

A DYNAMIC FRAMEWORK FOR STATISTICAL SELECTION PROBLEMS

Yijie Peng

School of Management
Fudan University
Shanghai, 200433, CHINA

Chun-Hung Chen

Systems Engineering and Operations Research
George Mason University
Fairfax, VA 22030, USA

Michael Fu

Robert H. Smith School of Business
Institute for Systems Research
University of Maryland
College Park, MD 20740, USA

Jianqiang Hu

School of Management
Fudan University
Shanghai, 200433, CHINA

ABSTRACT

For the statistical selection problem we formulate a general framework comprising both sequential sampling allocation and optimal design selection, for which the traditional probability of correct selection measure is inadequate. Therefore, we introduce the integrated probability of correct selection to better characterize the objective. In this framework, the usual selection policy of choosing the design with the largest sample mean as the estimate of the best is no longer optimal. Rather, the optimal selection policy is to choose the design that maximizes the posterior integrated probability of correct selection, which is a function of both the posterior mean and the correlation structure induced by the posterior variance.

1 INTRODUCTION

Statistical ranking and selection (R&S) aims to select the best of a finite set of alternatives, where best is determined with respect to the largest mean, and the mean is inferred via statistical sampling (Bechhofer et al. 1995). By analyzing the least favorable configuration, R&S can guarantee a predetermined lower bound on the probability of correct selection (PCS) (Kelton and Law 2000). The downside of R&S is that it usually allocates more replications than necessary to guarantee the PCS.

Alternatively, optimal computing budget allocation (OCBA) (Chen et al. 2000) maximizes the PCS subject to a finite amount of simulation budget, which usually leads to a substantially higher PCS for a specified simulation budget, or achieves the same level of PCS using a smaller simulation budget. OCBA assumes the parameters of the underlying designs are known, and relaxes the objective function with the Bonferroni inequality. In practice, sample statistics are used to get an initial estimate for the parameters. OCBA can also be understood in the Bayesian framework (Branke et al. 2007). Then, how to interpret the meaning of the prior distribution, and how it is related to the original framework are worthwhile questions. To construct a coherent theoretical framework is the first motivation of our work.

The optimal sampling allocation method studied in Bayesian framework is still formulated as a parametric optimization problem, not accounting for sequential allocation. Chen et al. (2006) found that for some special cases, the OCBA algorithm with perfect information may be far inferior to that of the sequential

OCBA with non-perfect information, implying that the sequential mechanism is critical. To incorporate the sequential mechanism in the theoretical framework is the second motivation of our work.

In this work, we propose a more general framework, incorporating the sequential mechanism. We define a sampling allocation and design selection (SADS) policy which comprises an allocation policy that decides allocation to which design sequentially, and a selection policy that decides the design selected as the “best” after exhausting the simulation replications.

In the optimal sample allocation literature, the selection policy is assumed to choose the design that has largest sample mean, but in our framework it is a part of the decision policy. With this generalization, we can easily show that there is no dominant policy. Therefore, we need to trade off the PCSs under different parameter settings of the underlying designs, and make a subjective definition of the “optimal” policy. One way to do this is to introduce a prior distribution on the space of the value set of the parameter, fix a policy and integrate the PCSs under all possible parameter settings, and define the value as integrated PCS (IPCS) under this policy. Then the optimal policy is defined as the policy that maximizes the IPCS.

Under a special case that the variances are assumed to be known, we show that the posterior variances, as well as the posterior means, influence the posterior IPCS, because of the induced correlation structure, and demonstrate that the policy that chooses the design with highest sample mean is not necessarily optimal. Since the optimal selection policy is the decision made in the last step, it influences the allocation policy, but isn’t affected by the latter.

The optimal allocation policy turns out to be the solution of a difficult stochastic dynamic programming (SDP) problem, for which simulation seems to be the only tractable approach for finding the optimal allocation policy. To avoid this time-consuming procedure, we propose and test a myopic suboptimal allocation policy utilizing the knowledge gradient (KG) studied in Frazier et al. (2008).

The rest of the paper is organized as follows. In Section 2, we define the optimization problem for the optimal sampling allocation and design selection policy, and review some results for the conjugate prior distribution. Section 3 investigates the structure of the optimal selection policy. The myopic suboptimal allocation policy is studied in Section 4. Numerical results are provided in Section 5 to test the performance of the proposed algorithms. The last section offers conclusions.

2 PROBLEM FORMULATION

We have k designs, the performance means μ_i , $i = 1, \dots, k$, of which are unknown. We can only observe the performance of design i by random experiments. The observations satisfy $(x_{1,l}, \dots, x_{k,l}) \sim Q(\cdot; \theta)$, i.i.d., $l \in \mathbb{Z}^+$, where $\theta \in \Theta$ comprise of all parameters of parametric family $Q(\cdot; \theta)$. The total simulation budget is T . Our objective is to select the best design from these k competing designs, based on the information collected from T samples allocated to the alternative designs.

