Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

BLACK-BOX SIMULATION-OPTIMIZATION WITH QUANTILE CONSTRAINTS: AN INVENTORY CASE STUDY

Ebru Angün[1] and Jack Kleijnen[2]

[1]Dept. of Industrial Eng., Galatasaray University, Istanbul, TURKEY
[2]Dept. of Mgmt., Tilburg University, Tilburg, NETHERLANDS

ABSTRACT

We apply a recent variant of "efficient global optimization" (EGO). EGO is closely related to Bayesian optimization (BO): both EGO and BO treat the simulation model as a black box, and use a Kriging metamodel or Gaussian process. The recent variant of EGO combines (i) EGO for unconstrained optimization, and (ii) the Karush-Kuhn-Tucker optimality conditions for constrained optimization. EGO sequentially searches for the global optimum. We apply this variant and a benchmark EGO variant to an (s,S) inventory model. We aim to minimize the mean inventory costs—excluding disservice costs—while satisfying a prespecified threshold for the 90%-quantile of the disservice level. Our numerical results imply that the mean inventory costs increase by 2.5% if management is risk-averse instead of risk-neutral—using the mean value. Comparing the two EGO variants shows that these variants do not give significantly different results, for this application.

1 INTRODUCTION

In this paper we apply Kleijnen et al. (2023)'s KKT-EGO algorithm that combines the popular *efficient global optimization* (EGO) algorithm and the *Karush-Kuhn-Tucker* (KKT) conditions. Angün and Kleijnen (2023) extends the KKT-EGO algorithm to cope with stochastic simulation with heterogenous aleatory noise. The latter publication assumes *risk-neutral* decision makers, who use expected values of the multiple simulation outputs—be it the goal output or the constrained outputs. In the current study, we assume *risk-averse* decision makers. We model this risk attitude through quantiles that characterize the tail of the *probability density function* (PDF) of the simulated outputs. Quantiles are also used to quantify value-at-risk in financial portfolio management.

EGO is closely related to *Bayesian optimization* (BO) and *machine learning* (ML), especially active learning. The KKT conditions are well-known (first-order necessary) optimality conditions in mathematical programming, but these conditions are not used in other EGO algorithms. EGO, BO, and ML use a *Kriging* or *Gaussian process* (GP) metamodel (approximation, emulator, surrogate) for the *input/output* (I/O) function of the underlying simulation model. EGO algorithms for optimization with output constraints are not guaranteed to converge to the global optima. We compare KKT-EGO with a benchmark EGO algorithm related to Carpio et al. (2018) and Gardner et al. (2014). EGO is a rapidly evolving field; see the many references in Baker et al. (2022), Garnett (2023), Wang and Yang (2023), Wang et al. (2023).

We apply Angün and Kleijnen (2023)'s KKT-EGO to the (s,S) inventory model that was originally defined in Bashyam and Fu (1998). We let w_1 denote *disservice* level per period; i.e., the percentage of total accumulated demand that is not satisfied from stock on-hand. The 90%-quantile of w_1—denoted by $q_{0.90}(w_1)$—implies that the probability is at least 90% that this percentage does not exceed $q_{0.90}(w_1)$. We assume that the managers are risk-neutral regarding the relevant inventory cost—denoted by w_0—which excludes the hard-to-quantify cost of out-of-stock, so we use a service-level constraint. Altogether, our goal is to minimize $E(w_0)$ (expected value or mean of w_0), while satisfying the constraint $q_{0.90}(w_1) \leq c_1$;

e.g., $c_1 = 0.10$. Our formulation with quantile constraints also covers probabilistic or chance constraints (see (4)). Obviously, this model has *aleatory* uncertainty, but no *epistemic* uncertainty; these two types of uncertainty are detailed in Kleijnen (2015). We shall present more details of this model, in Section 4.

Obviously, risk-aversion implies that $E(w_0)$ increases when satisfying (say) $q_{0.90}(w_1) \leq 0.10$ instead of $E(w_1) \leq 0.10$. However, it is not obvious *how much* $E(w_0)$ increases; i.e., how much s and S change, and how much they affect $E(w_0)$. We shall see that—in our example—$E(w_0)$ increases by 2.5% if management is risk-averse (modeled by $q_{0.90}(w_1)$) instead of risk-neutral (modeled by $E(w_1)$).

We also compare numerical results of KKT-EGO and the benchmark, regarding efficiency (convergence speed) and effectiveness (closeness to the true optimal solution). This comparison suggests that these variants do not give significantly different results, for this application (however, Angün and Kleijnen (2023) compares KKT-EGO and three competing algorithms including Carpio et al. (2018)'s, in three examples, and finds that KKT-EGO is relatively efficient and effective).

We organize the rest of this paper as follows. Section 2 reviews recent literature. Section 3 details the mathematical formulation of our inventory problem, and its solution via the KKT-EGO algorithm. Section 4 gives numerical results for the simulation optimization of our (s,S) model with a constraint for either $E(w_1)$ or $q_{0.90}(w_1)$, applying KKT-EGO and the benchmark. Section 5 summarizes our conclusions and possible future research topics.

2 LITERATURE REVIEW

Quantiles are frequently used as risk measures for different types of problems such as our simulation optimization problems and input uncertainty (or epistemic) problems discussed in Song et al. (2024).

Chang and Cuckler (2022) develops a simulation optimization method that minimizes a specific quantile while satisfying an upper bound for the total cost which is assumed analytically available. That publication applies its method to solve vehicle fleet sizing for an automated material handling system in a wafer fab in Taiwan, minimizing a specific quantile of transport time of wafer lots; obviously, this fleet sizing implies integer variables instead of continuous variables (Kriging assumes continuos variables). The method does not use EGO. Chang and Lin (2023) includes a chance constraint in a simulation-optimization method for system reliability via the allocation of redundant system components (obviously, the number of components is an integer); that method does not use EGO.

Hu et al. (2024) considers a constrained optimization of a quantile with input constraints but without output constraints; that publication does not estimate the quantiles and their derivatives through Kriging. Wang et al. (2022) discusses some algorithms for unconstrained optimization of loss functions. These algorithms should not require strict properties for the loss functions such as convexity, and should allow computationally expensive simulations. Therefore that article uses metamodels; namely Kriging models. Baker et al. (2022, Section 3.3.1) briefly discusses output quantiles, in a review of Kriging for analyzing stochastic simulation. Kroetz et al. (2020) discusses ordinary Kriging (OK) and EGO in the context of reliability analysis with risk.

3 PROBLEM FORMULATION

Section 3.1 discusses our quantile estimation. Focusing on these quantiles, Section 3.2 summarizes constrained optimization. Section 3.3 summarizes Kriging. Section 3.4 summarizes the KKT-EGO algorithm.

3.1 Quantile Estimation

Following Bashyam and Fu (1998, eq. 3), we define Y_p as the amount of demand not satisfied from on-hand stock in period p; we estimate the *disservice level* through the running *average* of P' periods:

$$w_{1;P'} = \frac{\sum_{p=1}^{P'} Y_p}{\sum_{p=1}^{P'} D_p} \text{ with } P' = 1,...,P = 30,000. \tag{1}$$

We let $F(w_1)$ denote the marginal *cumulative density function* (CDF) of $w_{1;P'}$; serial correlations among $w_{1;P'}$ do affect the joint CDF, but not the marginal CDF in the *steady state*. The standard definition of the $(1-\alpha)$-quantile of w_1 is then:

$$q_{1-\alpha}(w_1) = F^{-1}(1-\alpha) = inf\ [w_1 : F(w_1) \geq 1-\alpha]; \tag{2}$$

see Alexopoulos et al. (2019, p. 1162) and Lolos et al. (2023), which give more references.

