This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Robust Sampling Budget Allocation Under Deep Uncertainty

Michael Perry, Jie Xu, *Senior Member, IEEE*, Edward Huang, and Chun-Hung Chen, *Fellow, IEEE*

***Abstract*—A novel methodology is introduced for optimally allocating a sampling budget. Sampling budget allocation problems arise frequently in various settings. For example, in the design of complex engineering systems, given both the complexity of these systems and the imperfect information on new technologies, designers often face deep uncertainty as to system performance. Consequently, designers need to sample multiple alternative designs under a limited budget. This article proposes a minimax regret approach to allocate the sampling budget in the presence of deep uncertainty pertaining to system performance. The objective is to maximize the probability of selecting the design with the minimum–maximum regret under a limited sampling budget and imperfect information. To effectively solve the minimax regret problem, an approximation methodology that provides good solutions with quantifiable uncertainty is developed. The essence of the methodology, which has the added benefit of being generally applicable to any multilevel optimization, is that all but the first level of multilevel optimization can be eliminated via a response surface. By sampling many values of a higher level decision's variables, solving the next lower level optimization given those samples values, and calibrating a response surface to the objective function value eliminate one required optimization. Doing this repeatedly reduces the complexity of the multilevel optimization to a standard optimization. Regardless of the number of levels in the optimization, repeating this process ultimately leaves one with a single optimization whose objective function can be directly computed, given the highest level variables. Numerical experiments with two sampling allocation examples demonstrate both the benefit of the robust sampling budget allocation versus nonrobust formulations and the effectiveness of the proposed solution approach.**

***Index Terms*—Complex system design, minimax regret, response surface methodologies, sampling budget allocation.**

I. Introduction

IN THIS article, a novel methodology is introduced for allocating a sampling budget under deep uncertainty. Deep uncertainty is defined as a situation where decision makers lack probability distributions for some or all uncertain parameters [14]. Sampling budget allocation is an important problem arising in many different contexts. For example, in the design of complex systems, where the utility of design alternatives is often understood by running simulations that are computationally expensive, often requiring multiple days to generate a single sample, the designer must carefully allocate a limited sampling budget among many design alternatives. Classical approaches to sampling allocation for the system design typically assume the mean and variance of each design alternative's utility can be reliably estimated by a subject-matter expert (SME); an optimal sampling allocation is then determined using these estimates which maximizes the probability of identifying the best design, post-sampling. Decisions made under the classical framework are likely to be poor, however, if an SME fails to accurately estimate means and variances. Realistically, in the design of complex systems means and variances of alternative designs' utilities cannot be accurately estimated prior to generating samples for a variety of reasons: complex systems have many interacting parts, so new system designs are very unpredictable; even when implementing an existing design in a new environment, interactions with the environment may cause unpredictable behavior. For this reason, this article relaxes the assumption an SME can provide a point estimate of design utility means and variances and instead assumes an SME can only specify an uncertainty set in which each lies. Optimizing sampling budget allocation when means and variances are unknown, but within a given uncertainty set, presents a robust optimization (RO).

ROs are a class of multilevel optimizations where a decision maker needs to make a decision under uncertainty, but instead of assuming a probability distribution for the uncertain variables in the problem, it is assumed nature is an intelligent actor who picks these variables, so as to create the worst possible outcome for the decision maker. The simplest RO formulation contains a sequence of two optimizations ("bilevel"), where a decision is made in the first optimization, and nature then selects values for uncertain parameters to produce the worst case result as the second optimization. As discussed further in Section II, this formulation is often overly conservative and a common remedy is to formulate a trilevel minimax regret problem: 1) a decision is made; 2) nature then selects the uncertain parameters; and 3) finally, a notional decision maker selects values for the original decision variables that would have been optimal, given the realized values of the uncertain parameters. The actual decision maker's objective is to minimize a quantity called "regret," which is the difference between the realized utility and what could have been realized had the uncertain parameters been known.

Manuscript received May 30, 2021; revised September 18, 2021; accepted January 16, 2022. This work was supported in part by the National Science Foundation under Award FAIN 2123683 and Award DMS-1923145, and in part by the Air Force Office of Scientific Research under Award FA9550-19-1-0383. This article was recommended by Associate Editor L. Fang. *(Corresponding author: Michael Perry.)*

The authors are with the Department of Systems Engineering and Operations Research, George Mason University, Fairfax, VA 22030 USA (e-mail: mperry20@gmu.edu; jxu13@gmu.edu; chuang10@gmu.edu; cchen9@gmu.edu).

Digital Object Identifier 10.1109/TSMC.2022.3144363

Applying this notion to the sampling budget allocation problem, the decision maker is tasked to allocate samples among k designs available for implementation, each with an uncertain utility that is assumed to follow a normal distribution with unknown mean μ and standard deviation σ. All that can be elicited from an SME are uncertainty sets for μ and σ. Generating samples, i.e., running computer simulations, reduces the uncertainty in utility estimates, which in turn affects the probability, post-sampling, of correctly selecting the best design. However, the sampling allocation that maximizes the probability of correct selection (PCS) of the best design can be quite different depending on where μ and σ lie in their uncertainty sets. Thus, a methodology is needed to find the sampling allocation that minimizes the maximal difference between potential PCS and realized PCS (i.e., regret).

RO problems are notoriously difficult to solve, especially in the trilevel, minimax regret form. To tackle this challenge, this article develops an approximation technique that well approximates the exact solution. The essence of the methodology, described fully in Section III, is that all but the highest level of the optimization can be reduced by sampling many values of the higher level decision variables at a given level (which are assumed known during lower levels), solving the optimization given those sampled values, and then estimating a response surface that predicts the objective function value as a function of the higher level decision variables. This transforms each lower level optimization into an expression of the higher level decision variables, which can then be optimized over. The process is continued until all that is left is the highest level of the optimization, whose objective function is now a response surface with a value that can be directly computed, given the highest level decision variables.

The remainder of this article proceeds as follows. Section II presents a brief literature review of sampling budget allocation problems, robust and other multilevel optimizations, and response surface fitting. Section III describes the methodology and Section IV provides example problems, their solution, and comparisons with nonrobust benchmarks. Section V comments on issues to be addressed by future research as well as some concluding remarks.

