilde{\epsilon}_1 = \epsilon_1$, and*

$$\widetilde{X}_{k+1} = \frac{1}{2}\Big(\max\Big\{X_{k+1} - \epsilon_{k+1},\ \widetilde{X}_k - \widetilde{\epsilon}_k\Big\} + \min\Big\{X_{k+1} + \epsilon_{k+1},\ \widetilde{X}_k + \widetilde{\epsilon}_k\Big\}\Big), \tag{3.11a}$$

$$\widetilde{\epsilon}_{k+1} = \frac{1}{2}\Big(\min\Big\{X_{k+1} + \epsilon_{k+1},\ \widetilde{X}_k + \widetilde{\epsilon}_k\Big\} - \max\Big\{X_{k+1} - \epsilon_{k+1},\ \widetilde{X}_k - \widetilde{\epsilon}_k\Big\}\Big) \tag{3.11b}$$

for $k = 1, \ldots,$ satisfy the monotonicity relations (3.9).

The model construction (3.11) is illustrated in Figure 1.

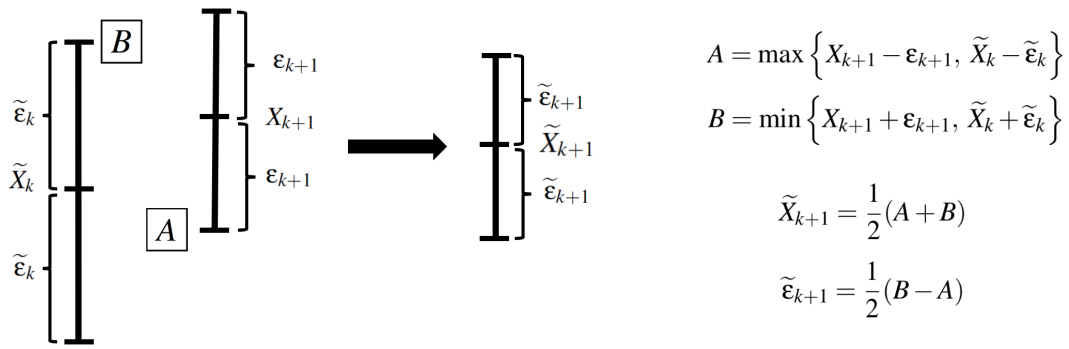

Figure 1. *Illustration of the model construction (3.11). The true function X is contained in the intervals $[\widetilde{X}_k - \widetilde{\epsilon}_k, \widetilde{X}_k + \widetilde{\epsilon}_k]$ and $[X_{k+1} - \epsilon_{k+1}, X_{k+1} + \epsilon_{k+1}]$. While the second interval is smaller, it is not contained in the first. The model (3.11) is constructed so that $[\widetilde{X}_{k+1} - \widetilde{\epsilon}_{k+1}, \widetilde{X}_{k+1} + \widetilde{\epsilon}_{k+1}]$ includes the true model and is nested.*

Proof. The proof is by induction. We have initialized $\widetilde{X}_1 = X_1$, and $\widetilde{\epsilon}_1 = \epsilon_1$, so $\widetilde{X}_1, \widetilde{\epsilon}_1$ satisfy (3.3), since by assumption, X_1, ϵ_1 satisfy (3.3).

Now, suppose that $(\widetilde{X}_1, \widetilde{\epsilon}_1), \ldots, (\widetilde{X}_k, \widetilde{\epsilon}_k)$ satisfy the monotonicity relations (3.9). Since $(\widetilde{X}_k, \widetilde{\epsilon}_k)$ and $(X_{k+1}, \epsilon_{k+1})$ satisfy (3.3),

$$\max\Big\{X_{k+1} - \epsilon_{k+1},\ \widetilde{X}_k - \widetilde{\epsilon}_k\Big\} \leq X \leq \min\Big\{X_{k+1} + \epsilon_{k+1},\ \widetilde{X}_k + \widetilde{\epsilon}_k\Big\}.$$

By construction of $\widetilde{X}_{k+1}$ and $\widetilde{\epsilon}_{k+1}$,

$$\begin{aligned}\widetilde{X}_k - \widetilde{\epsilon}_k &\le \max\left\{X_{k+1} - \epsilon_{k+1},\ \widetilde{X}_k - \widetilde{\epsilon}_k\right\} = \widetilde{X}_{k+1} - \widetilde{\epsilon}_{k+1} \\ &\le X \le \widetilde{X}_{k+1} + \widetilde{\epsilon}_{k+1} = \min\left\{X_{k+1} + \epsilon_{k+1},\ \widetilde{X}_k + \widetilde{\epsilon}_k\right\} \le \widetilde{X}_k + \widetilde{\epsilon}_k;\end{aligned}$$

i.e., the monotonicity relations (3.9) are satisfied for $(\widetilde{X}_1, \widetilde{\epsilon}_1), \ldots, (\widetilde{X}_{k+1}, \widetilde{\epsilon}_{k+1})$. ∎

The error (3.11b) satisfies

$$\widetilde{\epsilon}_{k+1} \le \min\{\widetilde{\epsilon}_k, \epsilon_{k+1}\} \quad \text{a.e. in } \Xi. \tag{3.12}$$

Let $\widetilde{\mathbb{G}}_\beta^k$ be the ϵ-risk region (3.4) associated with $\widetilde{X}_k$, $\widetilde{\epsilon}_k$. The estimate (3.12) implies that to achieve

$$\widetilde{\epsilon}_{k+1}(\xi) < \widetilde{\epsilon}_k(\xi) \quad \text{a.e. in } \widetilde{\mathbb{G}}_\beta^k \tag{3.13}$$

we only need to improve the model X_{k+1} in the small ϵ-risk region $\widetilde{\mathbb{G}}_\beta^k$—not in the entire parameter region Ξ—i.e., we only need that

$$\epsilon_{k+1}(\xi) \le \widetilde{\epsilon}_k(\xi) - \delta_k \quad \text{a.e. in } \widetilde{\mathbb{G}}_\beta^k \tag{3.14}$$

for some $\delta_k > 0$. We summarize the improvement result in the following theorem.

Theorem 3.3. *If $\widetilde{X}_k$, $k = 1, \ldots,$ are the models with corresponding error functions $\widetilde{\epsilon}_k$, $k = 1, \ldots,$ defined in (3.11a), (3.11b), and $\widetilde{\mathbb{G}}_\beta^k$, $k = 1, \ldots,$ are the ϵ-risk regions (3.4) associated with $\widetilde{X}_k$, $\widetilde{\epsilon}_k$, then*

$$\left|\mathrm{CVaR}_\beta[X] - \mathrm{CVaR}_\beta[\widetilde{X}_k]\right| \le \left(1 + \frac{1}{1-\beta}\right) \operatorname{ess\,sup}_{\xi \in \widetilde{\mathbb{G}}_\beta^k} \widetilde{\epsilon}_k(\xi), \quad k = 1, 2, \ldots, \tag{3.15}$$

and

$$\mathbb{G}_\beta[X] \subset \widetilde{\mathbb{G}}_\beta^{k+1} \subset \widetilde{\mathbb{G}}_\beta^k, \quad k = 1, 2, \ldots. \tag{3.16}$$

Moreover, if $\epsilon_{k+1}(\xi) \le \widetilde{\epsilon}_k(\xi) - \delta_k$ a.e. in $\widetilde{\mathbb{G}}_\beta^k$ for some $\delta_k > 0$, then

$$\operatorname{ess\,sup}_{\xi \in \widetilde{\mathbb{G}}_\beta^{k+1}} \widetilde{\epsilon}_{k+1}(\xi) \le \operatorname{ess\,sup}_{\xi \in \widetilde{\mathbb{G}}_\beta^k} \widetilde{\epsilon}_k(\xi) - \delta_k. \tag{3.17}$$

Proof. Since the models $\widetilde{X}_k$, $k = 1, 2, \ldots,$ satisfy the monotonicity relations (3.9), the error estimate (3.15) is just (3.7); see [3, Thm. 3.3]. The inclusions (3.16) follow from the arguments used to derive (3.10). The error reduction (3.17) follows from (3.12)–(3.14) and (3.16). ∎

Having defined new models $\widetilde{X}_k$ and errors $\widetilde{\varepsilon}_k$, we revisit Example 3.1. We show that for this example problem, the monotonicity of the ε-risk regions is now indeed satisfied.