Whereas Bashyam and Fu (1998) uses $E(w_{1;P'})$ with $P' \uparrow \infty$ and the estimator $w_{1;P'}$ with $P' = P = 30,000$, we use $q_{1-\alpha}(w_1)$ and the estimator $\widehat{q}_{1-\alpha}(w_1)$—computed as follows. We sort the P' observations $w_{1;P'}$ in ascending order, so $w_{1;(1)} \leq w_{1;(2)} \leq ... \leq w_{1;(P'-1)} \leq w_{1;(P')}$; i.e., we use the *order statistic* $w_{1;(\lceil (1-\alpha)P' \rceil)}$ (where $\lceil x \rceil$ denotes the ceiling function of the real number x). Then, combining $P' = 30,000$ and $\alpha = 0.10$ gives $\widehat{q}_{1-\alpha}(w_1) = w_{1;(27000)}$. Alexopoulos et al. (2019) also uses this quantile estimator based on order statistics, but focuses on *confidence intervals* (CIs) for quantiles based on a single replication, whereas we use multiple replications—as we shall see in Section 3.3. Chen and Kim (2016) discusses alternative estimators of quantiles and the related value-at-risk and conditional value-at-risk.

3.2 Optimization with a Constrained Quantile

We wish to solve the following constrained optimization problem for a specific (s,S) simulation model with the goal output $E(w_0)$ and the constrained output $q_{0.90}(w_1)$. We let c_1 denote the upper threshold for $q_{0.90}(w_1)$, and $\mathbf{x} = (x_1, x_2)'$ denote the $k = 2$ control or decision variables $x_1 = s$ and $x_2 = Q$ with $Q = S - s$. The problem has the *box constraints* $\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}$ with $\mathbf{l} = (l_1, l_2)'$ and $\mathbf{u} = (u_1, u_2)'$, which determine the *experimental area*. Altogether, this gives

$$\min_{\mathbf{x}}\ \{E[w_0(\mathbf{x})]:\ q_{1-\alpha}[w_1(\mathbf{x})] \leq c_1,\ \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}\}. \tag{3}$$

This problem formulation also covers *chance constraints*:

$$\text{Prob}[w_1(\mathbf{x}) \leq c_1] \geq 1-\alpha \iff q_{1-\alpha}[w_1(\mathbf{x})] \leq c_1. \tag{4}$$

However, a practical simulation model is a *black box*; i.e., $E[w_0(\mathbf{x})]$ and $q_{1-\alpha}[w_1(\mathbf{x})]$ in (3) are unknown I/O functions. Like many authors, we estimate these functions through Kriging metamodels, which give explicit approximations of the implicit I/O functions defined by the underlying simulation model. For our problem defined in (3) we use specific Kriging models, as follows.

3.3 Stochastic Kriging

Like most publications on EGO, we apply *univariate* Kriging instead of multivariate Kriging or co-Kriging; i.e., we ignore the correlations between different types of simulation outputs (in our case study, the correlation between $w_0(\mathbf{x})$ and $w_1(\mathbf{x})$; obviously, $\widehat{q}_{(1-\alpha)}[w_1(\mathbf{x})]$ depends on $w_1(\mathbf{x})$). Because the (s,S) simulation is stochastic, we apply SK. We use the formulas for the SK predictor $\widehat{y}$ and its variance $s^2(\widehat{y})$ that are derived in Ankenman et al. (2010). These formulas assume $m_i > 1$ replications for the old simulated input combination $\mathbf{x}_i$ ($i = 1, ..., n$ where n is initialized via (10)). By definition, these replications have a common fixed simulation-runlength $P = 30,000$, and they use non-overlapping streams of *pseudorandom numbers* (PRN), so they give m_i *identically and independently distributed* (IID) outputs $w_{0;r}(\mathbf{x}_i)$ and $\widehat{q}_{(1-\alpha);r}[w_1(\mathbf{x}_i)]$ with $r = 1, ..., m_i$. We find it convenient to denote these two outputs by $o_{0;r} = w_{0;r}$ and $o_{1;r} = \widehat{q}_{1-\alpha;r}(w_1)$. Then, these m_i replications give the averages

$$\overline{o}_h(\mathbf{x}_i) = \frac{\sum_{r=1}^{m_i} o_{h;r}(\mathbf{x}_i)}{m_i} \text{ with } h = 0, 1,\ i = 1, ..., n \tag{5}$$

and the unbiased estimators of the heterogeneous variances

$$s^2[o_h(\mathbf{x}_i)] = \frac{\sum_{r=1}^{m_i}[o_{h;r}(\mathbf{x}_i) - \overline{o}_h(\mathbf{x}_i)]^2}{(m_i - 1)};\ s^2[\overline{o}_h(\mathbf{x}_i)] = \frac{s^2[o(\mathbf{x}_i)]}{m_i}. \tag{6}$$

SK assumes that the so-called *intrinsic noise*—caused by the PRN—has a Gaussian (or normal) distribution with zero mean and heterogeneous variances. We apply SK to $\overline{w}_0(\mathbf{x}_i)$ and $\overline{\widehat{q}}_{1-\alpha}[w_1(\mathbf{x}_i)]$, which are averages of m_i IID observations $w_{0;r}(\mathbf{x}_i)$ and $\widehat{q}_{(1-\alpha);r}[w_1(\mathbf{x}_i)]$ so we assume that the central limit theorem applies; i.e., we assume that $\overline{w}_0(\mathbf{x}_i)$ and $\overline{\widehat{q}}_{1-\alpha}[w_1(\mathbf{x}_i)]$ are normally distributed. We denote the intrinsic noise by $e_{h;r}(\mathbf{x}_i)$—or briefly $e_{i;h;r}$—and its average by $\overline{e}_h(\mathbf{x}_i)$—or $\overline{e}_{i;h}$. Let $\mathbf{\Sigma}_{\overline{e};h}$ denote the covariance matrix of $\overline{e}_{i;h}$. Like most authors on SK, we do not apply common random numbers so $\mathbf{\Sigma}_{\overline{e};h}$ is an $n \times n$ diagonal matrix (say) $\widehat{\mathbf{\Sigma}}_{\overline{e};h}$ with the main-diagonal elements $s^2[\overline{o}_h(\mathbf{x}_i)]$ defined in (6).

Besides intrinsic noise, SK considers *extrinsic noise* which models how the SK outputs—of a given type h—at two points $\mathbf{x}$ and $\mathbf{x}'$ are more correlated, the "closer" $\mathbf{x}$ and $\mathbf{x}'$ are. This assumption is realistic if the simulation output functions are smooth. This noise implies that a SK model is a stationary GP with a covariance matrix that is determined by the correlation function or *kernel*. We select the *anisotropic Gaussian* (or squared exponential) kernel with the so–called *length-scale* hyperparameters $\boldsymbol{\theta}_h = (\theta_{h;1}, ..., \theta_{h;k})'$ where $\theta_{h;j} \geq 0$:

$$\rho(\boldsymbol{\theta}_h, \mathbf{x}, \mathbf{x}') = \prod_{j=1}^{k} \exp[-\theta_{h;j}(x_j - x'_j)^2], \tag{7}$$

which implies that the correlation between $o_h(\mathbf{x})$ and $o_h(\mathbf{x}')$ decreases exponentially, as $\mathbf{x}$ and $\mathbf{x}'$ are farther apart. This kernel is most popular in simulation, but we may replace it by another kernel; e.g., the quite popular *Matérn* class of kernels.