We formulate this statistical selection problem as a SADS policy. The allocation policy is a sequence of maps $\mathcal{A} = (a_1(\cdot), \dots, a_T(\cdot))$, $a_s(\cdot) : \mathcal{E}_s \mapsto \{1, \dots, k\}$, \mathcal{E}_s is all the information collected in step s . The selection policy is a map $\mathcal{D}(\cdot) : \mathcal{E}_T \mapsto \{1, \dots, k\}$. For fixed $\{\mu_1, \dots, \mu_k\} \subset \theta$, define $\mu_{[1]} > \dots > \mu_{[k]}$, where $[\cdot]$ denote the order statistics. The PCS is defined as

$$PCS(\mathcal{A}, \mathcal{D}, \theta) \triangleq Pr(\{\mathcal{D}(\mathcal{E}_T^{\mathcal{A}}) = [1]\} | \theta),$$

which is the probability that, with specific policy $(\mathcal{A}, \mathcal{D})$ and the underlying parameters given, after depleting all replications, the best design is successfully selected. Obviously, the distribution of \mathcal{E}_T is dependent on \mathcal{A} ; therefore, we write $\mathcal{E}_T^{\mathcal{A}}$ to specify this dependence. The brute-force equal allocation (EA) and OCBA are special cases with selection policy $\mathcal{D}^m \triangleq \arg \max_{i=1, \dots, k} \bar{m}_i^{(T)}$ and static allocation policy determined at the beginning.

To define an “optimal” policy, a natural definition is a policy $(\mathcal{A}^*, \mathcal{D}^*)$ that dominates the rest of policies $(\mathcal{A}, \mathcal{D})$ in PCS, which means $\forall \theta$,

$$Pr(\{\mathcal{D}^*(\mathcal{E}_T^{\mathcal{A}^*}) = [1]\} | \theta) \geq Pr(\{\mathcal{D}(\mathcal{E}_T^{\mathcal{A}}) = [1]\} | \theta).$$

Unfortunately, this “optimal” policy doesn’t exist, assuming that there exist $\theta, \theta' \in \Theta$ such that $\{\mu_1, \dots, \mu_k\} \subset \theta, \mu_{[1]} > \dots > \mu_{[k]}, \{\mu'_1, \dots, \mu'_k\} \subset \theta', \mu'_{(1)} > \dots > \mu'_{(k)}$, and $[1] \neq (1)$, i.e., there exist at least two configurations of parameter settings in question such that the indices of their respective best design are different.

Since there is no policy that can dominate the rest of the policies, we need to define a criterion under which there exists an “optimal” policy in some sense. The tradeoff of PCSs under different true parameters should be specified. To do this, we can introduce the prior distribution of parameters $P(\theta), \theta \in \Theta$. For practitioners, this prior distribution can be understood as the information gathered from previous study, past experience, expert opinion and individual preference. After specifying the prior, the “optimal” policy can be defined by the solution to a stochastic dynamic programming problem (SDP):

$$(\mathcal{A}^*, \mathcal{D}^*) \triangleq \arg \max_{\mathcal{A}, \mathcal{D}} E[\mathbf{1}\{\mathcal{D}(\mathcal{E}_T^{\mathcal{A}}) = [1]\}] \quad (1)$$

where

$$\begin{aligned} E[\mathbf{1}\{\mathcal{D}(\mathcal{E}_T^{\mathcal{A}}) = [1]\}] &= E \left[E[\mathbf{1}\{\mathcal{D}(\mathcal{E}_T^{\mathcal{A}}) = [1]\} \mid \theta] \right] \\ &= \int_{\theta \in \Theta} Pr(\{\mathcal{D}(\mathcal{E}_T^{\mathcal{A}}) = [1]\} \mid \theta) \cdot P(d\theta). \end{aligned} \quad (2)$$

Because the objective function (2) is the weighted average of the PCSs under different parameters, we call it the integrated PCS (IPCS).

Additionally, we denote $\mathcal{A}_s \triangleq (a_1(\cdot), \dots, a_s(\cdot))$, $\mathcal{A}_{s+} \triangleq (a_{s+1}(\cdot), \dots, a_T(\cdot))$, $\mathcal{E}_{s+} \triangleq \mathcal{E}_T \setminus \mathcal{E}_s$, and $\xi_s \triangleq \mathcal{E}_s \setminus \mathcal{E}_{s-1}$, $s = 0, \dots, T$, where $\mathcal{A}_{T+} = \emptyset$, $\mathcal{E}_{-1} = \emptyset$, and \mathcal{E}_0 is the information set before allocating any simulation replication, which can be an empty set or comprises of some prior information. Obviously, $\mathcal{A} = \mathcal{A}_T = \mathcal{A}_{0+} = \mathcal{A}_s \cup \mathcal{A}_{s+}$ and $\mathcal{E}_T = \mathcal{E}_{0+} = \mathcal{E}_s \cup \mathcal{E}_{s+}$. The distributions of \mathcal{E}_s , \mathcal{E}_{s+} , and ξ_s are dependent on \mathcal{A}_s , \mathcal{A}_{s+} , and a_s respectively. We write $\mathcal{E}_s^{\mathcal{A}_s}$, $\mathcal{E}_{s+}^{\mathcal{A}_{s+}}$, and $\xi_s^{a_s}$ to specify this dependence. With the additional notations, we can also write down the corresponding Bellman equation of (1), $s = 0, \dots, T-1$,