Example 3.4. Recall the setup from Example 3.1, where $X \geq 0$ is a nonnegative random variable and a surrogate model is $X_k = X + \frac{1}{k}(-1)^k X$ with error $\epsilon_k(\xi) = |X(\xi) - X_k(\xi)| = \frac{1}{k}X$. We now construct $\widetilde{X}_k$, $\widetilde{\varepsilon}_k$ following Lemma 3.2. We have

$$\widetilde{X}_1 = X_1 = X + 1(-1)^1 X = 0, \qquad \widetilde{\varepsilon}_1 = \varepsilon_1 = X,$$

and with $X \geq 0$ and evaluating equations (3.11a)–(3.11b), we find that, for this particular example, $\widetilde{X}_k = X$, $\widehat{\varepsilon}_k = 0$ for $k \geq 2$. Moreover, the first risk region is $\widetilde{\mathbb{G}}_\beta^1 = \{\xi \ : \ X \geq \mathrm{VaR}_\beta[-X]\}$ $= \Xi$ and the subsequent risk regions are $\widetilde{\mathbb{G}}_\beta^k = \{\xi \ : \ X(\xi) \geq \mathrm{VaR}_\beta[X]\} = \mathbb{G}_\beta[X]$, the true risk region of the full order model X, for $k \geq 2$. Consequently,

$$\widetilde{\mathbb{G}}_\beta^1 \supset \widetilde{\mathbb{G}}_\beta^2 = \widetilde{\mathbb{G}}_\beta^k = \mathbb{G}_\beta[X], \quad k \geq 2;$$

i.e., the risk regions are shrinking monotonically and contain the true risk region, as guaranteed by Theorem 3.3. The fact that the second adjusted risk region is already identical to the true risk region of the FOM X is particular to this artificial example.

3.3. Algorithm for surrogate-based CVaR$_\beta$ approximation. The previous results lead to the following Algorithm 3.1 that adaptively constructs models X_k based on estimates $\widetilde{\mathbb{G}}_\beta^k$ of the risk region $\mathbb{G}_\beta[X]$. As noted earlier, applying (2.8) and (2.9) with X and $\widehat{\mathbb{G}}$ replaced by $\widetilde{X}_k$ and $\widetilde{\mathbb{G}}_\beta^k \supset \mathbb{G}_\beta[\widetilde{X}_k]$ shows that we only need to evaluate $\widetilde{X}_k$ in the ϵ-risk region $\widetilde{\mathbb{G}}_\beta^k \subset \widetilde{\mathbb{G}}_\beta^{k-1}$ to evaluate $\mathrm{CVaR}_\beta[\widetilde{X}_k]$. Furthermore, X_{k+1} only needs to improve upon $\widetilde{X}_k$ in the ϵ-risk region $\widetilde{\mathbb{G}}_\beta^k$; i.e., we only need (3.14). Since $\widetilde{\mathbb{G}}_\beta^k$ tend to be small (in probability) subsets of the parameter space Ξ, the adaptive generation of the models by the previous algorithm can lead to large computational savings.

Algorithm 3.1 Surrogate-based CVaR$_\beta$ estimation.

Input: Desired error tolerance TOL, maximum number of iterations $k_{\max}$, risk-level $\beta \in (0,1)$.

Output: $\mathrm{CVaR}_\beta[\widetilde{X}_k]$ and $\widetilde{\epsilon}_k^G$ such that $|\mathrm{CVaR}_\beta[\widetilde{X}_k] - \mathrm{CVaR}_\beta[X]| \leq \widetilde{\epsilon}_k^G \leq \mathrm{TOL}$ or $k = k_{\max}$.

1: Set $k = 1$ and generate model $\widetilde{X}_1 = X_1$, $\widetilde{\epsilon}_1 = \epsilon_1$ with (3.3).
2: Compute $\mathrm{CVaR}_\beta[\widetilde{X}_1]$ and $\epsilon_1^G = \operatorname{ess\,sup}_{\xi \in \widetilde{\mathbb{G}}_\beta^1} \widetilde{\epsilon}_1(\xi)$.
3: **while** $\widetilde{\epsilon}_k^G > \mathrm{TOL}$ and $k < k_{\max}$ **do**
4: Compute model X_{k+1} and error function ϵ_{k+1} with (3.3) and (3.14).
5: Compute model $\widetilde{X}_{k+1}$ and error function $\widetilde{\epsilon}_{k+1}$ as in (3.11a) and (3.11b).
6: Compute $\mathrm{VaR}_\beta[\widetilde{X}_{k+1}]$, $\mathrm{CVaR}_\beta[\widetilde{X}_{k+1}]$, ϵ-risk region $\widetilde{\mathbb{G}}_\beta^{k+1}$, and error in ϵ-risk region

$$\widetilde{\epsilon}_k^G = \operatorname{ess\,sup}_{\xi \in \widetilde{\mathbb{G}}_\beta^{k+1}} \widetilde{\epsilon}_{k+1}(\xi).$$

7: Set $k = k + 1$ and continue.
8: **end while**

Before we address several implementation details that are important for the realization of

Algorithm 3.1 in combination with ROMs, we comment on the extension of our idea to the estimation of probability of failure from a QoI X.

Remark 3.5. There is a close relationship between probability of failure and the Value-at-Risk. If failure of a system is defined as $X(\xi) \geq X_0$, then the probability of failure is $\Pr[\mathbb{F}[X]]$, where $\mathbb{F}[X] := \{\xi \in \Xi \,:\, X(\xi) \geq X_0\}$ is the failure region. If (3.3) holds and $X_k(\xi) - \epsilon_k(\xi) \geq X_0$, then

$$X(\xi) \geq X_k(\xi) - \epsilon_k(\xi) \geq X_0.$$

Similarly, if $\xi \in \mathbb{F}[X]$, then

$$\epsilon_k(\xi) + X_k(\xi) \geq X(\xi) \geq X_0.$$

Hence, the failure region $\mathbb{F}[X]$ can be estimated as

$$\{\xi \in \Xi \,:\, X_k(\xi) - \epsilon_k(\xi) \geq X_0\} \subset \mathbb{F}[X] \subset \{\xi \in \Xi \,:\, X_k(\xi) + \epsilon_k(\xi) \geq X_0\}.$$

This can be used in the estimation of failure probability, as, e.g., in [1]. Since the models $\widetilde{X}_k$ and corresponding error functions $\widetilde{\epsilon}_k$ satisfy the monotonicity relations (3.9), we have that

$$\begin{aligned}\left\{\xi \in \Xi \,:\, \widetilde{X}_k(\xi) - \widetilde{\epsilon}_k(\xi) \geq X_0\right\} \subset \left\{\xi \in \Xi \,:\, \widetilde{X}_{k+1}(\xi) - \widetilde{\epsilon}_{k+1}(\xi) \geq X_0\right\} \subset \mathbb{F}[X],\\ \mathbb{F}[X] \subset \left\{\xi \in \Xi \,:\, \widetilde{X}_{k+1}(\xi) + \widetilde{\epsilon}_{k+1}(\xi) \geq X_0\right\} \subset \left\{\xi \in \Xi \,:\, \widetilde{X}_k(\xi) + \widetilde{\epsilon}_k(\xi) \geq X_0\right\}.\end{aligned}$$

Thus, the models $\tilde{X}_k$ and error bounds $\tilde{\epsilon}_k$ can be used for failure probability estimation as well and yield monotonely converging failure regions.

4. Implementation.

This section discusses an implementation of Algorithm 3.1 to estimate the CVaR_β of a QoI defined via (2.3) and a linear version of the state equation (2.2). The implementation uses projection-based ROMs and sampling-based estimation of VaR_β and CVaR_β for the ROMs. We begin by reviewing the basic form of projection-based ROMs and error estimates in subsection 4.1. The standard greedy sampling strategy and differences from our proposed adaptive sampling strategy are discussed in subsection 4.2. The combination of ROM adaptation and sampling-based CVaR_β computation is then presented in subsection 4.3.

4.1. Error estimation for projection-based ROMs.

We summarize results on error estimation for projection-based ROMs for linear parametric systems. These results are by now standard and can be found in, e.g., [9, 4, 15, 2]. Given $A(\xi) \in \mathbb{R}^{N\times N}$, $b(\xi) \in \mathbb{R}^n$, parameters $\xi \in \Xi$, and $s : \mathbb{R}^N \to \mathbb{R}$, we consider the FOM

$$A(\xi)y(\xi) = b(\xi) \qquad \text{for } \xi \in \Xi, \tag{4.1}$$

and corresponding QoI

$$X(\xi) = s(y(\xi)) \in \mathbb{R}. \tag{4.2}$$

This fits the framework of section 2.1 with $F(y, \xi) = A(\xi)y - b(\xi)$. We assume that

$$\|A(\xi)\| \leq \gamma, \ \|A(\xi)^{-1}\| \leq \alpha^{-1}. \tag{4.3}$$

We use α^{-1} to denote the upper bound for the inverse, since this notation is closer to what is used, e.g., in [9, 4, 15, 2], where (4.1) arises from the discretization of an elliptic PDE and α is related to coercivity constants of the PDE.