Like most authors on SK, we assume a constant mean (say) $\mu_h = E(y_h)$. The usual symbol for $Var(y_h)$ is τ_h^2, so the correlation matrix $\mathbf{R}_h = (\rho_{i;i';h})$ equals $\tau_h^{-2}\mathbf{\Sigma}_{M;h}$ with the covariance matrix $\mathbf{\Sigma}_{M;h} = (Cov(y_{h;i}, y_{h;i'}))$ where $y_{h;i} = y_h(\mathbf{x}_i)$ and $i' = 1, ..., n$. The correlations between the Kriging outputs of type h at the new point $\mathbf{x}_*$ and the n old points are $\boldsymbol{\rho}_h(\mathbf{x}_*) = \tau_h^{-2}\boldsymbol{\sigma}_{M;h}(\mathbf{x}_*)$ with the n-dimensional covariance vector $\boldsymbol{\sigma}_{M;h}(\mathbf{x}_*) = (\sigma_{h;*}) = (Cov(y_{h;*}, y_{h;i}))$ with $y_{h;*} = y_h(\mathbf{x}_*)$. Finally, $\mathbf{1}_n$ denotes the n-dimensional vector with all elements equal to 1. We use the *maximum likelihood estimators* (MLEs) for the hyperparameters (except $\widehat{\mathbf{\Sigma}}_{\overline{e};h}$), and let the symbol $\widehat{}$ denote these MLEs. Using these symbols, the SK predictors are

$$\widehat{y}_h(\mathbf{x}_*) = \widehat{\mu}_h + \widehat{\boldsymbol{\sigma}}_{M;h}(\mathbf{x}_*)'(\widehat{\mathbf{\Sigma}}_{M;h} + \widehat{\mathbf{\Sigma}}_{\overline{e};h})^{-1}(\overline{\mathbf{o}}_h - \widehat{\mu}_h \mathbf{1}_n), \tag{8}$$

and their estimated standard deviations are

$$\begin{aligned} &s^2[\widehat{y}_h(\mathbf{x}_*)] = \widehat{\tau}_h^2 - \widehat{\tau}_h^4 \boldsymbol{\rho}'_h(\mathbf{x}_*)[\widehat{\tau}_h^2 \mathbf{R}_h + \widehat{\mathbf{\Sigma}}_{\overline{e};h}]^{-1}\boldsymbol{\rho}_h(\mathbf{x}_*) + \widehat{\delta}_h^2[\mathbf{1}'_n(\widehat{\tau}_h^2 \mathbf{R}_h + \widehat{\mathbf{\Sigma}}_{\overline{e};h})^{-1}\mathbf{1}_n]^{-1} \\ &\text{with } \widehat{\delta}_h = 1 - \mathbf{1}'_n(\widehat{\tau}_h^2 \mathbf{R}_h + \widehat{\mathbf{\Sigma}}_{\overline{e};h})^{-1}\boldsymbol{\rho}_h(\mathbf{x}_*)\widehat{\tau}_h^2. \end{aligned} \tag{9}$$

We compute (8) and (9) using MATLAB functions for SK (developed by Gonzalez—see Acknowledgment). To compute $\widehat{\boldsymbol{\theta}}_h$ in (7), we must select a search area; we select the lower bound 0.001 and Gonzalez's upper bound $3^{1/2}$, for each $\theta_{h;j}$.

We let n_0 denote the *initial* or *pilot* number of simulated points. To select this n_0, we follow Angün and Kleijnen (2023):

$$n_0 = (k+1)(k+2)/2 \text{ if } k \leq 6; \text{ else } n_0 = 5k. \tag{10}$$

Our (s, S) example has $k = 2$, so (10) gives $n_0 = 6$. Like most authors on Kriging, we use *Latin hypercube sampling* (LHS) to select a space-filling design with $\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}$ (see (3)). We use LHS with *midpoints*, so the n_0 points projected onto the k axes are equidistant, which gives better $\widehat{\boldsymbol{\theta}}_h$. LHS with midpoints is an option in MATLAB's function *lhsdesign*.

Furthermore, we let m_0 denote the *initial number of replications* for point i with $i = 1, ..., n_0$. We start with $m_0 = 2$, which is the smallest value of m_0 that enables the computation of $s(o_h)$ (defined in (6)). Because our (s, S) model has a large simulation-runlength P, we expect a small $s(o_h)$ so $m_0 = 2$ might suffice.

We *validate* the SK models that are estimated from the initial n_0 points with m_i replications per point. We apply the popular *leave-one-out cross-validation* (LOO-CV), which means that we delete $\mathbf{x}_i$ (with $i = 1,...,n_0$) and its $\overline{o}_h(\mathbf{x}_i)$ and $s^2[\overline{o}_h(\mathbf{x}_i)]$, and use the remaining simulation I/O data to compute $\widehat{y}_{-i;h}$ and $s^2(\widehat{y}_{-i;h})$, etc. We select a prespecified error rate (say) α_{CV}; our numerical experiment uses $\alpha_{\text{CV}} = 0.20$. We use the *Bonferroni* inequality, accounting for two outputs (namely, $\overline{w}_0(\mathbf{x}_i)$ and $\overline{\widehat{q}}_{1-\alpha}[w_1(\mathbf{x}_i)]$) at n_0 points, so a two-sided CI uses $\alpha_{\text{CV}}/(2 \times 2 \times n_0)$. So, we use the $1-(\alpha_{\text{CV}}/(4n_0))$-quantile $z_{1-[\alpha_{\text{CV}}/(4n_0)]}$ of the standard normal distribution. Altogether, we reject the SK models if

$$\max_{i;h} \left[\frac{|\overline{o}_{i;h} - \widehat{y}_{-i;h}|}{\sqrt{s^2(\overline{o}_{i;h}) + s^2(\widehat{y}_{-}i;h)}} \right] > z_{1-[\alpha_{\text{CV}}/(4n_0)]}. \tag{11}$$

If (11) holds, then we add a single replication to those points that require additional replications according to the following procedure.

We apply Law (2005, p. 505)'s sequential procedure for estimating $E(o_h)$ with prespecified relative error γ (with $0 < \gamma < 1$) and prespecified confidence level $1-\alpha_m$; our experiment uses $\gamma = 0.10$ and $\alpha_m = 0.10$:

$$\widehat{m}_h = \min \left[r \geq m : \frac{t_{r-1;1-\alpha_m/2} s_h(r)/\sqrt{r}}{|\overline{o}_h(r)|} \leq \frac{\gamma}{1+\gamma} \right]. \tag{12}$$

Obviously, $s_h/\overline{o}_h$ in (12) is the coefficient of variation of output h. If $\overline{o}_h(r) \approx 0$, then (12) gives $\widehat{m}_h$ so high that we replace $\gamma/(1+\gamma)$ by the absolute error β and $\overline{o}_h(r)$ by the constant 1. Indeed, our (s,S) model tends to give $\overline{\widehat{q}}_{1-\alpha}[w_1(\mathbf{x})] \approx 0$ if $\mathbf{x}$ lies far away inside the feasible area so $w_{1;r}(\mathbf{x}) \approx 0$; we then specify $\beta = 0.01$.