$$E[\mathbf{1}\{\mathcal{D}^*(\mathcal{E}_s, \mathcal{E}_{s+}^{\mathcal{A}_{s+}}) = [1]\} \mid \mathcal{E}_s] = \max_{a_{s+1}} E_{\xi_s^{a_{s+1}}} [E[\mathbf{1}\{\mathcal{D}^*(\mathcal{E}_{s+1}, \mathcal{E}_{(s+1)+}^{\mathcal{A}_{(s+1)+}}) = [1]\} \mid \mathcal{E}_{s+1}]], \quad (3)$$

and we call the conditional expectation on the left hand side the posterior IPCS in step s . Since the reward function $E[\mathbf{1}\{\mathcal{D}^*(\mathcal{E}_s, \mathcal{E}_{s+}^{\mathcal{A}_{s+}}) = [1]\} \mid \mathcal{E}_T]$ in the last step is bounded and the decision variable $a_s(\cdot)$ only takes finite possible values, the Bellman equation (3) has a solution that is also the solution of the SDP problem (1).

The case that is easiest to understand is when $P(\theta)$ is uniformly distributed on a compact set Θ , in which we impose equal weights on the PCS under different parameters in the compact set. It can also be understood as a uninformative prior on a compact set. In this prior, $IPCS = 1/k$ for the trivial policy $(\mathcal{A}, \mathcal{D}_{tr})$.

The case considered henceforth is when $P(\theta)$ is the conjugate prior of distribution $Q(\cdot; \theta)$. In this case, the posterior distribution of θ falls into the same parametric family as the conjugate prior, and the parameters in the posterior distribution can be updated easily.

Assume $x_{i,l} \sim N(\mu_i, \sigma_i^2)$, i.i.d., $l \in \mathbb{Z}^+$, $i = 1, \dots, k$, and $\forall i, j, l, l', x_{il}$ and $x_{j'l'}$ are mutually independent, in this case, $\theta = (\mu_i, \sigma_i^2)_{i=1}^k$, and the conjugate prior is the normal-Gamma, $i = 1, \dots, k$,

$$NG(\mu_i, \lambda_i \mid \mu_i^{(0)}, \kappa_i^{(0)}, \alpha_i^{(0)}, \beta_i^{(0)}) \triangleq N(\mu_i \mid \mu_i^{(0)}, (\kappa_i^{(0)} \lambda_i)^{-1}) \Gamma(\lambda_i \mid \alpha_i^{(0)}, \beta_i^{(0)}), \quad (4)$$

where $\lambda_i \triangleq \sigma_i^{-2}$ which is called the precision. Parameters in prior distributions are called hyper-parameters, to distinguish them from parameters in the distribution of the underlying designs. The sufficient statistics

for a normal distribution are the sample mean and variance defined as

$$\begin{aligned}\bar{m}_i^{n_i} &\triangleq \frac{\sum_{l=1}^{n_i} x_{i,l}}{n_i}, \\ \bar{v}_i^{n_i} &\triangleq \frac{\sum_{l=1}^{n_i} (x_{i,l} - \bar{m}_i^{n_i})^2}{n_i}.\end{aligned}$$

The information collected in step $s = \sum_{i=1}^k n_i$ is $\mathcal{E}_s = (\bar{m}_1^{n_1}, \dots, \bar{m}_k^{n_k}, \bar{v}_1^{n_1}, \dots, \bar{v}_k^{n_k})$. We can also define $(s)_i \triangleq n_i$, $\bar{m}_i^{(s)} \triangleq \bar{m}_i^{n_i}$ and $\bar{v}_i^{(s)} \triangleq \bar{v}_i^{n_i}$, $i = 1, \dots, k$, which are the replication number, and sample mean and variance for design i after allocation of s simulation replications. Under this framework, from Berger (2004), we know the posterior distribution of (μ_i, λ_i) is also normal-Gamma, $i = 1, \dots, k$,

$$NG(\mu_i, \lambda_i | \mu_i^{(s)}, \kappa_i^{(s)}, \alpha_i^{(s)}, \beta_i^{(s)}),$$

where

$$\begin{aligned}\mu_i^{(s)} &\triangleq \frac{\kappa_i^{(0)} \mu_i^{(0)} + (s)_i \bar{m}_i^{(s)}}{\kappa_i^{(0)} + (s)_i}, \\ \kappa_i^{(s)} &\triangleq \kappa_i^{(0)} + (s)_i, \\ \alpha_i^{(s)} &\triangleq \alpha_i^{(0)} + (s)_i / 2, \\ \beta_i^{(s)} &\triangleq \beta_i^{(0)} + \frac{(s)_i \bar{v}_i^{(s)}}{2} + \frac{\kappa_i^{(0)} (s)_i (\bar{m}_i^{(s)} - \mu_i^{(0)})^2}{2(\kappa_i^{(0)} + (s)_i)}.\end{aligned}\tag{5}$$