The ROM is specified by a matrix $V_k \in \mathbb{R}^{N\times N_k}$ of rank N_k and is given by

$$V_k^T A(\xi)V_k y_k(\xi) = V_k^T b(\xi) \qquad \text{for } \xi \in \Xi, \tag{4.4}$$

and corresponding QoI

$$X_k(\xi) = s(V_k y_k(\xi)) \in \mathbb{R}. \tag{4.5}$$

We assume that the matrix V_k is such that (4.4) has a unique solution for all $\xi \in \Xi$. To simplify the presentation we also assume that the computation of quantities like $V_k^T A(\xi)V_k$, $A(\xi)V_k$, and $A(\xi)^T V_k$ for $\xi \in \Xi$ is computationally inexpensive, which is the case if $A(\xi)$ and $b(\xi)$ admit an affine parametric dependence; see, e.g, [2, sec. 2.3.5], [4, sec. 3.3], or [15, sec. 3.4].

The equations (4.1) and (4.4) imply the basic error estimate for the state

$$\|y(\xi) - V_k y_k(\xi)\| \le \alpha^{-1} \|A(\xi)V_k y_k(\xi) - b(\xi)\| \qquad \text{for } \xi \in \Xi. \tag{4.6}$$

If s is Lipschitz continuous, i.e., $|s(y) - s(z)| \le L\|y - z\|$ for all $y, z \in \mathbb{R}^N$, then the basic error estimate

$$|X(\xi) - X_k(\xi)| \le \epsilon_k(\xi) := \frac{L}{\alpha} \|A(\xi)V_k y_k(\xi) - b(\xi)\| \qquad \text{for } \xi \in \Xi \tag{4.7}$$

holds for the QoI. This is the realization of the bound (3.3). Improved error estimates for linear QoIs can be obtained based on solutions of a dual or adjoint equation; see, e.g, [2, sec. 2.3.4], [4, sec. 4], [9], or [15, sec. 3.6].

4.2. Greedy ROM construction and estimation of CVaR$_\beta$. In a standard greedy algorithm, the ROM specified by V_k is updated by computing the FOM solution (4.1) at $\xi^{(k)} = \arg\max_{\xi\in\Xi} \epsilon_k(\xi)$ and setting $V_{k+1} = [V_k, y(\xi^{(k)})]$. In practice, one often does not simply add the FOM solution $y(\xi^{(k)})$ as a column to V_k, but instead computes an orthonormal basis (see, e.g., [4, sec. 3.2.2] or [15, Chapter 7]).

In our recent work [3] we have used this greedy procedure and the resulting ROMs without adjustment. That is, we have used $\widetilde{X}_k = X_k$ and $\widetilde{\epsilon}_k = \epsilon_k$, which implies $\widetilde{\mathbb{G}}_\beta^k = \mathbb{G}_\beta^k$ and $\widetilde{\epsilon}_k^G = \epsilon_k^G$. While for each ROM a CVaR$_\beta$ error bound holds, this approach has two deficiencies. First, as discussed in subsection 3.2, the ROM CVaR$_\beta$ estimation error is not guaranteed to decrease as we go from ROM X_k to ROM X_{k+1}. Second, the standard greedy procedure seeks the maximum of $\epsilon_k(\xi)$ over the entire parameter space. Even though computation of $\epsilon_k(\xi)$ only requires ROM (4.4) solutions and FOM residual evaluations, these evaluations at a large number of points $\xi \in \Xi$ are still expensive. Moreover, the ROM error over the ϵ-risk region determines the ROM CVaR$_\beta$ estimation error; see Theorem 3.3. Limiting the greedy approach to this smaller set tends to decrease this error faster.

Our adaptive approach corrects these deficiencies: It uses the modified ROMs $\widetilde{X}_k$ and error bounds $\widetilde{\epsilon}_k$ introduced in Lemma 3.2 to guarantee monotonicity of the resulting ROM CVaR$_\beta$ estimation error, and it selects FOM snapshots by maximizing the current ROM error bound $\widetilde{\epsilon}_k$ only over the small ϵ-risk region $\widetilde{\mathbb{G}}_\beta^k$. The details are specified in the next section.

4.3. Adaptive ROM construction and estimation of CVaR_β. The sampling-based version of Algorithm 3.1 is presented in Algorithm 4.1 below. In each step k of the algorithm a projection-based ROM (4.4) of size $N_k \times N_k$ is computed, as well as the corresponding ROM QoI (4.5). To improve the ROM, snapshots of the FOM are computed using the greedy approach limited to the current estimate $\widetilde{\mathbb{G}}_\beta^k$ of the risk region. As (3.13) and (3.14) show, we only need to improve X_{k+1} in $\widetilde{\mathbb{G}}_\beta^k$ in order to improve the estimate of CVaR_β. Since we work with a discrete sample space Ξ_m, (3.13) implies (3.14) with some $\delta_k > 0$. Furthermore, we can easily check whether the condition $\max_{\xi \in \widetilde{\mathbb{G}}_k^\beta} \widetilde{\epsilon}_{k+1} < \widetilde{\epsilon}_k^G$ holds, which is sufficient for $\widetilde{\epsilon}_{k+1}^G$ to be less than $\widetilde{\epsilon}_k^G$, and is weaker than condition (3.13). We recommend using this last condition in practice because it can sometimes be achieved with fewer FOM snapshots than are needed to enforce (3.13). In Algorithm 4.1 we limit the number of snapshots that are added in each iteration by ℓ_{max}. Even though the (possibly pessimistic) error bound may not be reduced, the actual error may reduce. Finally, in Algorithm 4.1 we simply add the FOM solution $y(\xi^{(\ell)})$ to the current ROM basis, but in practice we compute orthogonal bases.

Algorithm 4.1 Adaptive construction of ROMs for CVaR_β estimation.

Input: Linear FOM (4.1) with (4.3) and Lipschitz continuous QoI (4.2). Parameter samples $\Xi_m = \{\xi^{(1)}, \ldots, \xi^{(m)}\}$ with probabilities $p^{(1)}, \ldots, p^{(m)}$. Risk level $\beta \in (0,1)$. Tolerance TOL.

Output: $\widehat{\mathrm{CVaR}}_\beta[\widetilde{X}_k]$ and $\widetilde{\epsilon}_k^G$ such that $|\widehat{\mathrm{CVaR}}_\beta[\widetilde{X}_k] - \widehat{\mathrm{CVaR}}_\beta[X]| \le \widetilde{\epsilon}_k^G \le \mathrm{TOL}$ or $k = k_{\max}$.