Because the simulation model gives multiple types of output o_h, we apply (12) such that this equation holds for all these outputs, at $\mathbf{x}_i$. So, the desired number of replications at $\mathbf{x}_i$ is

$$\widehat{m}(\mathbf{x}_i) = \max_h \ [\widehat{m}_h(\mathbf{x}_i)]. \tag{13}$$

Our replication rule stops adding replications for $\mathbf{x}_i$, as soon as $m(\mathbf{x}_i) \geq \widehat{m}(\mathbf{x}_i)$ holds.

After the initial stage, we also apply this replication rule to the point that a next iteration selects as the new point to be simulated; see the next section.

3.4 The KKT-EGO Algorithm

To solve our problem defined in (3), we apply Angün and Kleijnen (2023)'s KKT-EGO algorithm. EGO is a popular method that was originally developed for unconstrained optimization in deterministic simulation; see Jones et al. (1998). Because EGO methods (and other simulation-optimization methods) treat the simulation model as a black box, it is unknown whether the problem has *multiple* optima. Therefore, EGO balances *global* search and *local* search—or exploration of the whole experimental area versus exploitation of a local promising area.

EGO is sequential; i.e., it selects a new point (say) $\mathbf{x}_*$ to be simulated next—given the n old (already simulated) combinations $\mathbf{x}_i$ with $i = 1,...,n$; in the initial stage, $n = n_0$, and in the next stages n is updated one-by-one. To select this $\mathbf{x}_*$, EGO uses a Kriging model to estimate the *acquisition function* or *infill criterion*; the most popular criterion is the *expected improvement* (EI):

$$\text{EI}(\mathbf{x}_*) = E[\max \ (w_{0;\ \min} - \widehat{y}(\mathbf{x}_*), 0)] \text{ with } w_{0;\ \min} = \min_{1 \leq i \leq n} [w_0(\mathbf{x}_i)]. \tag{14}$$

EGO tries to select $\mathbf{x}_*$ that maximizes $\text{EI}(\mathbf{x}_*)$. Therefore, Jones et al. (1998) derives the following estimator where Φ and ϕ denote the CDF and the PDF of the standard normal variable z:

$$\widehat{\text{EI}}_0(\mathbf{x}_*) = (w_{0;\ \min} - \widehat{y}_0(\mathbf{x}_*)) \, \Phi \left(\frac{w_{0;\ \min} - \widehat{y}_0(\mathbf{x}_*)}{s[\widehat{y}_0(\mathbf{x}_*)]} \right) + s[\widehat{y}_0(\mathbf{x}_*)] \, \phi \left(\frac{w_{0;\ \min} - \widehat{y}_0(\mathbf{x}_*)}{s[\widehat{y}_0(\mathbf{x}_*)]} \right).$$

If $s[\widehat{y}_0(\mathbf{x}_*)]$ is high, then maximizing $\widehat{\text{EI}}_0(\mathbf{x}_*)$ stimulates exploration instead of exploitation.

For random simulation-optimization without output constraints, Angün and Kleijnen (2023) uses the estimated *modified EI* (MEI) which uses $\widehat{y}_{0;\ \min}$ instead of the unknown $w_{0;\ \min}$ and the prediction error s_{OK} of OK:

$$\widehat{\text{MEI}}_0(\mathbf{x}) = (\widehat{y}_{0;\ \min} - \widehat{y}_0(\mathbf{x}))\,\Phi\left(\frac{\widehat{y}_{0;\ \min} - \widehat{y}_0(\mathbf{x})}{s_{\text{OK}}[\widehat{y}_0(\mathbf{x})]}\right) + s_{\text{OK}}[\widehat{y}_0(\mathbf{x})]\,\phi\left(\frac{\widehat{y}_{0;\ \min} - \widehat{y}_0(\mathbf{x})}{s_{\text{OK}}[\widehat{y}_0(\mathbf{x})]}\right)$$
$$\text{with } \widehat{y}_{0;\ \min} = \min_{1\le i\le n}[\widehat{y}_0(\mathbf{x}_i)]. \tag{15}$$

In random simulation-optimization with constrained outputs, Angün and Kleijnen (2023) controls the probability of selecting an infeasible solution; i.e., accept $\mathbf{x}$ as a *feasible* point—with a prespecified α_{infe}—if

$$\widehat{y}_1(\mathbf{x}) + z_{(1-\alpha_{\text{infe}})}\,s[\widehat{y}_1(\mathbf{x})] \le c_1. \tag{16}$$

We assume that as n increases, $s[\widehat{y}_1(\mathbf{x})]$ tends to decrease and $\widehat{y}_1(\mathbf{x})$ converges to $q_{1-\alpha}[w_1(\mathbf{x})]$ so $q_{1-\alpha}[w_1(\mathbf{x})] \le c_1$ in (3) holds. In our experiment we select $\alpha_{\text{infe}} = 1\%$, so $z_{(1-\alpha_{\text{infe}})} = 2.3263$.

KKT-EGO *penalizes* $\widehat{\text{MEI}}(\mathbf{x})$ if the KKT conditions do not hold at this $\mathbf{x}$. In our problem defined in (3), these conditions imply the following two conditions where $\nabla_0(\mathbf{x})$ and $\nabla_1(\mathbf{x})$ denote the gradient of $E[w_0(\mathbf{x})]$ and $q_{1-\alpha}[w_1(\mathbf{x})]$: (i) $\mathbf{x}$ lies on the boundary of the feasible area (so the constraint is active or binding): $q_{1-\alpha}[w_1(\mathbf{x})] = c_1$; (ii) $-\nabla_0(\mathbf{x})$ and $\nabla_1(\mathbf{x})$ point into the same direction: $-\nabla_0(\mathbf{x}) = \lambda(\mathbf{x})\nabla_1(\mathbf{x})$ with $\lambda(\mathbf{x}) \ge 0$ where λ denotes the Lagrange multiplier. Obviously, (ii) implies that $E[w_0(\mathbf{x})]$ and $q_{1-\alpha}[w_1(\mathbf{x})]$ are differentiable; the differentiability of $q_{1-\alpha}[w_1(\mathbf{x})]$ implies that its CDF $F(w_1)$ has no jumps or kinks. However, the *black-box* simulation model implies (i) an unknown boundary, and (ii) unknown gradients. KKT-EGO solves these two problems as follows.

Sub (i): KKT-EGO infers that the output constraint is *binding* at $\mathbf{x}$ if the following two-sided CI with prespecified α_{BC} (where we let BC stand for "binding constraint") holds:

$$\frac{|\widehat{y}_1(\mathbf{x}) - c_1|}{s[\widehat{y}_1(\mathbf{x})]} \le z_{1-\alpha_{\text{BC}}/2}. \tag{17}$$

In our experiment we select $\alpha_{\text{BC}} = 20\%$, so $z_{1-\alpha_{\text{BC}}/2} = 1.2816$. Because in our (s,S) simulation, the input bounds $\mathbf{l}$ and $\mathbf{u}$ in (3) are rather arbitrary, we assume that the input constraints are not binding when searching for $\mathbf{x}$. Formulas for problems with multiple binding output or input constraints are derived in Angün and Kleijnen (2023).