We also know the posterior marginal distribution of μ_i is a t -distribution, $i = 1, \dots, k$,

$$T_{2\alpha_i^{(s)}}(\cdot | \mu_i^{(s)}, \vartheta_i^{(s)}),\tag{6}$$

where

$$\vartheta_i^{(s)} \triangleq \frac{\beta_i^{(s)}}{\alpha_i^{(s)} \kappa_i^{(s)}},$$

and the posterior predictive distribution of $x_{i,l}$, $l > (s)_i$ is a t -distribution, $i = 1, \dots, k$,

$$T_{2\alpha_i^{(s)}}(\cdot | \mu_i^{(s)}, \mathbf{v}_i^{(s)}),\tag{7}$$

where

$$\mathbf{v}_i^{(s)} \triangleq \frac{\beta_i^{(s)} (\kappa_i^{(s)} + 1)}{\alpha_i^{(s)} \kappa_i^{(s)}}.$$

In this case, we have

$$E[\mathbf{1}\{\mathcal{D}(\mathcal{E}_s, \mathcal{E}_{s^+}^{\mathcal{A}}) = [1]\} | \mathcal{E}_s] = \int_{\theta \in \Theta} Pr(\{\mathcal{D}(\mathcal{E}_s, \mathcal{E}_{s^+}^{\mathcal{A}}) = [1]\} | \theta) \cdot P(d\theta | \mathcal{E}_s),$$

where

$$P(\theta | \mathcal{E}_s) = \prod_{i=1}^k NG(\mu_i, \lambda_i | \mu_i^{(s)}, \kappa_i^{(s)}, \alpha_i^{(s)}, \beta_i^{(s)}).$$

We then illustrate the meaning of each input parameters in the prior distribution. From (4), the prior distribution of the inverse of true variance σ_i^{-2} is a gamma distribution with hyper-parameters $\alpha_i^{(0)}$ and $\beta_i^{(0)}$. The expectation of σ_i^2 exists when $\alpha_i^{(0)} > 1$, and its variance exists when $\alpha_i^{(0)} > 2$. For $\alpha_i^{(0)} > 2$, we have

$$\begin{aligned} E[\sigma_i^2] &= \frac{\beta_i^{(0)}}{\alpha_i^{(0)} - 1}, \\ \text{Var}(\sigma_i^2) &= \frac{(\beta_i^{(0)})^2}{(\alpha_i^{(0)} - 2)(\alpha_i^{(0)} - 1)^2}. \end{aligned} \quad (8)$$

From (4), the prior distribution of the true mean μ_i conditional on σ_i^2 is a normal distribution with mean $\mu_i^{(0)}$ and variance $(\kappa_i^{(0)})^{-1}\sigma_i^{-2}$. Therefore, we have

$$\begin{aligned} E[\mu_i] &= E[E[\mu_i|\sigma_i^2]] = \mu_i^{(0)}, \\ \text{Var}(\mu_i) &= E[\text{Var}(\mu_i|\sigma_i^2)] + \text{Var}(E[\mu_i|\sigma_i^2]), \end{aligned} \quad (9)$$

where

$$\begin{aligned} E[\text{Var}(\mu_i|\sigma_i^2)] &= E[(\kappa_i^{(0)})^{-1}\sigma_i^2] = \frac{\beta_i^{(0)}}{\kappa_i^{(0)}(\alpha_i^{(0)} - 1)}, \\ \text{Var}(E[\mu_i|\sigma_i^2]) &= \text{Var}(\mu_i^{(0)}) = 0. \end{aligned}$$

With these relationships, we can understand the meaning of the hyper-parameters in the prior distribution.

A special case is when the variance vector $(\sigma_1^2, \dots, \sigma_k^2)$ is known. Then the only unknown parameter is $\theta = (\mu_1, \dots, \mu_k)$. In this case, the conjugate prior $P(\theta)$ is the normal distribution, $i = 1, \dots, k$,

$$N(\mu_i|\mu_i^{(0)}, (\sigma_i^{(0)})^2).$$

The posterior distribution of μ_i is also a normal distribution, $i = 1, \dots, k$,

$$N(\mu_i|\mu_i^{(T)}, (\sigma_i^{(T)})^2), \quad (10)$$

where

$$\begin{aligned} \mu_i^{(T)} &\triangleq (\sigma_i^{(T)})^2 \left(\frac{\mu_i^{(0)}}{(\sigma_i^{(0)})^2} + \frac{(T)_i \bar{m}_i^{(T)}}{\sigma_i^2} \right), \\ (\sigma_i^{(T)})^2 &\triangleq \frac{\sigma_i^2 (\sigma_i^{(0)})^2}{(T)_i (\sigma_i^{(0)})^2 + \sigma_i^2}. \end{aligned} \quad (11)$$