II. Literature Review

While limited research has been done on robust approaches to the sampling allocation problem, the broader literature on this subject is quite rich. The Ranking & Selection (R&S) literature is concerned with the allocation of simulation samples to maximize PCS or provide a probabilistic guarantee on PCS [8], [13], [27], [28]. Optimal computing budget allocation (OCBA) [6], [7], [11], [29], [30] is a popular R&S algorithm that achieves the optimal asymptotic convergence rate for PCS when the total number of samples to be drawn goes to infinity. In this article, no asymptotic assumptions are made, as in general, it may be computationally impractical to generate enough samples for an asymptotic assumption to be an accurate reflection of reality, in the context of designing complex systems using computationally very expensive simulation models. Furthermore, an asymptotic assumption is not needed for the developed response surface methodology to be used. Robust approaches to the sampling allocation problem have not been widely explored, but recent examples include Gao *et al.* [12], Ungredda *et al.* [24], Wu *et al.* [26], and Zhu *et al.* [31], which again use an asymptotic assumption. The literature on nonasymptotic approaches to sampling allocation includes [9], [10], and [17] where the focus is to sequentially allocate one sample at a time with an objective to maximize the improvement achieved by the new sample.

As stated in Section I, RO is a subset of multilevel optimizations. Multilevel optimizations occur when a decision maker must act, and subsequent decisions by intelligent actors influence his realized utility. Multilevel optimizations are often found in game theory where actual intelligent actors are making decisions (e.g., [1], [5], and [23]), as well as in the RO setting where nature is assumed to be an intelligent actor. RO is an alternative to traditional stochastic optimization; rather than specifying probability distributions for uncertain parameters and solving a stochastic optimization, the decision maker assumes nature selects values for the uncertain parameters so as to create a worst case result for the decision maker, given his decision. This framework was developed to allow decision makers lacking sufficient information to specify probability distributions to find solutions that perform well in many scenarios. A good textbook on RO is [3]. The standard RO framework seeks to maximize the minimum value of a utility function, which often gives an overly conservative solution as detailed in [18]. A commonly used alternative formulation to address the conservatism of RO is to minimize the maximum value of regret. This is defined as the difference in utility when a decision maker is allowed to make a decision with perfect knowledge of the uncertain parameters, and the realized utility from making a decision without the knowledge of those parameters. As seen clearly in Section III, the minimax regret formulation requires three levels of optimization, making it a more difficult problem to solve. The proposed methodology is able to approximate the exact solution well, however, with limited additional difficulty compared to the standard robust formulation. We thus adopt the minimax regret formulation to obtain a more practically useful robust sampling budget allocation policy.

Even the smallest scale multilevel optimizations are known to be NP hard [2]. Extensive research has been performed to approximate their solutions. An excellent review of both exact and approximate solution techniques can be found in [15]. Of particular note is that most approximation methods for trilevel optimizations (of which minimax regret is an example) utilize fuzzy programming [22]. While useful for decision makers, these fuzzy methods seek to find satisfactory solutions without any claim of producing the true optimum in expectation. A key component of the methodology developed in this article is that it is not only intended to arrive at the true optimum in expectation but also produce a measurable error term so the decision maker is aware of how much error has been introduced.

This article's methodology relies on fitting a response surface to the inner levels of the optimization. Each time a response surface is fit, an error is introduced. Therefore, fitting surfaces accurately will be a critical component for any decision maker utilizing this method. Recall that when fitting a response surface to a nested optimization, one must draw samples of the model variables that are assumed known at the current level of optimization, optimize over the decision variables and record the objective function value, and repeat this process many times. This leads to two key considerations when fitting response surfaces: 1) drawing sufficient samples to accurately represent the space of unknown variables and 2) choosing a functional form for the response surface that will be easy to optimize when solving the more outer levels of the optimization. Choosing proper functional forms is problem specific, but drawing representative samples is a widely studied problem. Latin hypercube sampling (LHS), for example, is an efficient way to capture the entire space of unknown variables when they lie in a box set, and generalizing LHS to other than box sets is an open research area that has been studied in [16], [19], and [20].

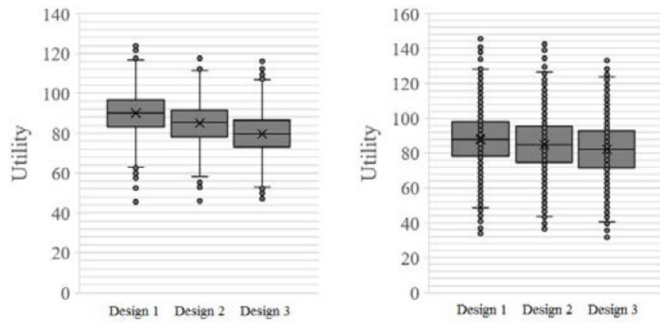

(a) (b)

Fig. 1. Box-and-whiskers plots for theoretical utilities, prior to sampling, are plotted for two possible realizations of μ and σ. In each Fig. 1(a) and (b), uncertainty sets for designs 1–3 are $\mu_1 \in [85, 95]$, $\mu_2 \in [80, 90]$, and $\mu_3 \in [75, 85]$, and $\sigma_1 \in [10, 15]$, $\sigma_2 \in [10, 15]$, and $\sigma_3 \in [10, 15]$. For a sampling allocation of $N_1 = 80$, $N_2 = 15$, and $N_3 = 5$, realized APCS is 95% [Fig. 1(a)] versus 57% [Fig. 1(b)]. (a) Realized $\mu = [90, 85, 80]$, $\sigma = [10, 10, 10]$. (b) Realized $\mu = [88, 85, 82]$, $\sigma = [15, 15, 15]$.

III. Approximation Model for Sampling Allocation With Minimax Regret

This section outlines modeling assumptions, defines the objective function, and formulates the minimax regret problem for that objective. Noting that no analytical solution to this particular minimax regret problem is known to exist, a method for approximating the model is then described.

A. Modeling Assumptions

Assume a decision maker needs to select a single design to implement among k alternatives. All designs have already been fully developed, but there is still uncertainty in how each would perform if implemented. This could be because the designs have never been employed, or because they have only been employed in a different setting than the decision maker intends to use them. It is therefore assumed all the decision maker has is an SME's best guess as to the lower and upper bounds of each design utility's mean and standard deviation, $\mu_i \in [\mu_i^l, \mu_i^u]$ and $\sigma_i \in [\sigma_i^l, \sigma_i^u]$, for $i = 1, 2, \ldots, k$. It is assumed these unknown utilities are normally distributed and the notation $L_i \sim N(\mu_i, \sigma_i)$ denotes the ith design's utility. The decision maker has a total of B samples to allocate to the k designs that will be used to reduce uncertainty before selecting a design. Prior to generating samples, the effect on μ_i cannot be predicted, but it can be assumed variance is reduced by the inverse of the number of samples allocated. That is

$$\sigma_i^2 \rightarrow \sigma_i^2 / N_i \tag{1}$$

where N_i is the number of samples allocated to design i.