Sub (ii): KKT-EGO estimates ∇_h via $\widehat{y}(\mathbf{x})$, which gives $\widehat{\nabla}_h(\mathbf{x}) = \nabla[\widehat{y}_h(\mathbf{x})] = (\partial[\widehat{y}_h(\mathbf{x})]/\partial x_j)'$. We let $c_{h;i}$ denote component i of $\mathbf{c}_h = (\widehat{\mathbf{\Sigma}}_{M;h} + \widehat{\mathbf{\Sigma}}_{\bar{e};h})^{-1}(\overline{\mathbf{o}}_h - \widehat{\mu}_h \mathbf{1}_n)$. Then, the kernel (7) gives

$$\frac{\partial[\widehat{y}_h(\mathbf{x})]}{\partial x_j} = -2\widehat{\tau}_h^2\,\widehat{\theta}_{h;j}\{\Sigma_{i=1}^n c_{h;i}(x_{*;j} - x_{i;j})\exp[\Sigma_{j'=1}^k - \widehat{\theta}_{h;j'}(x_{*;j'} - x_{i;j'})^2]\}. \tag{18}$$

Letting $\widetilde{}$ denote *least squares* (LS) estimators, we obtain the following LS estimator of $\lambda(\mathbf{x})$:

$$\widetilde{\lambda}(\mathbf{x}) = [\widehat{\nabla}_1(\mathbf{x})'\widehat{\nabla}_1(\mathbf{x})]^{-1}\widehat{\nabla}_1(\mathbf{x})'[-\widehat{\nabla}_0(\mathbf{x})].$$

This $\widetilde{\lambda}(\mathbf{x})$ gives the following LS model with the explained (dependent) variable $-\widetilde{\nabla}_0(\mathbf{x})$ and the explanatory (independent) variable $\widehat{\nabla}_1(\mathbf{x})$:

$$-\widetilde{\nabla}_0(\mathbf{x}) = \widetilde{\lambda}(\mathbf{x})\widehat{\nabla}_1(\mathbf{x}).$$

To quantify how well the KKT conditions hold, we compute the *angle* between $\widehat{\nabla}_0(\mathbf{x})$ and $\widetilde{\nabla}_0(\mathbf{x})$, which is measured by the following formula where $\bullet$ denotes the inner product of two vectors and $||.||$ denotes the l_2-norm:

$$\widetilde{\cos}(\mathbf{x}) = \frac{\widehat{\nabla}_0(\mathbf{x}) \bullet \widetilde{\nabla}_0(\mathbf{x})]}{||\widehat{\nabla}_0(\mathbf{x})]|| \times ||\widetilde{\nabla}_0(\mathbf{x})||}. \tag{19}$$

Ideally, $\widehat{\nabla}_0(\mathbf{x})$ and $\widetilde{\nabla}_0(\mathbf{x})$ point into exactly the same direction so their angle is zero and $\widetilde{\cos}(\mathbf{x}) = 1$. If $\widetilde{\lambda}(\mathbf{x}) < 0$ (so, $\widehat{\nabla}_0(\mathbf{x})$ and $\widetilde{\nabla}_0(\mathbf{x})$ point into opposite directions), then $\widetilde{\cos}(\mathbf{x}) < 0$. (This $\widetilde{\cos}(\mathbf{x})$ is related to R^2 with $0 \leq R^2 \leq 1$, which is a popular measure for fit in multiple regression, and is related to the Pearson correlation coefficient ρ with $-1 \leq \rho \leq 1$ in simple regression; perfect fit implies $R^2 = 1$ and in case of a single binding constraint $\rho = 1$.)

Sub (i) and (ii): Altogether, KKT-EGO searches for $\widehat{\mathbf{x}}_o$ that maximizes the acquisition function

$$a_{\text{KKT}}(\mathbf{x}) = \widehat{\text{MEI}}(\mathbf{x}) \times \widetilde{\cos}(\mathbf{x}): \ \widehat{y}_1(\mathbf{x}) + z_{(1-\alpha_{\text{BC}})} s[\widehat{y}_1(\mathbf{x})] \leq c_1 \tag{20}$$

where $\widehat{\text{MEI}}(\mathbf{x})$ equals $\widehat{\text{MEI}}_0(\mathbf{x})$ in (15) provided $\widehat{y}_{0;\,\min}$ is limited to $\mathbf{x}$ satisfying (16), and the search is limited to the estimated feasible area defined by the one-sided CI in (20). The factor $\widehat{\text{MEI}}(\mathbf{x})$ in (20) drives the search for $\widehat{\mathbf{x}}_o$ to the binding constraint, and the factor $\widetilde{\cos}(\mathbf{x})$ drives that search along that constraint to the point where the KKT conditions hold.

To find this $\widehat{\mathbf{x}}_o$, KKT-EGO uses MATLAB's function for finding the minimum of a constrained problem—called *fmincon*. Because fmincon is a local optimizer, we use $n_{\text{fmincon}} > 1$ starting points; in our experiment we select $n_{\text{fmincon}} = 20$. To sample these n_{fmincon} points in the experimental area $\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}$, we apply LHS without midpoints. For further details on the use of fmincon in KKT-EGO we refer to Angün and Kleijnen (2023).

After finding $\widehat{\mathbf{x}}_o$, KKT-EGO obtains simulation replications for $\widehat{\mathbf{x}}_o$—applying Law's procedure defined in (13). Next, KKT-EGO accepts $\widehat{\mathbf{x}}_o$ as feasible if $\mathbf{x} = \widehat{\mathbf{x}}_o$ satisfies the feasibility constraint (16). If KKT-EGO accepts this $\widehat{\mathbf{x}}_o$ as feasible, then $\widehat{\boldsymbol{x}}_o$ may give a better (lower) solution $\widehat{y}_{0;\,\min}$; KKT-EGO then uses this $\widehat{\mathbf{x}}_o$ to update the estimated optimal input (say) $\widehat{\mathbf{x}}_{\min}$. KKT-EGO terminates when it satisfies a prespecified stopping criterion; in our experiment, we (rather arbitrarily) stop the search after 95 iterations.

An obvious alternative for KKT-EGO uses the estimator of the *probability of feasibility* (PF) (also see (17)):

$$\widehat{\text{PF}}(\boldsymbol{x}) = \Phi\left(\frac{c_{h'} - \widehat{y}_1(\mathbf{x})}{s[\widehat{y}_1(\mathbf{x})]}\right).$$

Our (s,S) example has a single constraint; otherwise $\widehat{\text{PF}}(\mathbf{x})$ would depend on all constraints including nonbinding constraints, so $\widehat{\text{PF}}(\mathbf{x})$ would decrease as the number of constraints increases, and $\widehat{\text{PF}}(\mathbf{x})$ would treat all constraints as statistically independent. So, our example (with its single constraint) favors this alternative EGO method. Altogether, this alternative—which we call PF-EGO—uses the acquisition function

$$a_{\text{PF}}(\mathbf{x}) = \widehat{\text{MEI}}(\mathbf{x}) \times \widehat{\text{PF}}(\mathbf{x}). \tag{21}$$

PF-EGO is related to EGO in Carpio et al. (2018) and BO in Gardner et al. (2014). We shall return to PF-EGO, in the next section.

4 NUMERICAL EXAMPLE: A SPECIFIC (s,S) MODEL

The (s,S) model implies that a new order is placed as soon as the inventory position—defined as on-hand stock minus backorders plus outstanding orders—decreases below the reorder level s, and the size of this order is such that the inventory position increases to the order-up-to level S. Like Bashyam and Fu (1998), we assume that the inventory is monitored per period p with $p = 1, 2, ..., P$, where P terminates the simulation run. To find the optimal s and S when orders may cross in time, we need to apply simulation.