3 OPTIMAL SELECTION POLICY

Now we consider the optimal selection policy. Suppose that \mathcal{A}^* is found, and all the replications are exhausted following this allocation policy. Based on the collected information \mathcal{E}_T , we make the final selection decision \mathcal{D}^* based on \mathcal{E}_T , which \mathcal{D}^* maximize the posterior IPCS in step T :

$$\mathcal{D}^* = \arg \max_{\mathcal{D}} E[\mathbf{1}\{\mathcal{D}(\mathcal{E}_T) = [1]\}|\mathcal{E}_T], \quad (12)$$

where

$$\begin{aligned} E[\mathbf{1}\{\mathcal{D}(\mathcal{E}_T) = [1]\}|\mathcal{E}_T] &= \int_{\theta \in \Theta} \mathbf{1}\{\mathcal{D}(\mathcal{E}_T) = [1]\} \cdot P(d\theta|\mathcal{E}_T) \\ &= \text{Pr}(\mu_{\mathcal{D}} - \mu_j > 0, j \neq \mathcal{D}|\mathcal{E}_T). \end{aligned}$$

In the literature of optimal sampling allocation, the selection policy is assumed to be $\mathcal{D}^m = \arg \max_{i=1, \dots, k} \bar{m}_i^{(T)}$. Is \mathcal{D}^m the solution of (12)? If not, what other factors determine the optimal selection policy? From the large sample perspective, \mathcal{D}^m can be approximately the optimal selection policy, as the following proposition shows.

Lemma 1 If $x_{i,l} \sim N(\mu_i, \sigma_i^2)$, i.i.d., $l \in \mathbb{Z}^+$, $i = 1, \dots, k$, and $\theta = (\mu_i, \sigma_i^2)_{i=1}^k$ follows the normal-Gamma model defined in (4), then

- (i) $\mu_i | \mathcal{E}_T$, $i = 1, 2, \dots, k$, are mutually independent;
- (ii) for $i = 1, 2, \dots, k$, $\forall x$ and any finite $(T)_i$, we have $0 < Pr(\mu_i > x | \mathcal{E}_T) < 1$.

Remark. Statement (i) rules out the possibility that by simulating design i , we can gain information on design j . Statement (ii) basically says the support space of the posterior distribution of the mean of each design wouldn't shrink with only finite many observations.

Proposition 1 Suppose the sample space and parameter space are a subset set of Euclidean space, $(x_{1,l}, x_{2,l}, \dots, x_{k,l}) \sim Q(\cdot; \theta)$, i.i.d. $l \in \mathbb{Z}^+$, and $\theta \sim P(\cdot)$, satisfying conditions: $Q(\cdot; \theta) \neq Q(\cdot; \theta')$ whenever $\theta \neq \theta'$, and (i) (ii) in Lemma 1, then, as $T \rightarrow \infty$, we have

$$\begin{aligned} Pr(\mu_{\mathcal{D}^m} - \mu_j > 0, j \neq \mathcal{D}^m | \mathcal{E}_T^{\mathcal{A}^*}) &\rightarrow 1, \quad \mathcal{Q} - a.e., \\ Pr(\mu_i - \mu_j > 0, j \neq i | \mathcal{E}_T^{\mathcal{A}^*}) &\rightarrow 0, \quad \mathcal{Q} - a.e., \quad i \neq \mathcal{D}^m, \end{aligned}$$

therefore,

$$\mathcal{D}^*(\mathcal{E}_T^{\mathcal{A}^*}) - \mathcal{D}^m(\mathcal{E}_T^{\mathcal{A}^*}) \rightarrow 0, \quad \mathcal{Q} - a.e.,$$

where \mathcal{Q} denotes the distribution of $(\{x_{i,l} : i = 1, \dots, k, l = 1, 2, \dots\}, \theta)$.

Remark. The assumption of conditions (i) and (ii) is only to make sure that, under the optimal allocation policy, the number of replications allocated to each design would approach infinity as the total simulation budget grows to infinity. We believe that this conclusion holds for a wide range of distributions, although to give a more general but simple condition to describe this category of distributions is not an easy task. The identifiability condition is to apply Doob's consistency theorem (Van der Vaart 2000). For the normal-Gamma model introduced in Section 2, the identifiability condition is satisfied, and the consistency of Bayesian posterior measure can be directly obtained by utilizing the the form of posterior distribution of the normal-Gamma model without resorting to the Doob's consistency theorem.

Corollary 1 If $x_{i,l} \sim N(\mu_i, \sigma_i^2)$, i.i.d., $l \in \mathbb{Z}^+$, $i = 1, \dots, k$, and $\theta = (\mu_i, \sigma_i^2)_{i=1}^k$ follows the normal-Gamma model defined in (4), then the conclusion in Proposition 1 holds.

However, from the small sample perspective, \mathcal{D}^m may not be a good approximation to \mathcal{D}^* , and there should be some other factors. To see the structure of this problem clearly, consider a simple case that the variance vector $(\sigma_1^2, \dots, \sigma_k^2)$ is considered as a known value. In this case, $\mathcal{E}_T = (\mu_1^{(T)}, \dots, \mu_k^{(T)}, \sigma_1^{(T)}, \dots, \sigma_k^{(T)}) = (\bar{m}_1^{(T)}, \dots, \bar{m}_k^{(T)}, (T)_1, \dots, (T)_k)$.