Set $k = 1$ and generate $V_1 \in \mathbb{R}^{N \times N_1}$ and ROM (4.4), $\widetilde{X}_1(\xi) = X_1(\xi) = (V_1^T c(\xi))^T y_1(\xi)$ with error function $\widetilde{\epsilon}_1(\xi) = \epsilon_1(\xi)$ given by (4.7).
Set $\widetilde{\mathbb{G}}_\beta^0 = \Xi_m$.
while $k < k_{\max}$ **do**
  Call Algorithm 2.1 with $\Xi_m = \widetilde{\mathbb{G}}_\beta^{k-1}$, corresponding probabilities $p^{(j)}$, and $X = \widetilde{X}_k$ to compute $\widehat{\mathrm{VaR}}_\beta[\widetilde{X}_k]$, and $\widehat{\mathrm{CVaR}}_\beta[\widetilde{X}_k]$.
  Call Algorithm 2.1 with $\Xi_m = \widetilde{\mathbb{G}}_\beta^{k-1}$, corresponding probabilities $p^{(j)}$, and $X = \widetilde{X}_k - \widetilde{\epsilon}_k$ to compute $\widehat{\mathrm{VaR}}_\beta[\widetilde{X}_k - \widetilde{\epsilon}_k]$.
  Estimate $\widetilde{\mathbb{G}}_\beta^k = \{\xi^{(j)} \in \widetilde{\mathbb{G}}_\beta^{k-1} : \widetilde{X}_k(\xi^{(j)}) + \widetilde{\epsilon}_k(\xi^{(j)}) \ge \widehat{\mathrm{VaR}}_\beta[\widetilde{X}_k - \widetilde{\epsilon}_k]\}$ and set $\widetilde{\epsilon}_k^G = \max\{\widetilde{\epsilon}_k(\xi^{(j)}) : \xi^{(j)} \in \widetilde{\mathbb{G}}_\beta^k\}$.
  if $\widetilde{\epsilon}_k^G < \mathrm{TOL}$ **then**
    break
  end if
  Set $\ell = 1$ (number of snapshots to add) and $V_{k+1} = V_k$
  while $\ell < \ell_{max}$ **do**
    Compute the FOM solution $y(\xi^{(\ell)})$ at $\xi^{(\ell)} = \arg\max_{\xi \in \widetilde{\mathbb{G}}_\beta^k} \widetilde{\epsilon}_k(\xi)$.
    Update ROM matrix $V_{k+1} \leftarrow [V_{k+1}, y(\xi^{(\ell)})]$ and set $N_{k+1} = N_k + \ell$.
    Construct the new ROM of size N_{k+1} and evaluate $X_{k+1}(\xi^{(j)})$ and $\epsilon_{k+1}(\xi^{(j)})$ for $\xi^{(j)} \in \widetilde{\mathbb{G}}_\beta^k$.
    Compute model $\widetilde{X}_{k+1}(\xi^{(j)})$ and error function $\widetilde{\epsilon}_{k+1}(\xi^{(j)})$ as in (3.11a) and (3.11b) for $\xi^{(j)} \in \widetilde{\mathbb{G}}_\beta^k$.
    if $\widetilde{\epsilon}_{k+1}(\xi^{(j)}) < \widetilde{\epsilon}_k(\xi^{(j)})$ for $\xi^{(j)} \in \widetilde{\mathbb{G}}_\beta^k$ $\left(\text{or } \max \widetilde{\epsilon}_{k+1}(\xi) < \widetilde{\epsilon}_k^G \text{ for } \xi^{(j)} \in \widetilde{\mathbb{G}}_\beta^k\right)$ **then**
      break
    end if
    Set $\ell = \ell + 1$.
  end while
  Set $k = k + 1$ and continue.
end while

5. Numerical results. We now apply our Algorithm 4.1 to the so-called thermal fin problem with varying numbers of random variables. We describe the test problem in subsection 5.1 and discuss the format of our reported results in subsection 5.2. The results for the case of two, three, and six random variables are shown in subsection 5.3–subsection 5.5.

5.1. Thermal fin model. We consider a thermal fin with fixed geometry as shown in Figure 2, consisting of a vertical post with horizontal fins attached. We briefly review the problem here and refer the reader to [10, 14] for more details. In particular, [14, sec. 3] discusses the efficiency of the derived reduced-basis (RB) error bounds for the thermal fin problem. The thermal fin consists of four horizontal subfins with width $L = 2.5$, thickness $t = 0.25$, and a fin post with unit width and height four. The fin is parametrized by the fin conductivities $k_i, i = 1, \ldots, 4$, and post conductivity k_0, as well as the Biot number Bi which is a nondimensionalized heat transfer coefficient for thermal transfer from the fins to the surrounding air. Thus, the system parameters are $[k_0,\ k_1, k_2,\ k_3,\ k_4, Bi] \in [0.1,\ 1] \times [0.1,\ 2]^4 \times [0.01,\ 0.1]$. In our experiments some or all of these parameters play the role of the random variables $\boldsymbol{\xi}$, which are uniformly distributed in the parameter space above. The system is governed by an elliptic PDE in two spatial dimensions $x = [x_1, x_2]^T$ whose solution is the temperature field $= y(x, \xi)$. We consider cases when only k_0 and Bi are random (subsection 5.3), k_0, k_1, and Bi are random (subsection 5.4), and finally, when all six parameters are random (subsection 5.5).

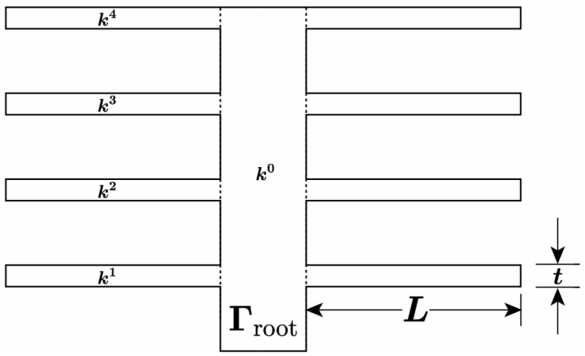

Figure 2. *Thermal fin geometry and model parameters.*

The fin conducts heat away from the root Γ_{root}, so the lower the root temperature, the more effective the thermal fin. Thus, as a QoI we consider the average temperature at the root, i.e.,

$$X(\xi) = \int_{\Gamma_{\text{root}}} y(x, \xi) \mathrm{d}x.$$

The FOM is a finite element discretization with $N = 4,760$ degrees of freedom. The ROMs are RB approximations y_k; see [14] for details of RB methods for the thermal fin problem. The ROM-based estimates are compared to a FOM-sampling-based estimation of $\text{CVaR}_\beta[X]$ using Algorithm 2.1.

We consider the problem with two random variables, three random variables, and six random variables, as specified in sections 5.3–5.5 below. The CVaR_β estimates and corresponding

confidence interval (CI) widths computed with several sample sizes $|\Xi_m|$ using the FOM are shown in Table 1.

Table 1

$CVaR_\beta$ estimates for $\beta = 0.99$ and corresponding confidence interval (CI) widths computed with several sample sizes $|\Xi_m|$. For $|\Xi_m| = 5,000$ samples the CI widths are less than 5% of the CVaR estimates.

	$\widehat{\text{CVaR}}_\beta$	Width CI	$\lvert\Xi_m\rvert$
2 RV	12.404	0.437	5,000
2 RV	11.956	0.326	10,000
2 RV	11.984	0.232	20,000
3 RV	10.379	0.405	5,000
3 RV	10.187	0.274	10,000
3 RV	10.546	0.194	20,000
6 RV	10.435	0.421	5,000
6 RV	10.510	0.296	10,000
6 RV	10.419	0.189	20,000

Since the CI widths are less than 5% of the CVaR estimates computed with $5,000$ samples, we use $|\Xi_m| = 5,000$ samples in the following computations.

Since the ROM needs to approximate the FOM on these sets of samples, we use them as training sets to construct the ROMs. The thermal fin model and the RB ROM fit exactly into the framework of subsection 4.1. We use the error bound (4.7) in the adaptive CVaR_β approximation below. The risk level β is set to

$$\beta = 0.99.$$

In the following sections we report the numerical results obtained with the adaptive Algorithm 4.1 and with the greedy approach outlined in subsection 4.2. The latter corresponds to Algorithm 4.1 with $\widetilde{X}_k = X_k$, $\widetilde{\epsilon}_k = \epsilon_k$, $\widetilde{\mathbb{G}}^k_\beta = \mathbb{G}^k_\beta$, and $\widetilde{\epsilon}^G_k = \epsilon^G_k$. Moreover, in the latter case, in step 12 we compute the FOM solution $y(\xi^{(\ell)})$ at $\xi^{(\ell)} = \arg\max_{\xi\in\Xi_m} \epsilon_k(\xi)$ to update the ROM X_k. In steps 4 and 5 we call Algorithm 2.1 with the full set Ξ_m of parameters. Since computation of $\arg\max_{\xi\in\Xi_m} \epsilon_k(\xi)$ in step 12 already requires computation of X_k and ϵ_k at all parameters in Ξ_m, this modification of steps 4 and 5 is insignificant.

5.2. Overview of reported data. We report the results of the CVaR_β estimation using the adaptive and the greedy approaches in Table 2–Table 7 in subsection 5.3–subsection 5.5 below. Each table contains the same information, which we discuss for convenience here:

- $\widehat{\text{CVaR}}_\beta$ reports the sampling-based CVaR_β estimates for the FOM or the kth ROM;
- "Width CI" is the width of the CI (2.14) of the sampling-based CVaR_β estimate using the FOM or the kth ROM;
- "Abs error" is $|\widehat{\text{CVaR}}_\beta[X] - \widehat{\text{CVaR}}_\beta[X_k]|$, i.e., the error between estimates with the FOM and the kth ROM (via the adaptive or greedy approach);
- ϵ^G_k and $\widetilde{\epsilon}^G_k$ are the CVaR_β error bounds computed using the ROM X_k / modified ROM $\widetilde{X}_k$;
- $|\mathbb{G}^k_\beta|$ and $|\widetilde{\mathbb{G}}^k_\beta|$ denote the percentage of "volume" measured in probability occupied by the ϵ-risk region for the ROM X_k / $\widetilde{X}_k$ within the parameter region Ξ;

- N_k is the size of the kth ROM;
- $|\Xi_m|$ is the number of samples at which the current ROM has to be evaluated.