The decision maker's objective is to allocate samples to ensure the best design is selected with high probability. The quantity of interest is therefore the PCS, which is defined as the probability the selected design has a higher mean than all others. PCS can be calculated via

$$\text{PCS} = P(L_b > L_i \ \ \forall\, i \neq b) \tag{2}$$

where the subscript b indicates the selected design, which will be that with $\mu_b > \mu_i \ \forall\, i$.

Using (2) as an objective function will make most optimization problems intractable, so as in Chen *et al.* [7], a lower bound on PCS using the Bonferroni inequality [4] will be used in its place. This bound is called the approximate PCS (APCS)

$$\begin{aligned} \text{PCS} &= P(L_b > L_i \ \forall\, i \neq b) \\ &\geq 1 - \sum_{i \neq b} P(L_i > L_b) \\ &= 1 - \sum_{i \neq b} \Phi\left(\frac{\mu_i - \mu_b}{\sqrt{\sigma_b^2 / N_b + \sigma_i^2 / N_i}} \right) \\ &= \text{APCS}. \end{aligned} \tag{3}$$

B. Minimax Regret Optimization for APCS

Notice the distinction between the problem formulated thus far and the traditional sampling allocation problem as in Chen *et al.* [7], where samples are allocated to maximize APCS when $L_i \sim N(\mu_i, \sigma_i)$, and μ_i and σ_i are assumed known for all designs i. The traditional problem solves

$$\begin{aligned} \max_{N} \ & \text{APCS}(N \mid \mu, \sigma) \\ \text{s.t.} \ & \sum_i N_i = B \end{aligned} \tag{4}$$

where B is the sampling budget. This article, on the other hand, still assumes $L_i \sim N(\mu_i, \sigma_i)$, but does not assume μ_i and σ_i are known; all that assumes is they lie in the known uncertainty sets.

As depicted in Fig. 1, for a given sampling allocation, the realized value of APCS could vary considerably depending on the realized values of μ_i and σ_i, as drawn from their uncertainty sets. The uncertainty sets are such that means can be quite different and variances relatively low [Fig. 1(a)], or lead to means that are tightly packed and variances that are high [Fig. 1(b)]. These two cases lead to drastically different sampling strategies. Because all the decision maker has is uncertainty sets for μ_i and σ_i, RO must be used, and for the reasons discussed in Section II, a minimax regret solution concept should be used rather than the (generally more conservative) traditional robust formulation.

Regret problems are intended to minimize the difference between the realized value of an objective function, and the value that could have been realized had one known in advance the value of unknown variables, μ_i and σ_i, for all designs i. Using the vector notation $N := \{N_1, N_2, \ldots, N_k\}$, and similarly for the other problem variables, the minimax regret model, denoted MMR, for the sampling allocation problem is formally defined as follows:

$$
\begin{aligned}
&\text{MMR}:\\
&\min_{N}\ \max_{\mu,\sigma} \mathit{Regret}(N\,|\mu,\sigma)\\
&\text{s.t.}\\
&\quad \text{Regret}(N|\mu,\sigma) = \text{APCS}\left(N^*|\mu,\sigma\right) - \text{APCS}(N,\mu,\sigma)\\
&\quad \text{APCS}\left(N^*|\mu,\sigma\right)\\
&\quad = \max_{N}\ \text{APCS}\left(N\,|\mu,\sigma\right)\ \text{s.t.} \sum_{i=1}^{k} N_i = B,\ N_i \geq 1\ \forall\, i\\
&\quad \mu_i \in \left[\mu_i^l, \mu_i^u\right]\ \forall\, i\\
&\quad \sigma_i \in \left[\sigma_i^l, \sigma_i^u\right]\ \forall\, i\\
&\quad \sum_{i=1}^{k} N_i = B\\
&\quad N_i \geq 1\ \forall\, i.
\end{aligned}
\tag{5}
$$

APCS(N,μ,σ) is simply the quantity in (3) evaluated for given values of μ, σ, and N, and APCS$(N^*|\mu,\sigma)$ is APCS computed using the optimal choice of N^* for known values of μ and σ. The constraint $N_i \geq 1\ \forall\, i$ ensures the uncertainty in a particular design cannot decrease below its initial, the presampling level. Throughout this article, it is assumed the sampling budget is finite, yet sufficiently large so that MMR can be solved using continuous values of N. The methodology presented here applies to small sampling budgets as well, but integer optimization techniques would need to be used when estimating response surfaces.

MMR is a trilevel optimization. Regret problems (like other multilevel optimizations) are solved by first solving the innermost optimization and working backward. While an analytical solution for APCS$(N^*|\mu,\sigma)$ is known if an asymptotic assumption is used (see Chen *et al.* [7]), it is not then known how to solve $\max_{\mu,\sigma}$ APCS$(N^*|\mu,\sigma)$ − APCS(N,μ,σ), and thus an approximation technique is needed.

C. Estimation of Response Surfaces

To overcome the difficulty in solving MMR analytically, this article begins by estimating a response surface from $\{\mu, \sigma\}$ to the objective value of APCS$(N^*|\mu,\sigma)$ using Latin hypercube samples for $\{\mu, \sigma\}$; denote this surface $f_3(\mu, \sigma)$. Next, $f_3(\mu, \sigma)$ is substituted for APCS$(N^*|\mu,\sigma)$ and a second response surface is estimated, using Monte Carlo samples of N to predict $\max_{\mu,\sigma} f_3(\mu, \sigma)$ − APCS(N,μ,σ); denote this surface $f_1(N)$. Finally, $f_1(N)$ is optimized subject to $\sum_{i=1}^{k} N_i = B$ to get an approximate solution to MMR.