As we can see in (11), hyper-parameters $\mu_i^{(0)}$ and $\sigma_i^{(0)}$ in the prior distribution affect the posterior mean and variance $\mu_i^{(T)}, \sigma_i^{(T)}$, $i = 1, \dots, k$. As the following lemmas and propositions show that even under this special case defined above, $\mathcal{D}^{pm} \triangleq \arg \max_{i=1, \dots, k} \mu_i^{(T)}$ is not necessarily the optimal selection policy.

Let's first introduce a notation

$$cor(X, Y) \triangleq \frac{E[(X - EX)(Y - EY)]}{E[(X - EX)^2]^{\frac{1}{2}} E[(Y - EY)^2]^{\frac{1}{2}}}$$

where X, Y are two r.v.'s

Lemma 2 Suppose $y_j \sim N(0, \sigma_j^2)$, $j = 1, \dots, k$, $y_{j'}$, y_j mutually independent, $j \neq j'$, and

$$\begin{aligned} \rho_{j,j'}^{(i)} &\triangleq \text{cor}(y_i - y_{j+1}, y_1 - y_{j'+1}), \quad j, j' = 1, \dots, i-1, \\ \rho_{j,j'}^{(i)} &\triangleq \text{cor}(y_i - y_{j+1}, y_1 - y_{j'+1}), \quad j, j' = i, \dots, k-1. \end{aligned} \quad (13)$$

Then $\text{Pr}(y_i - y_j > 0, j \neq i)$ is fully determined by $\rho_{j,j'}$, $j \neq j'$, $j, j' = 1, \dots, k-1$.

Although it's much more complicated to analyze the joint influence of posterior mean and variance on the IPCS, Lemma 2 shows that, when the posterior mean of each design are equal, the IPCS is determined by the correlation structure induced by the posterior variance of the selected design. The following lemma shows that selecting the design with larger posterior variance will lead to larger correlation between difference pairs.

Lemma 3 Suppose $y_j \sim N(0, \sigma_j^2)$, $j = 1, \dots, k$, $y_{j'}$, y_j mutually independent, $j \neq j'$, and $\rho_{j,j'}^{(i)}$, $i = 1, \dots, k$, $j, j' \neq i$, are defined as (13). If $\sigma_i^2 > \sigma_{i'}^2$, then

$$\rho_{j,j'}^{(i)} > \rho_{j,j'}^{(i')}, \quad j \neq j', \quad j, j' = 1, \dots, k-1.$$

Lemma 4 Suppose $z_i \sim N(0, 1)$, $i = 1, 2$, and $\text{cor}(z_1, z_2) = \rho$ ($0 < \rho < 1$). Then we have $\text{Pr}(z_1 > 0, z_2 > 0)$ is increasing with respect to ρ .

Remark. Actually, the quadrant probability can be given analytically by

$$\text{Pr}(z_1 > 0, z_2 > 0) = \frac{1}{4} + \frac{\arcsin \rho}{2\pi}$$

(see Rose and Smith 2002, p. 231). By this result, the conclusion of Lemma 4 is quite obvious.

Lemma 4 proves that for two correlated standard normal distributed random variables, the probability that they have the same sign increases with the correlation value. With the conclusion of this lemma, we can easily show that for three designs, if they have equal posterior means, selecting the design with larger posterior variance will lead to larger IPCS, as it is stated in the statement (i) of Proposition 2. Usually, the posterior mean of different designs wouldn't be equal, and, in this case, the IPCS is a function of both the posterior mean and the correlation structure induced by the posterior variance of the selected design. We can imagine that the optimal selection policy is not just choosing the design with largest posterior mean but a function of both posterior mean and the correlation. The rigorous statement which proves that the policy choosing largest posterior mean is not the optimal selection policy in general is illustrated in Proposition 3.

Proposition 2 Suppose $y_j \sim N(0, \sigma_j^2)$, $j = 1, \dots, k$, y_j , $y_{j'}$ mutually independent, $j \neq j'$.

(i) For $k = 3$, if $\sigma_1^2 > \sigma_2^2 > \sigma_3^2$, then

$$\text{Pr}(y_1 - y_j > 0, j = 2, 3) > \text{Pr}(y_2 - y_j > 0, j = 1, 3) > \text{Pr}(y_3 - y_j > 0, j = 1, 2);$$

(ii) If $\sigma_i^2 \rightarrow \infty$, $i = 1, \dots, k$, then

$$\text{Pr}(y_i - y_j > 0, j \neq 1) \rightarrow \frac{1}{2}, \quad j \neq i$$

and, for $i' \neq i$,

$$\text{Pr}(y_{i'} - y_j > 0, j \neq i') \rightarrow \frac{1}{2} \cdot \text{Pr}(y_i - y_j > 0, j \neq i, i').$$