5.3. Results for two random variables. We start with a problem with two random variables $\boldsymbol{\xi} = (k_0, Bi)$ uniformly distributed in $\Xi = [0.1, 1] \times [0.01, 0.1]$. Having two random variables allows us to visualize both the risk regions and the error estimates. We fix $k_1 = k_2 = k_3 = k_4 = 0.1$.

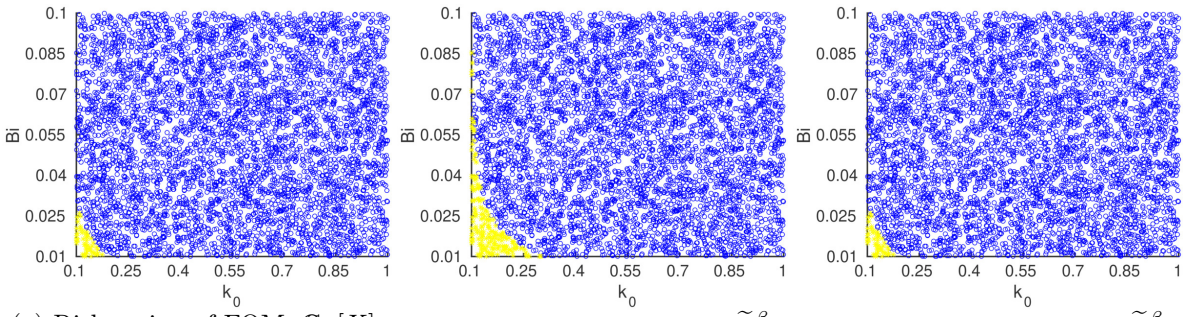

(a) Risk region of FOM, $\mathbb{G}_\beta[X]$. (b) ϵ-risk region of ROM 1, $\widetilde{\mathbb{G}}_1^\beta$. (c) ϵ-risk region of ROM 4, $\widetilde{\mathbb{G}}_4^\beta$.

Figure 3. *Risk regions shown in light yellow for thermal fin problem with two random variables and $\beta = 0.99$. The ε-risk regions for the ROMs are designed to contain the FOM risk region. The smaller the ROM error, the closer the ε-risk regions to the true FOM risk region.*

The reference value $\widehat{\mathrm{CVaR}}_\beta[X]$ is estimated with $m = 5,000$ Monte Carlo samples in Ξ. These samples, Ξ_m, also serve as input for Algorithm 4.1 with corresponding probabilities $p^{(j)} \equiv 1/m$, $j = 1, \ldots, m$. The risk region $\widehat{\mathbb{G}}_\beta[X]$ is shown in light yellow in Figure 3a. The ε-risk regions $\widetilde{\mathbb{G}}_\beta^k$ for the ROMs are designed to contain the FOM risk region, and are closer to the FOM risk region $\widehat{\mathbb{G}}_\beta[X]$ the smaller the ROM error is.

The error in the FOM estimate $\widehat{\mathrm{CVaR}}_\beta[X]$ is quantified by the CI width (2.14). We want a ROM estimate of the same quality. Therefore, we apply Algorithm 4.1 with tolerance

$$\mathrm{TOL} = 10^{-1} \times (\text{CI width}),$$

i.e., 10% of the current estimate of the width of the CI for $\widehat{\mathrm{CVaR}}_\beta[X]$.

Initially, Ξ_m is the set of $5,000$ Monte Carlo samples. The initial ROM basis V_1 is generated with a single $N_1 = 1$ snapshot of the FOM at a randomly selected $\xi \in \Xi_m$. The error function $\widetilde{\epsilon}_1(\xi) = \epsilon_1(\xi)$ evaluated at the samples is plotted in Figure 4a. To construct the next ROM we consider only the samples and the corresponding error values in the risk region $\widetilde{\mathbb{G}}_\beta^1$ plotted in Figure 3b. More generally, in step k we add a snapshot taken at a sample corresponding to the largest value of $\widetilde{\epsilon}_k(\xi)$ in $\widetilde{\mathbb{G}}_\beta^k$. For the newly constructed ROM $\widetilde{X}_{k+1}$ and its error function $\widetilde{\epsilon}_{k+1}$ we check whether $\widetilde{\epsilon}_{k+1}^G < \widetilde{\epsilon}_k^G$. If this is not the case, we add another FOM snapshot to the basis V_{k+1}. In the current example we found that $\widetilde{\epsilon}_{k+1}^G < \widetilde{\epsilon}_k^G$ is always satisfied after the addition of a single FOM snapshot.

In our adaptive framework, reported in Table 2, we only need to evaluate $\widetilde{X}_k$ and $\widetilde{\epsilon}_k$ in the current ϵ-risk region $\Xi_m = \widetilde{\mathbb{G}}_\beta^k$. For example, to build $\widetilde{X}_2$ we consider $8,128$ (and not the full

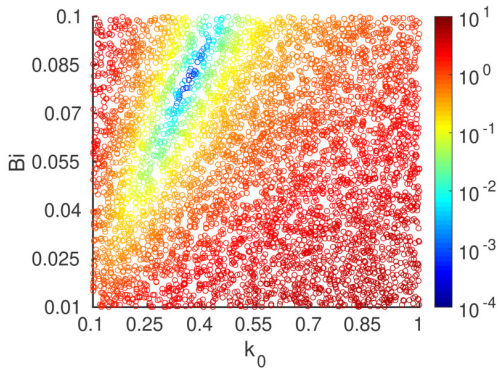

(a) Error of ROM 1, $\widetilde{\epsilon}_1(\xi)$ (both).

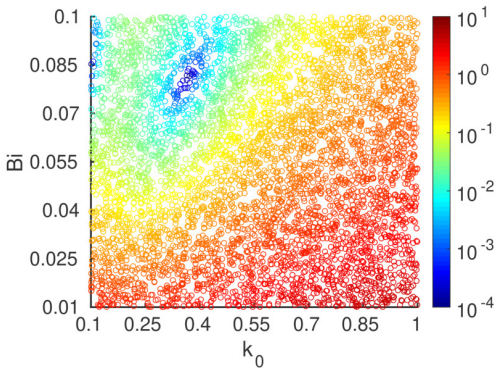

(b) Error of ROM 2, $\widetilde{\epsilon}_2(\xi)$ (both).

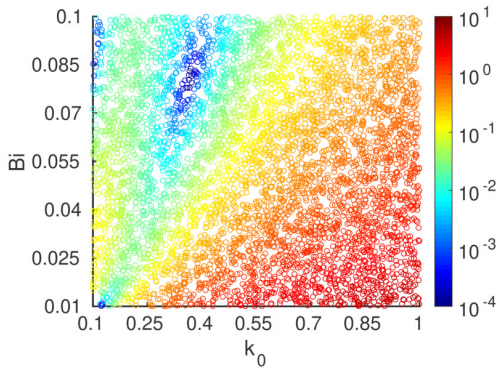

(c) Error of ROM 3, $\widetilde{\epsilon}_3(\xi)$ (adaptive).

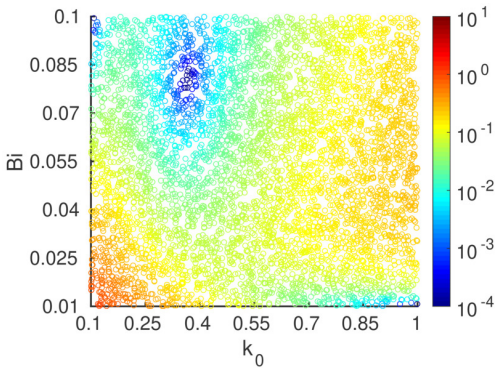

(d) Error of ROM 3, $\epsilon_3(\xi)$ (greedy).

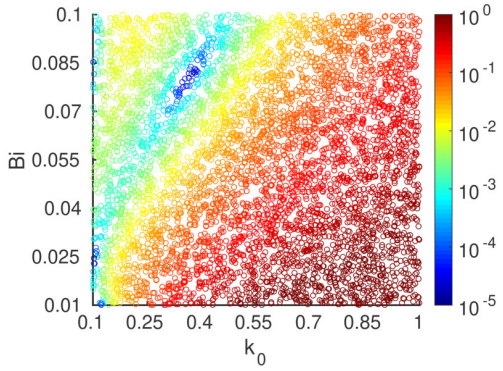

(e) Error of ROM 4, $\widetilde{\epsilon}_4(\xi)$ (adaptive).