To make the methodology described in the preceding paragraph more transparent, and to make the generalizability to other nested optimizations clear, MMR is restated as follows:

$$
\begin{aligned}
&\text{MMR}:\\
&\min_{N}\ z_1(N)\\
&\text{s.t.}\\
&\quad \sum_{i=1}^{k} N_i = B\\
&\quad N_i \geq 1\ \forall\, i
\end{aligned}
\tag{6}
$$

where…

$$
\begin{aligned}
&z_1(N) = \max_{\mu,\sigma}\ z_2(\mu, \sigma|N)\\
&\text{s.t.}\\
&\quad \mu_i \in \left[\mu_i^l, \mu_i^u\right]\ \forall\, i\\
&\quad \sigma_i \in \left[\sigma_i^l, \sigma_i^u\right]\ \forall\, i
\end{aligned}
\tag{6.1}
$$

where…

$$
\begin{aligned}
&z_2(\mu, \sigma|N)\\
&= z_3(\mu, \sigma) - 1 + \sum_{i=b} \left(\frac{\mu_i - \mu_b}{\sqrt{\sigma_b^2/N_b + \sigma_i^2/N_i}}\right)
\end{aligned}
\tag{6.2}
$$

and where…

$$
\begin{aligned}
&z_3(\mu, \sigma) = \max_{N}\ 1 - \sum_{i\neq b} \left(\frac{\mu_i - \mu_b}{\sqrt{\sigma_b^2/N_b + \sigma_i^2/N_i}}\right)\\
&\text{s.t.}\\
&\quad \sum_{i=1}^{k} N_i = B\\
&\quad N_i \geq 1\ \forall\, i.
\end{aligned}
\tag{6.3}
$$

The first step in approximating this model as a single minimization is to estimate a response surface for $z_3(\mu, \sigma)$; this will allow $z_2(\mu, \sigma|N)$ to be rewritten as a function of μ, σ, and N that can be maximized without reference to another optimization. To accomplish this, Latin hypercube samples over the decision space $\mu_i \in [\mu_i^l, \mu_i^u]$, $\sigma_i \in [\sigma_i^l, \sigma_i^u]\ \forall\, i$ are generated, the optimization $z_3(\mu, \sigma)$ is solved for each sample using nonlinear optimization techniques (this article used sequential least squares, as implemented by Python's SciPy package), and the results are stored. The following response surface can then be fit to the samples:

$$
z_3(\mu, \sigma) = \beta_{3,0} + \sum_{i} \beta_{3,i} \cdot \left(\frac{-\mu_b}{\sqrt{\sigma_b^2 + \sigma_i^2}}\right) + \varepsilon_3 \tag{7}
$$

where $\varepsilon_3 \sim N(0, \sigma_3)$ and β_3 values are constant coefficients.

This surface was chosen for the following three general reasons.

1) It is clearly based on the structure of the objective function, APCS(N,μ,σ).
2) It is simple enough such that when inserting it into the objective function for $z_1(N)$, it can be repeatedly

optimized over the decision variables μ and σ, for many samples of N.

3) Furthermore, it happens that the error term ε_3 is independent of the decision variables μ and σ; this was verified for the examples in Section IV via the White test [25]. While not required for the methodology to be implemented, this property means a decision maker can ignore ε_3 when maximizing $z_2(\mu, \sigma|N)$, regardless of risk preferences.

In the examples presented in Section IV, model (7) is seen to fit quite well so alternative specifications were not explored. With this response surface in hand, the next step in approximating MMR is to express $z_2(\mu, \sigma|N)$ as a direct expression of (7). Using the definition of $z_2(\mu, \sigma|N)$ in (...) this gives

$$z_2(\mu, \sigma|N) = \beta_{3,0} + \sum_i \beta_{3,i} \cdot \Phi\left(\frac{\mu_i - \mu_b}{\sqrt{\sigma_b^2 + \sigma_i^2}}\right) - \left(1 + \sum_{i \neq b} \Phi\left(\frac{\mu_i - \mu_b}{\sqrt{\sigma_b^2/N_b + \sigma_i^2/N_i}}\right)\right) + \varepsilon_3. \quad (8)$$

Next, Monte Carlo samples for N are generated such that $\sum_{i=1}^{k} N_i = B$, sequential least squares are used to maximize (8) over $\{\mu, \sigma\}$, and results are stored. The following response surface for $z_1(N)$ is then estimated:

$$z_1(N) = \beta_{1,0} + \sum_i \beta_{1,i} \cdot \sqrt{\frac{N_i}{B}} + \varepsilon_3 + \varepsilon_1 \quad (9)$$

where again $\varepsilon_1 \sim N(0, \sigma_1)$ and β_1 values are constants.

Note that while ε_3 error term from the response surface calibrated to $z_3(\mu, \sigma)$ carries over *to* $z_1(N)$, this term does not complicate the analysis because ε_3 is a random variable independent of all decision variables. The choice of response surface used the same criteria as for $z_3(\mu, \sigma)$, though bullet 2 is now less important as, at this highest level of optimization, $z_1(N)$ will only need to be optimized once. To complete the approximation of MMR, (9) is minimized over N, which is a standard, single-level optimization

$$\widehat{\text{MMR}}: \quad \min_N \; \beta_{1,0} + \sum_i \beta_{1,i} \cdot \sqrt{\frac{N_i}{B}} + \varepsilon_3 + \varepsilon_1$$
$$\text{s.t.} \quad \sum_{i=1}^{k} N_i = B$$
$$N_i \geq 1 \;\forall\, i. \quad (10)$$

As noted previously, ε_1 and ε_3 are independent of any decision variables and thus do not affect the optimal solution of $\widehat{\text{MMR}}$. They are nonetheless explicitly written in the objective function to maintain cognizance that this approximation of MMR is subject to error.

IV. Example Implementing the Response Surface Methodology for MMR

Four examples of interested are now presented, all with $k = 6$ and $B = 100$. This first example represents a scenario most likely to be encountered by a decision maker, where mean uncertainty sets exhibit moderate overlap, and variances are likewise moderate. The other examples look at other reasonable cases, consider higher variances, mean uncertainty sets that are near uniform, and a situation where the highest and lowest means have the highest variances. These latter situations pose additional challenges to robust decision making. Near uniform means, for example, is commonly used in the sample allocation literature as the most challenging test case for a selection methodology. High variances also make the allocation decision more difficult as this makes it more probable empirical means taken from small samples will be misleading. The parameters used are summarized in Table I, and all four examples are repeated with $k = 12$ and $k = 24$. Results are summarized in Table II, while Table III lists the response surface fitting errors for all examples.

TABLE I
Example Parameters

	Design 1	2	3	4	5	6
Moderate overlap in means						
μ^u	105	100	95	90	85	80
μ^l	95	90	85	80	75	70
Near uniformity in means						
μ^u	105	103	101	99	97	95
μ^l	95	93	91	89	87	85
Higher/lower means have higher variance						
σ^u	20	14	8	8	14	20
σ^l	10	7	4	4	7	10
Moderate variance: $\sigma_i^u = 20, \sigma_i^l = 10\ \forall\, i$						
High variance: $\sigma_i^u = 40, \sigma_i^l = 20\ \forall\, i$						

Note. These same parameters are used when $k = 12$ and $k = 24$, with natural extensions. Mean upper- and lower-bounds continue declining at the same rates for successive designs. For the case when variances are higher for the highest and lowest means, the first and last designs have $\sigma^u = 20$ and $\sigma^l = 10$, while the middle designs have $\sigma^u = 8$ and $\sigma^l = 4$.