For our numerical experiment, we use a PC with multiple cores and parallel software that enable 12 restarts of KKT-EGO such that each restart samples its own initial design. Such parallel computing implies that the restarts do not increase wall-clock time. Within each of these restarts, we use 20 non-parallel restarts of fmincon.

In our specific (s,S) discrete-event simulation, demand D is exponentially distributed with mean 100, and lead time L is integer-valued Poisson distributed with mean 6. Orders are received at the beginning of each period p; demand for this period is subtracted, and an order review is carried out.

To quantify the consequences of risk aversion, we compare the results of an (s,S) simulation model with (i) a constraint for the expected disservice level, and (ii) a quantile constraint. In Section 1 we have already mentioned that risk-aversion—modeled through a "high" quantile such as $q_{0.90;1}$—implies that s and S change. Now we add that s and S may increase, so we should explore an experimental area with higher maximum values for s and S. Therefore we now keep the same minimum value for s as Angün and Kleijnen (2023) uses; namely, $s_{\min} = 600$. For the maximum of s we select $s_{\max} = 2,400$ (instead of 1,200). Furthermore, besides $x_1 = s$ we define $x_2 = Q = S - s$. To select the bounds for Q, we use the *economic order quantity* (EOQ), which equals 85. We vary Q between EOQ/8 = 10.625 and 8×EOQ = 680. So, we use the box constraints $600 \leq s \leq 2,400$ and $10.625 \leq Q \leq 680$.

Table 1: Initial design for s and Q with $S = s + Q$ in restart 1 of macroreplication 1, with its average outputs $\overline{w}_0$ and $\overline{\widehat{q}}_{0.90;1}$ and their estimated standard deviations s; * denotes $\overline{\widehat{q}}_{0.90;1}(\mathbf{x}_i)$ significantly higher than c_1 = 0.1; number of replications m.

s	Q	$S = s+Q$	$\overline{w}_0$	$s(\overline{w}_0)$	$\overline{\widehat{q}}_{0.90;1}$	$s(\overline{\widehat{q}}_{0.90;1})$	m
1950	66.4	2016.4	1423.3	0.9682	0.0014	0.0002	14
1650	624.2	2274.2	1412.1	1.2475	0.0044	0.0005	14
2250	178.0	2428.0	1787.3	0.9550	0.0000	0.0000	3
1350	289.5	1639.5	943.0	1.7477	0.0277	0.0015	4
750	401.1	1151.1	483.8	0.7007	0.2970*	0.0058	3
1050	512.7	1562.7	777.4	5.0430	0.0890	0.0046	3

As we discussed in Section 3.3, we initially select n_0 = 6 and m_0 = 2. Table 1 displays results for restart 1 of macroreplication 1. Its columns 1 and 2 display an example of $\mathbf{X}_{6\times2}$, which denotes the initial design matrix with n_0 = 6 combinations of k = 2 inputs. This $\mathbf{X}_{6\times2}$ is selected through LHS-with-midpoints for s and Q (so $S = s+Q$ in column 3), so n_0 = 6 and $600 \leq s \leq 2,400$ implies that the smallest midpoint for s is 750 (see cell (5,1)). Columns 4 and 5 display $\overline{w}_0$ and $s(\overline{w}_0)$, computed from m replications per (s,S) combination. Likewise, columns 6 and 7 display $\overline{\widehat{q}}_{0.90;1}$ and $s(\overline{\widehat{q}}_{0.90;1})$, which determine a one-sided CI such that the symbol * denotes a $\overline{\widehat{q}}_{0.90;1}$-value that is significantly higher than c_1 = 0.10. Actually, $\alpha_{\text{infe}} = 1\%$ implies that $\overline{\widehat{q}}_{0.90;1}$ is significantly higher than c_1, in combination 5, which has the lowest s and S. This combination also gives the lowest $\overline{w}_0$ (low s and S implies low average inventory). The other five combinations give low $\overline{\widehat{q}}_{0.90;1}$-values, so they are points *inside* the feasible area. The last column displays m, which depends on $s_h/\overline{o}_h$; e.g., the combinations 1 and 6 give $s_1/\overline{o}_1$ equal to 0.0002 / 0.0014 = 0.14 and 0.0046 / 0.0890 = 0.05 so m equals 14 and 3, respectively. Combination 3 gives $\overline{\widehat{q}}_{0.90;1} = 0.0000$, so we use the absolute error $\beta = 0.01$ instead of γ. The smallest m is 3 and the highest m is 14 (obviously, estimating the 90% quantile requires more simulation observations than estimating the mean).

After this initial design—or iteration 0—each iteration of KKT-EGO estimates $\widehat{\mathbf{x}}_{\text{o}} = (\widehat{s}_{\text{o}}, \widehat{S}_{\text{o}})$ and the corresponding outputs $\overline{w}_0$ and $\overline{\widehat{q}}_{0.90;1}$. Furthermore, we decide to obtain 10 *macroreplications*, which sample different $\mathbf{X}_{6\times2}$. We note that all n_0 points require $m > m_0 = 2$ because $s(\overline{\widehat{q}}_{0.90;1})$ is relatively high.

The *true* optimal $(s_{\text{o}}, Q_{\text{o}})$ combination is unknown; see again Bashyam and Fu (1998). Therefore, we use *brute force*; i.e., we obtain "enough" macroreplications for several (s,Q) combinations in the area $1150 \leq s \leq 1215$ and $20 \leq Q \leq 90$; see Figure 2a that we discuss below (this area is much smaller than the original experimental area $600 \leq s \leq 2,400$ and $10 \leq Q \leq 680$; see Table 1). For s we consider fourteen values that lie five units apart (i.e., 1150, 1155, ..., 1210, 1215); for Q we consider fifteen values five units apart (i.e., 20, ..., 90). So, we consider $14 \times 15 = 210$ points. For each point we select the number of replications via Law's (12) with α_m = 10% and γ = 1%. Altogether, we estimate that the "true" optimal values are s_{o} = 1190 and Q_{o} = 60, which give $q_{0.10}$ = 0.0998 and w_0 = 660.

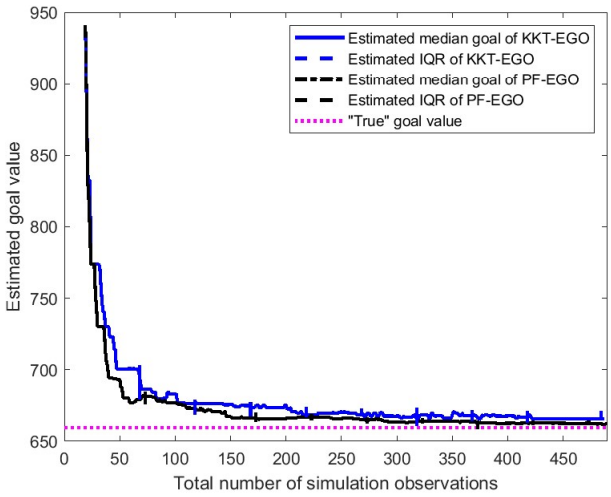

Figure 1: Convergence plots for KKT-EGO and PF-EGO with medians and IQRs estimated from ten macroreplications