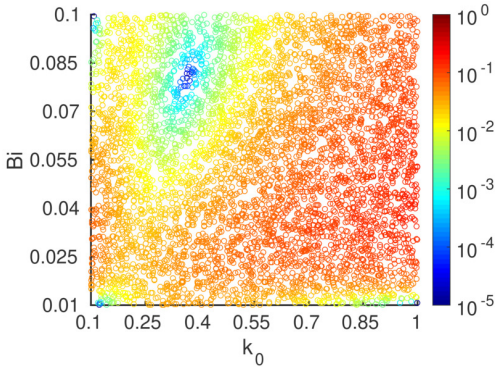

(f) Error of ROM 4, $\epsilon_4(\xi)$ (greedy).

Figure 4. *Error functions $\widetilde{\epsilon}_k(\xi)$ for the ROMs obtained at different steps of Algorithm 4.1 and error functions $\epsilon(\xi)$ obtained with a greedy approach evaluated at samples. Note the different magnitudes on the color bars. Both approaches reduce the error, but error reduction for the adaptive approach is focused more on the risk region.*

Table 2

Results for the adaptive algorithm for the thermal fin problem with two random variables and $\beta = 0.99$. The sizes of the ϵ-risk region $|\widetilde{\mathbb{G}}_\beta^k|$ and of the error bound $\widetilde{\epsilon}_k^G$ decrease monotonically. The current ROM needs to be evaluated at a decreasing number $|\Xi_m|$ of samples, which approaches $1\% = (1-\beta) * 100\%$ of the original number of samples.

	$\widehat{\mathrm{CVaR}}_\beta$	Width CI	Abs error	$\widetilde{\epsilon}_k^G$	$\lvert\widetilde{\mathbb{G}}_\beta^k\rvert$	N_k	$\lvert\Xi_m\rvert$
FOM	12.404	0.437	—	—	—	—	5,000
ROM1	11.381	0.354	1.0238	3.3645	3.60	1	5,000
ROM2	11.486	0.360	0.9185	1.6908	2.44	2	180
ROM3	12.360	0.432	0.0445	0.1461	1.12	3	122
ROM4	12.401	0.438	0.0032	0.0191	1.02	4	56

5,000) samples as candidates for the snapshot selection. These are the only samples that we use in Algorithm 2.1 to evaluate $\mathrm{VaR}_\beta[\widetilde{X}_2]$, $\mathrm{CVaR}_\beta[\widetilde{X}_2]$, and $\widetilde{\mathbb{G}}_\beta^2$. As we continue, the number of samples at which we need to evaluate the current ROM gets closer to $1\% = (1-\beta) * 100\%$ of the size of the initial set Ξ_m.

We contrast the results obtained with adaptive Algorithm 4.1 to those obtained with the greedy approach described in subsection 4.2 and at the end of subsection 5.1. We start with the same initial snapshot; i.e., the initial ROM X_1 is the same. The results for the greedy approach are reported in Table 3. As mentioned before, in each iteration we add a snapshot corresponding to the largest value of $\epsilon_k(\xi)$ at all original samples. Thus all ROMs X_k and error bounds ϵ_k need to be evaluated at all $|\Xi_m| = 5,000$ samples. Although there is no guarantee, in this case the greedy approach also happens to monotonically decrease the size of the ϵ-risk region $\mathbb{G}_\beta^k$ and the error bound ϵ_k^G. However, the error does not decrease as fast as with the adaptive approach.

Table 3

Results for the greedy approach for the thermal fin problem with two random variables and $\beta = 0.99$. Although this cannot be guaranteed, in this case the size of the ϵ-risk region $|\mathbb{G}_\beta^k|$ and the error bound ϵ_k^G happen to decrease monotonically. In each step the current ROM has to be evaluated at all $|\Xi_m| = 5,000$ samples.

	$\widehat{\mathrm{CVaR}}_\beta$	Width CI	Abs error	ϵ_k^G	$\lvert\mathbb{G}_\beta^k\rvert$	N_k	$\lvert\Xi_m\rvert$
FOM	12.404	0.437	—	—	—	—	5,000
ROM1	11.381	0.354	1.0238	3.3645	3.60	1	5,000
ROM2	11.644	0.353	0.7605	1.1809	2.34	2	5,000
ROM3	11.796	0.363	0.6081	1.0494	1.76	3	5,000
ROM4	12.386	0.437	0.0188	0.0680	1.06	4	5,000
ROM5	12.387	0.436	0.0170	0.0666	1.04	5	5,000
ROM6	12.403	0.438	0.0016	0.0057	1.02	6	5,000

The snapshots selected by Algorithm 4.1 and by the greedy approach are shown in Figure 5. Our proposed adaptive algorithm selects FOM snapshots in the current ϵ-risk region, which is close to the original risk region. In contrast, the standard greedy algorithm selects FOM snapshots in the original parameter region. For example, the third snapshot is far outside the risk region; see Figure 5b. In this example, selecting the next snapshot globally in the entire parameter region still gives a good reduction of the ROM error in the ϵ-risk region ϵ_k^G. The greedy algorithm only needs two additional steps to reach the CVaR_β tolerance, compared to

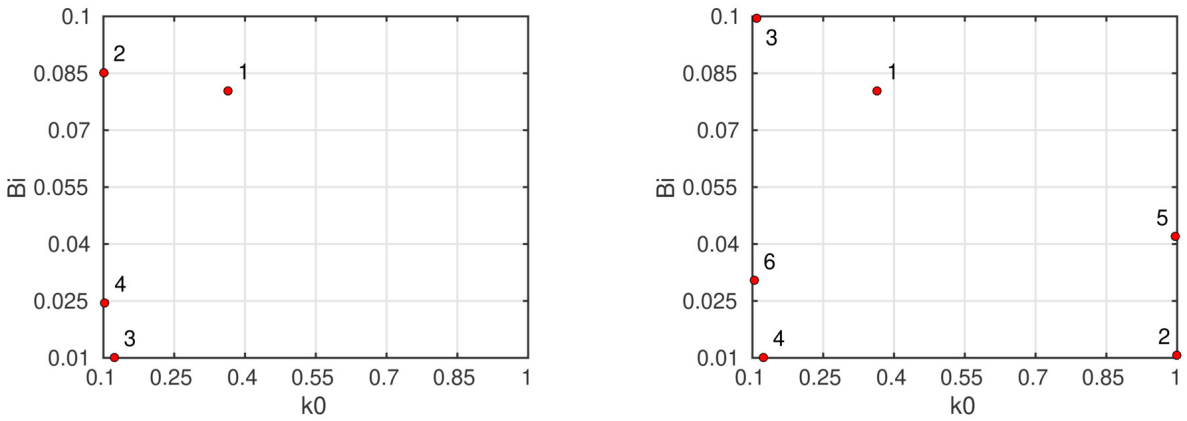

(a) Adaptive selection. (b) Greedy selection.

Figure 5. *Snapshots for ROM construction generated by the adaptive and greedy approaches for the thermal fin problem with two random variables and $\beta = 0.99$. The adaptive approach tends to select snapshots near the risk region.*

our adaptive algorithms. A big difference is in the expense of ROM evaluations; see the last columns of Table 2 and Table 3.

5.4. Results for three random variables. Now we consider the problem with $k_1 = k_2 = k_3 = k_4$ and three random variables $\boldsymbol{\xi} = (k_0, k_1, Bi)$ uniformly distributed in $\Xi = [0.1, 1] \times [0.1, 2] \times [0.01, 0.1]$. Again, we use $5,000$ Monte Carlo samples.

The results for the adaptive approach and the greedy approach are presented in Table 4 and Table 5, respectively. The format of these tables is identical to that of Table 2 and Table 3, respectively.

Table 4

Results for adaptive algorithm for the thermal fin problem with three random variables and $\beta = 0.99$.