Each example implements the response surface methodology by generating 1000 samples of $\{\mu, \sigma\}$ to estimate f_3, and 1000 samples of N to estimate f_1. Each surface is estimated using linear least squares. These sample sizes are conservative for the response surface coefficients to converge but were not computationally burdensome.

For each example, the effectiveness of the methodology was evaluated by generating 4000 uniformly distributed samples of $\{\mu, \sigma\}$ from their uncertainty sets and computing the resultant regret values for the solution to (10), which is denoted as $N_{\widehat{MR}}$. The upper quantiles of regret generated in this way give a strong estimation of the maximum regret possible; if this estimator of maximum regret is small, that is evidence the methodology is performing well. For comparison, two reasonable benchmarks are used. The first benchmark is the sampling allocation obtained by assuming all means and variances equal their "most likely" values, defined as the

TABLE II
REGRET QUANTILES

	95th	99th	99.9th	Comment
Moderate overlap in means				
Moderate variance				
$k = 6$	**.06***,.07,.12	**.07***,.10,.14	**.08***,.16,.15	*Ex. A*
$k = 12$	.19,**.15***,.41	.22,**.20***,.46	**.25***,.27,.50	
$k = 24$	.19,**.05***,.39	.23,**.09***,.45	.27,**.14***,.51	
High variance	95th	99th	99.9th	
$k = 6$	**.09***,.14,.14	**.11***,.19,.16	**.14***,.23,.18	*Ex. B*
$k = 12$	**.20***,.41,.42	**.24***,.47,.46	**.27***,.53,.50	
$k = 24$	.30,**.29***,.56	**.34***,.43,.63	**.40***,.48,.70	
Highest/lowest means have highest variance	95th	99th	99.9th	
$k = 6$	**.06***,.07,.16	**.07***,.11,.20	**.10***,.15,.24	*Ex. D*
$k = 12$	**.44***,.45,1.04	**.53***,.56,1.16	**.61***,.70,1.29	
$k = 24$	.16,**.07***,.39	.21,**.13***,.45	.24,**.20***,.55	
Near uniform means				
Moderate variance	90th	95th	99th	
$k = 6$	**.13***,.21,.14	**.160***,.26,.163	**.19***,.30,.20	*Ex. C*
$k = 12$	**.19***,.23,.36	**.25***,.31,.41	**.31***,.41,.46	
$k = 24$	.35,**.23***,.67	.44,**.35***,.76	**.52***,.61,.83	

Note. Each entry lists the regret quantiles for N_{MMR}, N_{ML}, N_{uni}, in that order. The bold face and * indicates the best value among the three.

TABLE III
RESPONSE SURFACE ERRORS

	σ_1	σ_3	Comment
Moderate overlap in means			
Moderate variance			
$k = 6$	0.1685	0.0857	*Ex. A*
$k = 12$	0.3751	0.1407	
$k = 24$	0.4840	0.0980	
High variance	σ_1	σ_3	
$k = 6$	0.2108	0.0555	*Ex. B*
$k = 12$	0.4129	0.0773	
$k = 24$	0.5752	0.0872	
Highest/lowest means have highest variance	σ_1	σ_3	
$k = 6$	0.1659	**0.0931**	*Ex. D*
$k = 12$	0.8705	0.1605	
$k = 24$	0.5800	0.1005	
Near uniform means			
Moderate variance	σ_1	σ_3	
$k = 6$	0.1949	0.0559	*Ex. C*
$k = 12$	0.5413	0.1393	
$k = 24$	**.8888**	0.1478	

midpoints of uncertainty sets. Γ nally, this benchmark

$$N_{\mathrm{ML}} = \operatorname{argmax}_N 1 - {}_{i=2}^{6} \left(\frac{\overline{\mu}_i - \overline{\mu}_{b_{\mathrm{mid}}}}{\overline{\sigma}^2_{b_{\mathrm{mid}}} / N_{b_{\mathrm{mid}}} + \overline{\sigma}^2_i / N_i} \right)$$

$$\text{s.t.} \quad \underset{i}{} N_i = B$$

$$\overline{\mu}_i = \frac{\mu^u{}_i - \mu^l_i}{2} \ \forall\, i$$

$$\overline{\sigma}_i = \frac{\sigma^u_i - \sigma^l_i}{2} \ \forall\, i$$

$$b_{\mathrm{mid}} := i | \overline{\mu}_i > \overline{\mu}_j \ \forall\, j = i. \tag{11}$$

The second benchmark is the uniform sampling allocation, $N_{\mathrm{uni}} = \{B/k, B/k, \ldots, B/k\}$. When facing deep uncertainty, this may perform rather well in minimizing maximal regret. N_{MR} will be compared to these two benchmarks by computing the upper quantiles of regret for each solution.

A. Example: Moderate Overlap in Means, Moderate Variance

Solving (10) yields $N_{\mathrm{MR}} = [25, 22, 18, 17, 11, 8]$. The solution to (11) is $N_{\mathrm{ML}} = [39, 28, 16, 8, 5, 3]$. The response surface methodology has clearly led to a more uniform allocation of samples, though is still far from a pure uniform allocation. Fig. 2 plots the upper quantiles for N_{MR}, N_{ML}, and N_{uni}. As seen, N_{MR} outperforms both N_{ML} and

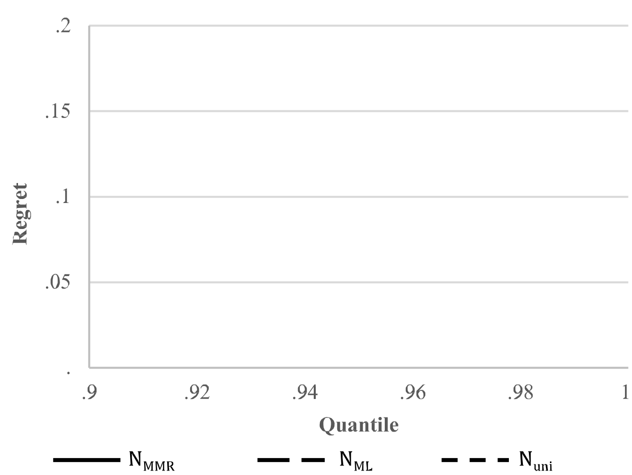

Fig. 2. Regret quantiles for Example A.