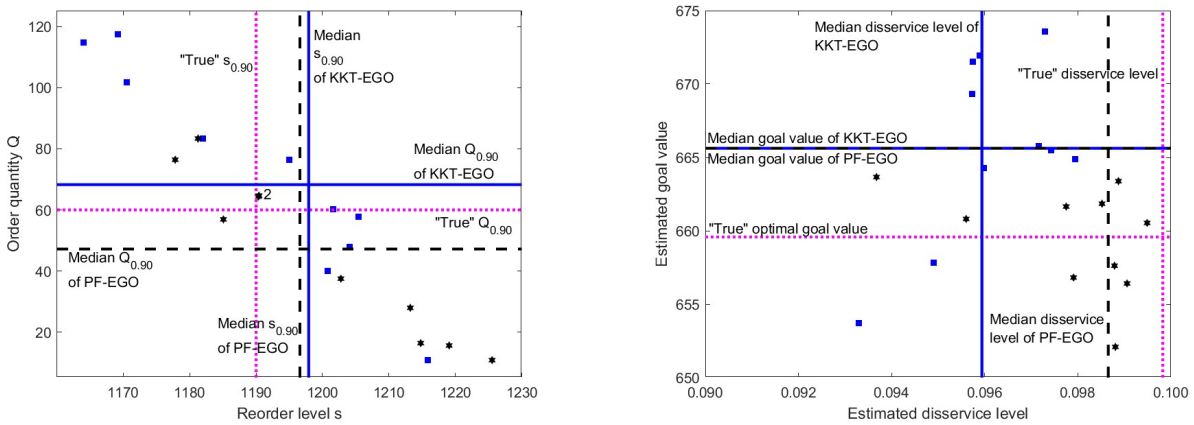

(a) (b)

Figure 2: (a) Final $(\widehat{s}_o, \widehat{Q}_o)$ for KKT-EGO and PF-EGO in ten macroreplications; (b) Final $(\overline{\widehat{q}}_{0.90;1}, \overline{w}_0)$ for KKT-EGO and PF-EGO in ten macroreplications

Figure 1 displays the *convergence plots* for KKT-EGO and PF-EGO. The long horizontal line denotes our *brute force* estimate of the true unknown optimal goal value, which is—with 90% confidence—within 1% of the true value. The plots display *median* goal values—estimated from ten macroreplications—at selected total number of simulations required by the iterations of KKT-EGO and PF-EGO, respectively. Vertical lines display estimated *interquartile ranges* (IQRs). Because these IQRs overlap for the two EGO algorithms, we infer that these algorithms do not give significantly different results, for this (s,S) model.

Fig. 2a displays the *final* $(\widehat{s}_o, \widehat{Q}_o)$ values—after 95 iterations—for the ten macroreplications, corresponding with Fig. 1. The symbol "2" means that two values are so close that they coincide, given the scale of the graph. We estimate the medians of these ten $(\widehat{s}_o, \widehat{Q}_o)$ values; KKT-EGO gives the solid (blue) lines, and PF-EGO gives the dashed (black) lines. Actually, our problem defined in (3) implies that we

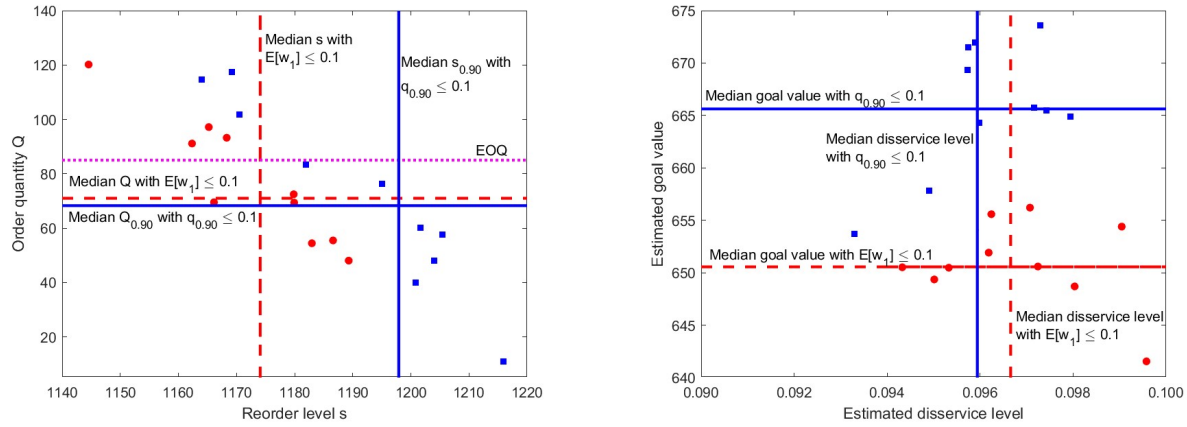

(a) (b)

Figure 3: (a) Final $(\widehat{s}_o, \widehat{Q}_o)$ for quantile constraint versus mean constraint, for KKT-EGO in ten macroreplications; (b) Disservice versus costs for the quantile constraint and the mean constraint

are more interested in $E(w_0)$ and the corresponding $q_{0.90;1}(w_1)$ than in $\widehat{s}_o$ and $\widehat{Q}_o$; so, we proceed to the next graph.

Fig. 2b displays $(\overline{\widehat{q}}_{0.90;1}, \overline{w}_0)$ for KKT-EGO and PF-EGO, corresponding with Fig. 2a. KT-EGO gives an estimated median for $\overline{\widehat{q}}_{0.90;1}$ that is slightly smaller than the PF-EGO median and the "true" value; yet, both median goal values are very close. Next we present the effects of a quantile constraint instead of the mean constraint, focusing on KKT-EGO.

Fig. 3a displays the final $(\widehat{s}_o, \widehat{Q}_o)$ for the $q_{0.90;1}$ quantile constraint in problem (3)—blue squares—and the $E[w_1(\mathbf{x})]$ constraint—red dots—for KKT-EGO. The two constraints give virtually the same estimated median for $\widehat{Q}_o$. If the constraint changes from $q_{0.90}[w_1(\mathbf{x})] \leq 0.1$ to $E[w_1(\mathbf{x})] \leq 0.1$, then $\widehat{s}_o$ decreases.

Fig. 3b displays disservice versus costs for the quantile constraint and the mean constraint. KKT-EGO gives disservice slightly smaller than c = 0.10%, as required by both constraints (KKT-EGO stays on the "safe" side of the boundary; see $\alpha_{\text{infe}} = 1\%$ in (16)). Obviously, costs increase for the risk-averse quantile constraint (this increase may be prohibitive in supermarket inventory management with its small profit margin).

5 CONCLUSIONS AND FUTURE RESEARCH

In this paper we explained how it is simple to apply Angün and Kleijnen (2023)'s KKT-EGO algorithm to solve optimization problems with quantile constraints instead of mean constraints; i.e., the black-box simulation has a quantile as one of its outputs, and KKT-EGO tries to ensure that this output satisfies its prespecified threshold. To illustrate KKT-EGO for quantile constraints, we investigated an inventory system that is controlled by an (s,S) model specified in Bashyam and Fu (1998). We defined risk-aversion as the requirement that the 90% quantile—instead of the mean—of the disservice level remain below 0.10. To estimate the values of the decision variables s and S that satisfy this service-level constraint while minimizing the mean inventory costs, we applied KKT-EGO and a benchmark; namely, PF-EGO. Our numerical results showed that the risk-averse requirement increases cost slightly. It is up to management to decide whether this cost increase is acceptable.