	$\widehat{\text{CVaR}}_\beta$	Width CI	Abs error	$\widetilde{\epsilon}_k^G$	$\lvert\widetilde{\mathbb{G}}_\beta^k\rvert$	N_k	$\lvert\Xi_m\rvert$
FOM	10.379	0.405	—	—	—	—	5,000
ROM1	8.292	0.477	2.0870	30.3903	19.88	1	5,000
ROM2	10.008	0.449	0.3718	10.1849	5.46	2	994
ROM3	10.281	0.423	0.0985	3.5377	2.00	3	273
ROM4	10.326	0.413	0.0534	0.2997	1.18	4	100
ROM5	10.357	0.411	0.0225	0.1305	1.08	5	59
ROM6	10.376	0.405	0.0035	0.0429	1.02	6	54
ROM7	10.378	0.405	0.0009	0.0140	1.02	7	51

The snapshots selected by both approaches are shown in Figure 6. We start with a randomly selected initial sample, which is chosen to be the same for both approaches (sample 1 in Figure 6a and Figure 6b). The second sample happens to be the same in both the adaptive and greedy approaches. Due to our suggested ROM modification (3.11a), ROM $\widetilde{X}_2$ in the adaptive case has a smaller bound $\widetilde{\epsilon}_2^G$ than ROM X_2 in the greedy case, ϵ_2^G. The third snapshot is different for the two approaches. However, the third snapshot selected by the

Table 5

Results for the greedy approach for the thermal fin problem with three random variables and $\beta = 0.99$.

	$\widehat{\text{CVaR}}_\beta$	Width CI	Abs error	ϵ_k^G	$\lvert\mathbb{G}_\beta^k\rvert$	N_k	$\lvert\Xi_m\rvert$
FOM	10.379	0.405	—	—	—	—	5,000
ROM1	8.292	0.477	2.0870	30.3903	19.88	1	5,000
ROM2	10.008	0.449	0.3718	11.1808	5.82	2	5,000
ROM3	10.294	0.418	0.0852	3.5377	2.00	3	5,000
ROM4	10.326	0.413	0.0533	0.2997	1.18	4	5,000
ROM5	10.362	0.409	0.0174	0.1792	1.08	5	5,000
ROM6	10.366	0.409	0.0137	0.0806	1.06	6	5,000
ROM7	10.368	0.409	0.0114	0.0815	1.08	7	5,000
ROM8	10.378	0.405	0.0010	0.0087	1.02	8	5,000

greedy approach happens to lie in the ϵ-risk region $\mathbb{G}_\beta^2$ of ROM X_2. (Of course, the third snapshot selected by the adaptive approach will always be chosen in ϵ-risk region $\widetilde{\mathbb{G}}_\beta^2$ of ROM $\widetilde{X}_2$.) In this case, the resulting ROM $\widetilde{X}_3$ in the adaptive case has a larger bound $\widetilde{\epsilon}_3^G$ than the bound ϵ_3^G for ROM X_3 in the greedy case. This can happen, since we compute the next snapshot based on an error bound of the current model, and not based on the error of the new model. In the majority of cases, however, the error bound $\widetilde{\epsilon}_k^G$ for the ROM constructed with the adaptive approach is smaller than the error bound ϵ_k^G for the ROM constructed with the greedy approach.

By construction, the error bound $\widetilde{\epsilon}_k^G$ in the adaptive approach decreases monotonically. This may not be true for the greedy approach. In fact, as can be seen from Table 5, between ROM 6 and ROM 7 we observe an increase in the estimate of ϵ_k^G.

A major strength of our proposed adaptive method is that the ROMs $\widetilde{X}_k$ and their error bounds $\widetilde{\epsilon}_k$ have to be evaluated only at a small number $\lvert\Xi_m\rvert$ of the total samples, whereas in the greedy approach all ROMs and their error bounds have to be evaluated at all 5,000 samples. This leads to significant computational savings for the adaptive ROM construction and CVaR_β estimation.

5.5. Results for six random variables. Finally, we let all six parameters be random, $\boldsymbol{\xi} = (k_0, k_1, k_2, k_3, k_4, Bi)$ uniformly distributed in $\Xi = [0.1, 1] \times [0.1, 2]^4 \times [0.01, 0.1]$. Again, we use $5,000$ Monte Carlo samples.

Results for $\beta = 0.99$ are presented in Table 6 and Table 7. We omit some of the rows in both tables in the interest of saving space. In the greedy case we once more observe an increase in ϵ_k^G between subsequent iterations (see rows corresponding to ROM 10 and ROM 11 in Table 7).

6. Conclusions. We have presented an extension of our recent work [3] that systematically and efficiently improves a ROM to obtain a better ROM-based CVaR estimate. A key ingredient to make efficient use of ROM is the structure of CVaR, which only depends on samples in a small but a priori unknown region of the parameter space. ROMs are used to approximate this region, and new ROMs only need to be better than the previous ROM in these approximate regions. However, to guarantee that this approach monotonically improves the CVaR estimate, we had to introduce a new way to combine previously constructed ROMs

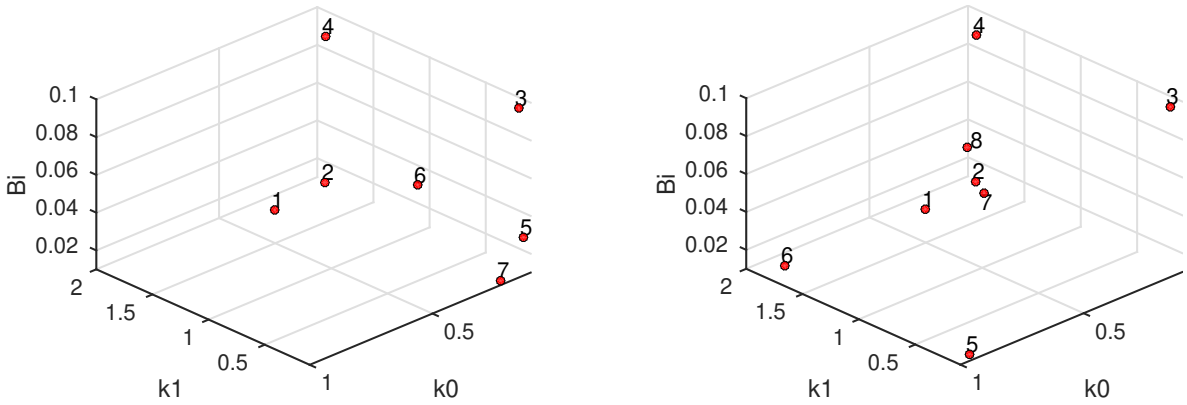

(a) Adaptive selection. (b) Greedy selection.

Figure 6. *Snapshots for ROM construction for the thermal fin problem with three random variables and* $\beta = 0.99$.

Table 6

Results for the adaptive algorithm for the thermal fin problem with six random variables and $\beta = 0.99$.

	$\widehat{\text{CVaR}}_\beta$	Width CI	Abs error	$\widetilde{\epsilon}_k^G$	$\lvert\widetilde{\mathbb{G}}_\beta^k\rvert$	N_k	$\lvert\Xi_m\rvert$
FOM	10.435	0.421	—	—	—	—	5,000
ROM1	9.386	0.388	1.0492	14.5163	15.08	1	5,000
ROM2	9.872	0.449	0.5630	11.6548	7.98	2	754
ROM3	10.201	0.403	0.2335	2.6354	2.42	3	399
ROM4	10.310	0.408	0.1249	0.7235	1.42	4	121
ROM5	10.363	0.416	0.0717	0.3908	1.34	5	71
ROM6	10.424	0.420	0.0110	0.2941	1.14	6	67
ROM7	10.430	0.421	0.0044	0.1314	1.02	7	57
ROM8	10.432	0.421	0.0026	0.0557	1.02	8	51
ROM9	10.433	0.421	0.0019	0.0285	1.02	9	51

Table 7

Results for the greedy procedure for the thermal fin problem with six random variables and $\beta = 0.99$.

	$\widehat{\text{CVaR}}_\beta$	Width CI	Abs error	ϵ_k^G	$\lvert\mathbb{G}_\beta^k\rvert$	N_k	$\lvert\Xi_m\rvert$
FOM	10.435	0.421	—	—	—	—	5,000
ROM1	9.386	0.388	1.0492	14.5163	15.08	1	5,000
ROM2	9.872	0.449	0.5623	12.4641	8.42	2	5,000
ROM3	10.206	0.401	0.2292	2.6354	2.48	3	5,000
ROM4	10.271	0.403	0.1634	1.9756	1.88	4	5,000
ROM5	10.349	0.413	0.0854	1.5134	1.68	5	5,000
ROM6	10.385	0.419	0.0496	0.8382	1.34	6	5,000
ROM7	10.398	0.421	0.0369	0.8645	1.32	7	5,000
ROM8	10.420	0.423	0.0144	0.2083	1.14	8	5,000
ROM9	10.421	0.423	0.0136	0.1854	1.12	9	5,000
ROM10	10.430	0.422	0.0052	0.0683	1.08	10	5,000
ROM11	10.430	0.422	0.0046	0.0680	1.08	11	5,000
ROM12	10.430	0.422	0.0043	0.0616	1.08	12	5,000
ROM13	10.431	0.422	0.0041	0.0655	1.06	13	5,000
ROM14	10.432	0.422	0.0032	0.0556	1.08	14	5,000
ROM15	10.433	0.422	0.0017	0.0266	1.06	15	5,000

into new adaptive ROMs. We have provided error estimates and demonstrated the benefits of our approach on a numerical example for the CVaR estimation of a QoI governed by an elliptic differential equation.