N_{uni} by a wide margin. The 99.9th quantiles of the regret in APCS for the three strategies are 0.0828, 0.1600, and 0.1525, respectively. Recalling the definition of APCS as a lower bound on the probability of selecting the best design option, post-sampling, these results show the response surface methodology can reduce a decision maker's potential regret by $1 - 0.0825/0.1525 = 45.7\%$ over the next-best benchmark.

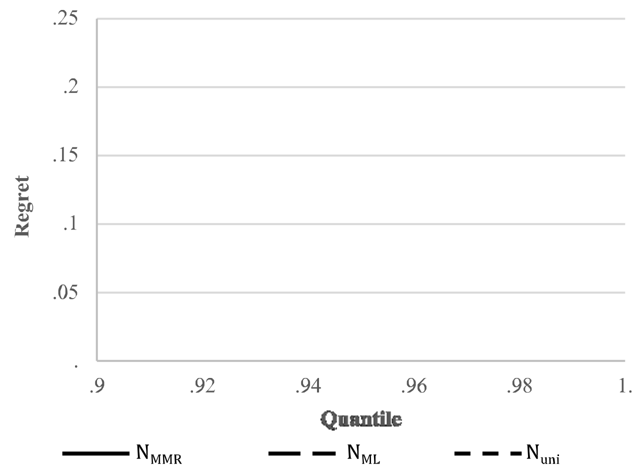

Fig. 3. Regret quantiles for Example B.

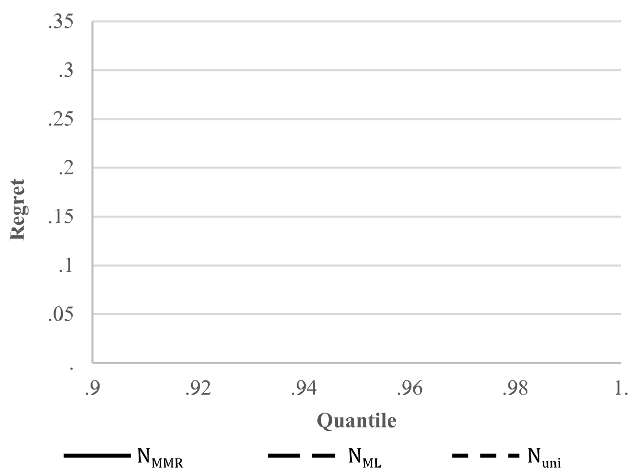

Fig. 4. Regret quantiles for Example C.

B. Example: Moderate Overlap in Means, High Variance

Increasing variances from $\sigma^u = 20$ and $\sigma^l = 10$ to $\sigma^u = 40$ and $\sigma^l = 20$ creates an interesting dynamic between N_{MMR}, N_{ML}, and N_{uni}, as seen in Fig. 3. N_{MMR} still vastly outperforms each N_{ML} and N_{uni}, but N_{uni} now outperforms N_{ML}. The relative improvement of N_{uni} is not surprising, as higher variances, and wider ranges in variance, creates greater potential for N_{ML} to perform poorly. The improvement of N_{MMR} over N_{uni} is still substantial, marking a 25.9% reduction in regret at the 99.9th quantile. This is understandably lower than the improvement in Example A, given the increased difficulty of allocating samples in a high variance scenario.

C. Example: Near Uniformity in Means, Moderate Variance

Example B motivates further exploration into N_{uni} as a viable strategy, and thus this example considers mean uncertainty sets that are close to identical. In this case, both N_{MMR} and N_{uni} lead to substantially smaller maximal regrets than N_{ML}. However, the maximal regret for N_{MMR} and N_{uni} is indistinguishable, as seen in Fig. 4. This does not invalidate the methodology of this article but rather reflects that uniform sampling is a good strategy when the designs are hard to distinguish. Consider, for instance, that the following realized means and variances are possible:

$$\mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5 = \mu_6 = 95$$
$$\sigma_1 = \sigma_2 = \sigma_3 = \sigma_4 = \sigma_5 = \sigma_6 \in [10, 20].$$

Thus, given the information provided by SMEs, it is possible the designs are all identical in mean and variance.

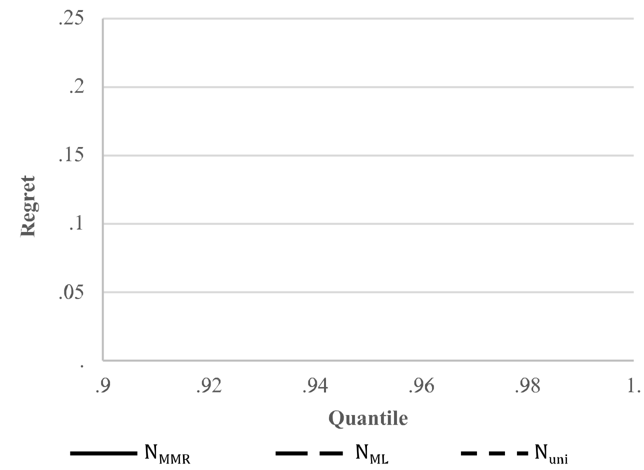

Fig. 5. Regret quantiles for Example D.

D. Example: Moderate Overlap in Means, While Higher/Lower Means Have Higher Variances

This final example is an interesting case, as a decision maker is incentivized to devote samples to designs with lower mean uncertainty bounds because those designs also have higher variances (refer to Table I). Not allocating samples to designs with high mean uncertainty bounds is risky, as these also have high variances. As Fig. 5 shows, the approximation method still outperforms each of the benchmarks. At the 99.9th quantile, N_{MMR} reduces regret by 30.6% compared to the next-best benchmark. The actual solutions N_{MMR} and N_{ML} are worth noting in this example: $N_{\text{MMR}} = [35, 21, 15, 7, 9, 12]$ and $N_{\text{ML}} = [51, 34, 5, 2, 3, 4]$. Each allocates the most samples to design 1, and successively less to designs 2–4, before allocating successively more to designs 5 and 6, reflecting the higher variance of these designs. N_{MMR}, however, is closer to uniform sampling, which evidently improves maximal regret while still being better than pure uniformity.

These same four examples were repeated with $k = 12$ and $k = 24$, expanding the sampling budgets to $B = 200$ and $B = 400$, respectively. The sampling budget was increased with k so that uniform sampling would still allocate a reasonable number of samples to each design. The uncertainty sets for designs 7–24 are computed by naturally extending the pattern used when $k = 6$. That is, mean upper and lower bounds continue declining at the same rates for successive designs. For the case when variances are higher for the highest and lowest means, the first and last designs have $\sigma^u = 20$ and $\sigma^l = 10$, while the middle designs have $\sigma^u = 8$ and $\sigma^l = 4$.