Proposition 3 Suppose $(\Omega, \mathcal{P}, \mathcal{M})$ is the probability space. Define

$$\Delta \triangleq \{\omega \in \Omega : \sigma_1^{(T)} = \dots = \sigma_k^{(T)}\}, \quad (14)$$

and

$$\mathcal{E}_T(\omega) \triangleq \{(\mu_1^{(T)}(\omega), \dots, \mu_k^{(T)}(\omega), \sigma_1^{(T)}(\omega), \dots, \sigma_k^{(T)}(\omega)) : \omega \in \Omega\}.$$

(i) For $k = 2, \forall \omega \in \Omega$,

$$\mathcal{D}^*(\mathcal{E}_T(\omega)) = \arg \max_{i=1,2} \mu_i^{(T)};$$

(ii) For $k > 2, \forall \omega \in \Delta$,

$$\mathcal{D}^*(\mathcal{E}_T(\omega)) = \arg \max_{i=1,\dots,k} \mu_i^{(T)},$$

and, if $Pr(\Delta^c) > 0$, there exists $\Xi \subset \Delta^c$ and $P(\Xi) > 0$ such that, $\forall \omega \in \Xi$,

$$\mathcal{D}^*(\mathcal{E}_T(\omega)) \neq \arg \max_{i=1,\dots,k} \mu_i^{(T)}.$$

Remark. The probability of the event (14) that all posterior variances are equal depends on the hyper-parameters, the true variances of each design, and the allocation policy \mathcal{A} , as shown from (11). For example, if $\sigma_i^{(0)} = \sigma_j^{(0)}, \sigma_i = \sigma_j, i, j = 1, \dots, k$, and the allocation policy is the equal allocation policy \mathcal{A}^e , we have $Pr(\Delta) = 1$, but, if we let the allocation policy depend also on the samples collected at each step, e.g. OCBA, it is quite possible that $Pr(\Delta) = 0$.

As we can see from the analysis, \mathcal{D}^* generally doesn't equal \mathcal{D}^m , and the correlation structure induced by the posterior variance of the selected design also has big impact on the selection policy, especially from the small sample perspective. As we can see in (ii) of Proposition 2, in this extreme case, due to the induced correlation structure, design i only needs to compete with one design, while alternative design i' needs to compete with all designs. The following example captures the idea behind the conclusions in this section.

Example. Suppose $\mu_1 \sim N(\mu_1^{(T)}, (\sigma_1^{(T)})^2), \mu_2 \sim N(\mu_2^{(T)}, (\sigma_2^{(T)})^2), \mu_3 \sim N(\mu_3^{(T)}, (\sigma_3^{(T)})^2)$, and $(\mu_1^{(T)}, (\sigma_1^{(T)})^2) = (0.1, 1), (\mu_2^{(T)}, (\sigma_2^{(T)})^2) = (0, 10^2), (\mu_3^{(T)}, (\sigma_3^{(T)})^2) = (-0.1, 1)$. Then, we have

$$(\mu_1 - \mu_2, \mu_1 - \mu_3) \sim N\left((0.1, 0.2), \begin{pmatrix} 101 & 1 \\ 1 & 2 \end{pmatrix}\right),$$

and

$$(\mu_2 - \mu_1, \mu_2 - \mu_3) \sim N\left((-0.1, 0.1), \begin{pmatrix} 101 & 100 \\ 100 & 101 \end{pmatrix}\right).$$

By numerical calculation, we have

$$Pr(\mu_1 - \mu_2 > 0, \mu_1 - \mu_3 > 0) = 0.292 < 0.478 = Pr(\mu_2 - \mu_1 > 0, \mu_2 - \mu_3 > 0).$$

Therefore, we should choose design 2 which gives us highest posterior IPCS, although design 1 has largest posterior mean.

Although we only rigorously prove that the increasing the correlation would lead to increasing of the posterior IPCS when $k = 3$ and the designs have equal posterior mean, the numerical experiments often indicate that this conclusion would hold in numerous application examples (see Fu et al. 2007 and Peng et al. 2013). From Lemma 2 and 3, we know the larger the posterior variance of the selected design

the higher the posterior IPCS. Then we can have a rule of thumb that selecting the design with larger posterior mean and variance will lead to higher posterior IPCS.

For this special case defined above, the optimal selection policy \mathcal{D}^* defined by (12) can be easily calculated. We can first rule out the designs that have both the posterior mean and variance smaller than some other design based on the analysis above, which will increase the efficiency of solving (12). For the case when $\sigma_i, i = 1, \dots, k$ are unknown, the posterior distribution should be the t -distribution (6), instead of the normal distribution (10). Because $T_{2\alpha_i^{(r)}}(\cdot) \rightarrow N(\cdot)$ in distribution as $(T)_i \rightarrow \infty$, for $(T)_i$ large enough, all the conclusions made in this section could be understood as some approximate descriptions in the case that $\sigma_i, i = 1, \dots, k$ are unknown, and we also use the normal approximation to solve (12) for the sake of simplicity.