Our approach requires the construction of ROMs with error bounds. In many examples it is difficult to find error bounds, and instead one may only have asymptotic bounds or estimates. Extension of our approach to such cases would expand the rigorous and systematic use of ROMs for CVaR estimation.

Acknowledgments. The computational results in section 5 used MATLAB code for the implementation of the thermal fin problem written by Prof. Martin Grepl and Elizabeth Qian. The authors would like to thank Prof. Martin Grepl and Elizabeth Qian for making the code available. The authors also thank the referees for their comments, which have led to improvements in the presentation.

REFERENCES

[1] P. Chen and A. Quarteroni, *Accurate and efficient evaluation of failure probability for partial different equations with random input data*, Comput. Methods Appl. Mech. Engrg., 267 (2013), pp. 233–260, https://doi.org/10.1016/j.cma.2013.08.016.

[2] B. Haasdonk, *Chapter* 2: *Reduced basis methods for parametrized PDEs—a tutorial introduction for stationary and instationary problems*, in Model Reduction and Approximation: Theory and Algorithms, P. Benner, A. Cohen, M. Ohlberger, and K. Willcox, eds., Computational Science and Engineering 15, SIAM, Philadelphia, 2017, pp. 65–136, https://doi.org/10.1137/1.9781611974829.ch2.

[3] M. Heinkenschloss, B. Kramer, T. Takhtaganov, and K. Willcox, *Conditional-value-at-risk estimation via reduced-order models*, SIAM/ASA J. Uncertain. Quantif., 6 (2018), pp. 1395–1423, https://doi.org/10.1137/17M1160069.

[4] J. S. Hesthaven, G. Rozza, and B. Stamm, *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, Springer Briefs in Mathematics, Springer, New York, 2015, https://doi.org/10.1007/978-3-319-22470-1.

[5] L. J. Hong, Z. Hu, and G. Liu, *Monte Carlo methods for value-at-risk and conditional value-at-risk: A review*, ACM Trans. Model. Comput. Simul., 24 (2014), 22, https://doi.org/10.1145/2661631.

[6] D. P. Kouri, M. Heinkenschloss, D. Ridzal, and B. G. van Bloemen Waanders, *Inexact objective function evaluations in a trust-region algorithm for PDE-constrained optimization under uncertainty*, SIAM J. Sci. Comput., 36 (2014), pp. A3011–A3029, https://doi.org/10.1137/140955665.

[7] D. P. Kouri and T. M. Surowiec, *Risk-averse PDE-constrained optimization using the conditional value-at-risk*, SIAM J. Optim., 26 (2016), pp. 365–396, https://doi.org/10.1137/140954556.

[8] P. Krokhmal, J. Palmquist, and S. Uryasev, *Portfolio optimization with conditional value-at-risk objective and constraints*, J. Risk, 4 (2002), pp. 11–27, https://doi.org/10.21314/JOR.2002.057.

[9] L. Machiels, Y. Maday, and A. T. Patera, *Output bounds for reduced-order approximations of elliptic partial differential equations*, Comput. Methods Appl. Mech. Engrg., 190 (2001), pp. 3413–3426, https://doi.org/10.1016/S0045-7825(00)00275-9.

[10] L. Machiels, Y. Maday, A. T. Patera, and D. V. Rovas, *A blackbox reduced-basis output bound method for shape optimization*, in Proceedings of the 12th International Conference on Domain Decomposition Methods, Chiba, Japan, T. Chan, T. Kako, and H. K. O. Pironneau, eds., DDM.org, 2001, pp. 429–436.

[11] R. Mansini, W. Ogryczak, and M. G. Speranza, *Conditional value at risk and related linear programming models for portfolio optimization*, Ann. Oper. Res., 152 (2007), pp. 227–256, https://doi.org/10.1007/s10479-006-0142-4.

[12] M. Norton, V. Khokhlov, and S. Uryasev, *Calculating CVaR and bPOE for common probability distributions with application to portfolio optimization and density estimation*, Ann. Oper. Res., 2019, pp. 1–35.

[13] B. Peherstorfer, K. Willcox, and M. D. Gunzburger, *Survey of multifidelity methods in uncertainty propagation, inference, and optimization*, SIAM Rev., 60 (2018), pp. 550–591, https://doi.org/10.1137/16M1082469.

[14] C. Prud'homme, D. V. Rovas, K. Veroy, L. Machiels, Y. Maday, A. T. Patera, and G. Turinici, *Reliable real-time solution of parametrized partial differential equations: Reduced-basis output bound methods*, J. Fluids Engrg., 124 (2002), pp. 70–80, https://doi.org/10.1115/1.1448332.

[15] A. Quarteroni, A. Manzoni, and F. Negri, *Reduced Basis Methods for Partial Differential Equations. An Introduction*, Unitext 92, Springer, Cham, 2016, https://doi.org/10.1007/978-3-319-15431-2.

[16] R. T. Rockafellar and J. O. Royset, *Engineering decisions under risk averseness*, ASCE-ASME J. Risk Uncertain. Engrg. Syst. A Civil Engrg., 1 (2015), p. 04015003, https://doi.org/10.1061/AJRUA6.0000816.

[17] R. T. Rockafellar and S. Uryasev, *Optimization of conditional value-at-risk*, J. Risk, 2 (2000), pp. 21–41, https://doi.org/10.21314/JOR.2000.038.

[18] R. T. Rockafellar and S. Uryasev, *Conditional value-at-risk for general loss distributions*, J. Banking Finance, 26 (2002), pp. 1443–1471, https://doi.org/10.1016/S0378-4266(02)00271-6.

[19] J. O. Royset, L. Bonfiglio, G. Vernengo, and S. Brizzolara, *Risk-adaptive set-based design and applications to shaping a hydrofoil*, J. Mech. Des., 139 (2017), 101403, https://doi.org/10.1115/1.4037623.

[20] T. Takhtaganov, *Efficient Estimation of Coherent Risk Measures for Risk-Averse Optimization Problems Governed by Partial Differential Equations with Random Inputs*, Ph.D. thesis, Department of Computational and Applied Mathematics, Rice University, Houston, TX, 2017.

[21] H. Yang and M. Gunzburger, *Algorithms and analyses for stochastic optimization for turbofan noise reduction using parallel reduced-order modeling*, Comput. Methods Appl. Mech. Engrg., 319 (2017), pp. 217–239, https://doi.org/10.1016/j.cma.2017.02.030.

[22] M. J. Zahr, K. T. Carlberg, and D. P. Kouri, *An efficient, globally convergent method for optimization under uncertainty using adaptive model reduction and sparse grids*, SIAM/ASA J. Uncertain. Quantif., 7 (2019), pp. 877–912, https://doi.org/10.1137/18M1220996.

[23] W. Zhang, H. Rahimian, and G. Bayraksan, *Decomposition algorithms for risk-averse multistage stochastic programs with application to water allocation under uncertainty*, INFORMS J. Comput., 28 (2016), pp. 385–404, https://doi.org/10.1287/ijoc.2015.0684.

[24] Z. Zou, D. P. Kouri, and W. Aquino, *An adaptive sampling approach for solving PDEs with uncertain inputs and evaluating risk*, in Proceedings of the 19th AIAA Non-Deterministic Approaches Conference, AIAA SciTech Forum (AIAA 2017-1325), 2017, https://doi.org/10.2514/6.2017-1325.

[25] Z. Zou, D. P. Kouri, and W. Aquino, *A locally adapted reduced basis method for solving risk-averse PDE-constrained optimization problems*, in Proceedings of the 2018 AIAA Non-Deterministic Approaches Conference (Kissimmee, FL), AIAA SciTech Forum (AIAA 2018-2174), 2018, https://doi.org/10.2514/6.2018-2174.

[26] Z. Zou, D. P. Kouri, and W. Aquino, *An adaptive local reduced basis method for solving PDEs with uncertain inputs and evaluating risk*, Comput. Methods Appl. Mech. Engrg., 345 (2019), pp. 302–322, https://doi.org/10.1016/j.cma.2018.10.028.