For example, when $k = 12$, and moderate overlap in mean uncertainty sets and variances that are higher for the highest and lowest means are used, the uncertainty sets are

$$\begin{aligned}\mu^u &= 105, 95, 90, \ldots, 55, 50\\ \mu^l &= 95, 85, 80, \ldots, 45, 40\\ \sigma^u &= 20, 17.6, 15.2, \ldots, 8, 8, 10.4, 12.8, \ldots, 20\\ \sigma^l &= 10, 8.8, 7.6, \ldots, 4, 4, 5.2, 6.4, \ldots, 20.\end{aligned}$$

The results are summarized in Table II, where regret quantiles are given for the 95th, the 99th, and the 99.9th quantiles. Continued outperformance of the benchmarks is seen for $k = 12$, but performance deteriorates when $k = 24$. When $k = 24$, the response surface methodology outperforms the benchmarks in only two of the four examples. This can be explained by the increase in response surface error, as seen in Table III. Implementing more elaborate response surfaces with less error (which consequently are expected to be more difficult to optimize) is left as a key point of future research. For the time being, this article has provided evidence that, if a response surface can be found with suitably small errors, then the response surface methodology for multilevel optimizations is a powerful decision-making tool.

V. Conclusion and Future Research

This article developed a robust sampling budget allocation policy via an MMR formulation, solved effectively using a novel response surface-based approximation approach. For designers of complex systems who often face deep uncertainty when exploring alternative design concepts, the developed procedure may provide valuable guidance on how to assign a limited sampling budget to select the best design available. Very little research has previously been done on robust solutions to sampling allocation problems, and to the best of our knowledge, this is the first article to utilize a minimax regret approach, which is more difficult than a standard max–min RO. The analysis was motivated by the need in complex system's design to allocate simulation sampling budgets intelligently, where information prior to running any simulations is often limited and thus robust solutions are valuable. The methodology has applicability beyond sampling allocation problems or even RO; it can be used more generally for any multilevel optimization.

There are several promising topics to study in the future. One such area is the incorporation of more elaborate response surfaces into the methodology. The response surfaces in this article were transformed linear regressions, which had two desirable properties. First, they were easy to optimize, which is an essential property of the lower level optimizations if one is to estimate the higher levels in a reasonable amount of time. Fast global optimization, in general, is a rich area of research, but specialized techniques to quickly optimize the more sophisticated response surfaces (such as neural networks, boosted regressions, and so on) will enhance the usefulness of response surfaces for multilevel optimization. The second useful property of the response surfaces used in this article is they exhibited no heteroskedasticity. This implied even risk-averse decision makers could ignore error terms in the response surface approximation to MMR. If errors instead depend on the decision variables, it would be straightforward enough to simply penalize decisions that introduce high error. However, a more sophisticated approach would be able to select samples from the decision space such that response surface errors are either: 1) uniform or 2) large only in sections of the decision space known to provide bad solutions.

Related to the desire for optimizations that can be solved quickly is the need to get a representative sample of the decision space with as few sample values as possible. This article used Latin hypercube samples for $\{\mu, \sigma\}$, as this is a well-developed method for efficiently sampling from a box set. When sampling from the simplex decision space for N, traditional LHS did not apply and the unbiased, though inefficient, method of drawing Monte Carlo samples was used. Extending LHS to constrained sample spaces is an open research area that will be important for improving the efficiency of the methodology presented in this article.

References

[1] D. Banks, F. Petralia, and S. Wang, "Adversarial risk analysis: Borel games," *Appl. Stochastic Models Bus. Ind.*, vol. 27, no. 2, pp. 72–86, 2011.

[2] O. Ben-Ayed and C. E. Blair, "Computational difficulties of bilevel linear programming," *Oper. Res.*, vol. 38, no. 3, pp. 556–560, 1990.

[3] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization* (Princeton Series in Applied Mathematics). Princeton, NJ, USA: Princeton Univ. Press, 2009.

[4] C. E. Bonferroni, *Teoria Statistica Delle Classi e Calcolo Delle Probabilita*, Florence, Italy: Seeber, 1936, pp. 1–62.

[5] G. G. Brown, W. M. Carlyle, and R. K. Wood, *Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change*. Washington, DC, USA: Nat. Acad. Press, 2008. [Online]. Available: https://doi.org/10.17226/12206

[6] C.-H. Chen, D. He, M. Fu, and L. H. Lee, "Efficient simulation budget allocation for selecting an optimal subset," *INFORMS J. Comput.*, vol. 20, no. 4, pp. 579–595, 2008.

[7] C.-H. Chen, J. Lin, E. Yücesan, and S. E. Chick, "Simulation budget allocation for further enhancing the efficiency of ordinal optimization," *Discr. Event Dyn. Syst. Theory Appl.*, vol. 10, pp. 251–270, Jul. 2000.

[8] C.-H. Chen and L. H. Lee, *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*. Singapore: World Sci., 2010.

[9] P. I. Frazier, W. B. Powell, and S. Dayanik, "A knowledge-gradient policy for sequential information collection," *SIAM J. Control Optim.*, vol. 47, no. 5, pp. 2410–2439, 2008.

[10] P. Frazier, W. Powell, and S. Dayanik, "The knowledge-gradient policy for correlated normal beliefs," *INFORMS J. Comput.*, vol. 21, no. 4, pp. 599–613, 2009.

[11] M. C. Fu, *Handbook of Simulation Optimization* (International Series in Operations Research & Management Science), vol. 216. New York, NY, USA: Springer, 2015.

[12] S. Gao, H. Xiao, E. Zhou, and W. Chen, "Robust ranking and selection with optimal computing budget allocation," *Automatica*, vol. 81, pp. 30–36, Jul. 2017. [Online]. Available: https://doi.org/10.1016/j.automatica.2017.03.019

[13] S.-H. Kim and B. L. Nelson, "Selecting the best system," in *Handbooks in Operations Research and Management Science: Simulation*, S. G. Henderson and B. L. Nelson, Eds. New York, NY, USA: Elsevier, 2006.

[14] R. J. Lempert, S. W. Popper, and S. C. Bankes, *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. Santa Monica, CA, USA : RAND Pardee Center, 2003.

[15] J. Lu, J. Han, Y. Hu, and G. Zhang, "Multilevel decision-making: A survey," *Inf. Sci.*, vols. 346–347, pp. 463–487, Jun. 2016.