In future research we may apply more alternative EGO algorithms, to more examples. Furthermore, we may investigate how to incorporate uncertainty of the estimated gradients into KKT-EGO. Finally, we may investigate epistemic uncertainty (besides aleatory uncertainty). This uncertainty is investigated in Parmar

et al. (2022)—but not in the context of optimization. Both epistemic uncertainty and aleatory uncertainty are investigated in Wauters (2024), in robust optimization.

ACKNOWLEDGEMENT

We thank the three anonymous reviewers for their very useful comments on a previous version. We also thank Dr. Sebastian Rojas Gonzalez (Postdoctoral research fellow, FWO, Belgium) for providing his MATLAB code for SK, which he developed while he was a Ph.D. student at KU Leuven, Belgium. This work has been financially supported by Galatasaray University Research Fund with project number FBA-2024-1261.

REFERENCES

Alexopoulos, C., D. Goldsman, A. C. Mokashi, K-W. Tien, and J. R. Wilson. 2019. "Sequest: a sequential procedure for estimating quantiles in steady-state simulations". *Operations Research* 67(4): 1162–1183.

Angün, E. and J. Kleijnen. 2023. "Constrained optimization in random simulation: efficient global optimization and Karush-Kuhn-Tucker conditions". CentER Discussion Paper 2023-010, https://www.tilburguniversity.edu/research/economics-and-management/publications/discussion-paper.

Ankenman, B., B. Nelson, and J. Staum. 2010. "Stochastic kriging for simulation metamodeling". *Operations Research* 58(2): 371–382.

Baker, E., P. Barbillon, A. Fadikar, R. B. Gramacy, R. Herbei, D. Higdon, J. Huang, L. R. Johnson, P. Ma, A. Mondal, B. Pires, J. Sacks, and V. Sokolov. 2022. "Analyzing stochastic computer models: a review with opportunities". *Statistical Science* 37(1): 64–89.

Bashyam, S. and M. C. Fu. 1998. "Optimization of (s,S) inventory systems with random lead times and a service level constraint". *Management Science* 44: 243–256.

Carpio, R., R. C. Giordano, and A. R. Secchi. 2018. "Enhanced surrogate assisted framework for constrained global optimization of expensive black-box functions". *Comput. Chem. Eng.* 118: 91–102, https://doi.org/10.1016/j.compchemeng.2018.06.027.

Chang, K. H. and R. Cuckler. 2022. "Applying simulation optimization for agile vehicle fleet sizing of automated material handling systems in semiconductor manufacturing". *Asia-Pacific Journal of Operational Research* 39(2 2150018): 22 pages, https://doi.org/10.1142/S0217595921500184.

Chang, K. H. and C. P. Lin. 2023. "An efficient simulation optimization method for the redundancy allocation problem with a chance constraint". *Journal of the Operational Research Society*, https://doi.org/10.1080/01605682.2023.2272860.

Chen, X. and K. Kim. 2016. "Efficient VaR and CVaR measurement via stochastic kriging". *INFORMS Journal on Computing* 28(4): 629–644.

Gardner, J. R., M. J. Kusner, Z. E. Xu, K. Q. Weinberger, and J. P. Cunningham. 2014. "Bayesian optimization with inequality constraints". *ICML*: 937–945.

Garnett, R. 2023. *Bayesian Optimization*. Cambridge University Press.

Hu, J., M. Song, and M. C. Fu. 2024. "Quantile optimization via multiple-timescale local search for black-box functions". *Operations Research*, https://doi.org/10.1287/opre.2022.0534.

Jones, D., M. Schonlau, and W. Welch. 1998. "Efficient global optimization of expensive blackbox functions". *Journal of Global Optimization* 13: 455–492.

Kleijnen, J. P. C. 2015. *Design and Analysis of Simulation Experiments*. 2nd ed. Springer.

Kleijnen, J. P. C., E. Angün, I. van Nieuwenhuyse, and W. C. M. van Beers. 2023. "Constrained optimization in simulation: efficient global optimization and Karush-Kuhn-Tucker conditions". CentER Discussion Paper 2022-015, https://pure.uvt.nl/ws/portalfiles/portal/62445340/2022_015.pdf.

Kroetz, H. M., M. Moustapha, A. T. Beck, and B. Sudret. 2020. "A two-level Kriging-based approach with active learning for solving time-variant risk optimization problems". *Reliability Engineering and System Safety* 203: 107033.

Law, A. M. 2015. *Simulation Modeling and Analysis*. 5th ed. Boston: McGraw-Hill.

Lolos, A., C. Alexopoulos, D. Goldsman, K. D. Dingeç, A. C. Mokashi, and J. R. Wilson. 2023. "A fixed-sample-size method for estimating steady-state quantiles". In *2023 Winter Simulation Conference (WSC)*, 457–468 https://doi.org/10.1109/WSC60868.2023.10407872.

Parmar, D., L. E. Morgan, S. M. Sanchez, A. C. Titman, and R. A. Williams. 2022. "Input uncertainty quantification for quantiles". In *2022 Winter Simulation Conference (WSC)*, 97–108 https://doi.org/10.1109/WSC57314.2022.10015272.

Song, E., H. Lam, and R. R. Barton. 2024. "A shrinkage approach to improve direct bootstrap resampling under input uncertainty". *INFORMS Journal on Computing*, https://doi.org/10.1287/ijoc.2022.0044.

Wang, H. and K. Yang. 2023. "Bayesian optimization". *Many-Criteria Optimization and Decision Analysis*, edited by D. Brockhoff, M. Emmerich, B. Naujoks, and R. Purshouse, Springer.

Wang, S., S. H. Ng, and W. B. Haskell. 2022. "A multilevel simulation optimization approach for quantile functions". *INFORMS Journal on Computing* 34(1): 569–585.

Wang, X., Y. Jin, S, Schmitt, and M. Olhofer. 2023. "Recent advances in Bayesian optimization". *ACM Computing Surveys* 55(13s): 1–36.

Wauters, J. 2024. "ERGO-II: an improved Bayesian optimization technique for robust design with multiple objectives, failed evaluations and stochastic parameters". *Journal of Mechanical Design*, https://doi.org/10.1115/1.4064674.

BIOGRAPHIES

Ebru Angün is Associate Professor of Operations Research in the Department of Industrial Engineering at Galatasaray University in Istanbul. Her research interests include simulation optimization, stochastic optimization, and applications in disaster management. She has been an author and reviewer for many international journals, and a project reviewer for the National Science Foundation of Turkey and the Research Foundation Flanders (FWO) in Belgium. She spent one year at the Georgia Institute of Technology, as a visiting scholar. Also see https://avesis.gsu.edu.tr/eangun and https://www.researchgate.net/profile/Ebru-Angun.

Jack P.C. Kleijnen is Emeritus Professor of Simulation and Information Systems at Tilburg University, where he is still active in the Department of Management and in the Center for Economic Research (CentER). His research focuses on the statistical design and analysis of simulation experiments. He is an author and reviewer for many international journals. He was a consultant for several organizations in the USA and Europe, and served on many international editorial boards and scientific committees. He spent several years in the USA, at universities and companies. In 2008 he became a Knight in the Order of the Netherlands Lion, and in 2005 he received the Lifetime Professional Achievement Award from the Simulation Society of the Institute for Operations Research and Management Science (INFORMS). Also see https://sites.google.com/view/kleijnen/home.