[16] V. L. Mulder, S. de Bruin, and M. E. Schaepman, "Representing major soil variability at regional scale by constrained Latin Hypercube Sampling of remote sensing data," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 21, pp. 301–310, Apr. 2013.

[17] Y. Peng and M. C. Fu, "Myopic allocation policy with asymptotically optimal sampling rate," *IEEE Trans. Autom. Control*, vol. 62, no. 4, pp. 2041–2047, Apr. 2017.

[18] G. Perakis and G. Roels, "Regret in the newsvendor model with partial information," *Oper. Res.*, vol. 56, no. 1, pp. 188–203, 2008. [Online]. Available: https://doi.org/10.1287/opre.1070.0486

[19] M. Petelet, B. Iooss, O. Asserin, and A. Loredo, "Latin hypercube sampling with inequality constraints," *AStA Adv. Stat. Anal.*, vol. 94, no. 4, pp. 325–339, Dec. 2010. [Online]. Available: https://doi.org/10.1007/s10182-010-0144-z

[20] P. Roudier, D. E. Beaudette, and A. Hewitt, "A conditioned Latin hypercube sampling algorithm incorporating operational constraints," in *Digital Soil Assessments and Beyond*. Boca Raton, FL, USA: CRC Press, 2012, pp. 227–231.

[21] I. O. Ryzhov, W. B. Powell, and P. I. Frazier, "The knowledge gradient algorithm for a general class of online learning problems," *Oper. Res.*, vol. 60, no. 1, pp. 180–195, 2012.

[22] H.-S. Shih, Y.-J. Lai, and E. S. Lee, "Fuzzy approach for multi-level programming problems," *Comput. Oper. Res.*, vol. 23, no. 1, pp. 73–91, Jan. 1996. [Online]. Available: https://doi.org/10.1016/0305-0548(95)00007-9

[23] H. von Stackelberg, D. Bazin, L. Urch, and R. Hill, *Market Structure and Equilibrium*. Heidelberg, Germany: Springer, 2011.

[24] J. Ungredda, M. Pearce, and J. Branke, "Bayesian optimisation vs. input uncertainty reduction," 2020, *arxiv:2006.00643*.

[25] H. White, "A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity," *Econometrica J. Econometric Soc.*, vol. 48, no. 4, pp. 817–838, 1980.

[26] D. Wu, H. Zhu, and E. Zhou, "A Bayesian risk approach to data-driven stochastic optimization: Formulations and asymptotics," 2018, *arXiv:1609.08665*.

[27] J. Xu, E. Huang, L. Hsieh, L. H. Lee, Q. S. Jia, and C. H. Chen, "Simulation optimization in the era of industrial 4.0 and the industrial Internet," *J. Simulat.*, vol. 10, no. 4, pp. 310–320, Nov. 2016.

[28] J. Xu, E. Huang, C. H. Chen, and L. H. Lee, "Simulation optimization: A review and exploration in the new era of cloud computing and big data," *Asia-Pacific J. Oper. Res.*, vol. 32, no. 3, 2015, Art. no. 1550019.

[29] S. Zhang, L.-H. Lee, E.-P. Chew, J. Xu, and C.-H. Chen, "A simulation budget allocation procedure for enhancing the efficiency of optimal subset selection," *IEEE Trans. Autom. Control*, vol. 61, no. 1, pp. 62–75, Jan. 2016.

[30] S. Zhang, J. Xu, L.-H. Lee, E.-P. Chew, E.-P. Wong, and C.-H. Chen, "Optimal computing budget allocation for particle swarm optimization in stochastic optimization," *IEEE Trans. Evol. Comput.*, vol. 21, no. 2, pp. 206–219, Apr. 2017.

[31] H. Zhu, T. Liu, and E. Zhou, "Risk quantification in stochastic simulation under input uncertainty," *ACM Trans. Model. Comput. Simulat.*, vol. 30, no. 1, pp. 1–24, 2020.

Michael Perry received the B.A. degree in mathematics from the University of British Columbia, Vancouver, BC, Canada, in 2007, the master's degree in financial engineering from the University of California at Berkeley, Berkeley, CA, USA, in 2015, the M.A. degree in defense and strategic studies from the U.S. Naval War College, Newport, RI, USA, in 2020, and the Ph.D. degree in systems engineering and operations research from George Mason University, Fairfax, VA, USA, in 2021.

His research interests include game theory, sampling-based optimization, and decision analysis.

Jie Xu (Senior Member, IEEE) received the B.S. degree in electrical engineering from Nanjing University, Nanjing, China, in 1999, the M.E. degree in electrical engineering from Shanghai Jiaotong University, Shanghai, China, in 2002, the M.S. degree in computer science from SUNY-Buffalo, Buffalo, NY, USA, in 2004, and the Ph.D. degree in industrial engineering and management sciences from Northwestern University, Evanston, IL, USA, in 2009.

He is an Associate Professor of Systems Engineering and Operations Research with George Mason University, Fairfax, VA, USA. His main research interests include stochastic simulation and optimization and their applications.

Edward Huang received the B.S. degree in industrial engineering from National Tsinghua University, Hsinchu, Taiwan, in 2001, and the Ph.D. degree in industrial and systems engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2011.

He is an Associate Professor with the Department of Systems Engineering and Operations Research, George Mason University, Fairfax, VA, USA. His current research interests include planning and scheduling, robust system design, and facility design.

Dr. Huang has served as an Associate Editor for *Advanced Engineering Informatics* and *Asia–Pacific Journal of Operational Research*.

Chun-Hung Chen (Fellow, IEEE) received the Ph.D. degree in decision and control from Harvard University, Cambridge, MA, USA, in 1994.

He is a Professor of Systems Engineering and Operations Research with George Mason University (GMU), Fairfax, VA, USA. Before joining GMU, he was an Assistant Professor with the University of Pennsylvania, Philadelphia, PA, USA. He was also a Professor of Electrical Engineering and Industrial Engineering with National Taiwan University, Taipei, Taiwan, from 2011 to 2014. He is the author of two books, including a best seller: *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation* (World Scientific, 2011).

Prof. Chen has served as a Department Editor for *IIE Transactions* and *Asia–Pacific Journal of Operational Research*, an Associate Editor for IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING, IEEE TRANSACTIONS ON AUTOMATIC CONTROL, and *Journal of Systems Science and Systems Engineering*, an Area Editor for *Journal of Simulation Modeling Practice and Theory*, and an Advisory Editor for *International Journal of Simulation and Process Modeling* and *Journal of Traffic and Transportation Engineering*.