4 MYOPIC ALLOCATION POLICY

A myopic policy is the policy that looks just one step ahead, based on the current information about the distribution. The KG policy is a myopic policy studied in Frazier et al. (2008). It maximizes the value of information for the posterior mean, which basically solves a myopic optimization problem, for $s = 1, \dots, T$, and $i = 1, \dots, k$

$$a_s \in \arg \max_{i=1, \dots, k} E \left[\max \left(\mu_i^{(s)}, \max_{j \neq i} \mu_j^{(s-1)} \right) \mid \mathcal{E}_{s-1}, a_s = i \right], \quad (15)$$

where

$$\mu_i^{(s)} = \frac{\kappa_i^{(s-1)}}{\kappa_i^{(s)}} \mu_i^{(s-1)} + \frac{1}{\kappa_i^{(s)}} x_{i,(s)i}. \quad (16)$$

From (7), we know the predictive distribution of $x_{i,(s)i}$ conditional on \mathcal{E}_{s-1} is a t -distribution

$$T_{2\alpha_i^{(s-1)}} \left(\cdot \mid \mu_i^{(s-1)}, v_i^{(s-1)} \right).$$

To simplify the notation, define $v_i \triangleq 2\alpha_i^{(s-1)}$, for which the dependence on s is omitted. To ensure the expectation in (15) exists, we need $v_i, i = 1, \dots, k$. If we let $\alpha_i^{(0)} > \frac{1}{2}, i = 1, \dots, k$, in the prior distribution, (15) is well defined for all $s = 1, \dots, T$.

As shown in Section 3, both the posterior mean and variance would lead to higher IPCS. Based on this fact, we also want to incorporate posterior variance into the value of information. To do this, we derive the second-order KG policy which maximize the value of information for the linear combination of the posterior mean and variance. The myopic optimization problem for second-order KG policy is, for $s = 1, \dots, T$, and $i = 1, \dots, k$,

$$a_s \in \arg \max_{i=1, \dots, k} E \left[\max \left(\mu_i^{(s)} + \gamma \vartheta_i^{(s)}, \max_{j \neq i} \left\{ \mu_j^{(s-1)} + \gamma \vartheta_j^{(s-1)} \right\} \right) \mid \mathcal{E}_{s-1}, a_s = i \right], \quad (17)$$

where

$$\vartheta_i^{(s)} = \eta_i + \eta'_i x_{i,(s)i} + \eta''_i x_{i,(s)i}^2,$$

and

$$\begin{aligned}\eta_i &\triangleq \frac{(s-1)_i(\bar{m}_i^{(s-1)})^2}{2(s)_i\alpha_i^{(s)}\kappa_i^{(s)}} + \frac{\kappa_i^{(0)}(s-1)_i^2(\bar{m}_i^{(s-1)} - \mu_i^{(0)})^2 + \kappa_i^{(0)}(\mu_i^{(0)})^2}{2(s)_i\alpha_i^{(s)}(\kappa_i^{(s)})^2} \\ &\quad + \frac{\beta_i^{(0)} + \frac{(s-1)_i}{2}\bar{v}_i^{(s-1)}}{\alpha_i^{(s)}\kappa_i^{(s)}} - \frac{\mu_i^{(0)}\kappa_i^{(0)}(s-1)_i(\bar{m}_i^{(s-1)} - \mu_i^{(0)})}{(s)_i\alpha_i^{(s)}(\kappa_i^{(s)})^2}, \\ \eta_i' &\triangleq \frac{\kappa_i^{(0)}(s-1)_i(\bar{m}_i^{(s-1)} - \mu_i^{(0)}) - \kappa_i^{(0)}\mu_i^{(0)}}{(s)_i\alpha_i^{(s)}(\kappa_i^{(s)})^2} - \frac{(s-1)_i\bar{m}_i^{(s-1)}}{(s)_i\alpha_i^{(s)}\kappa_i^{(s)}}, \\ \eta_i'' &\triangleq \frac{\kappa_i^{(0)}}{2(s)_i\alpha_i^{(s)}(\kappa_i^{(s)})^2} + \frac{(s-1)_i}{2(s)_i\alpha_i^{(s)}\kappa_i^{(s)}}.\end{aligned}$$

These terms are known after conditioning on $(\mathcal{E}_{s-1}, a_s = i)$. Let's introduce some additional notations:

$$\zeta_i \triangleq -\frac{\mu_i^{(s-1)} - \max_{j \neq i} \mu_j^{(s-1)}}{\vartheta_i^{(s-1)}},$$

and

$$\begin{aligned}c_i' &\triangleq \max_{j \neq i} \{\mu_j^{(s-1)} + \gamma \vartheta_j^{(s-1)}\}, \\ c_i'' &\triangleq \mu_i^{(s-1)} + \gamma[\eta_i + \eta_i' \mu_i^{(s-1)} + \eta_i'' (\mu_i^{(s-1)})^2], \\ b_i &\triangleq \vartheta_i^{(s-1)} + \gamma v_i^{(s-1)} (\eta_i' + 2\eta_i'' \mu_i^{(s-1)}),\end{aligned}$$

$c_i \triangleq c_i'' - c_i'$, $d_i \triangleq \gamma \eta_i'' (v_i^{(s-1)})^2$. Let $\underline{u}_i, \bar{u}_i$ be the two roots of the second order equation, if the equation has two roots,

$$d_i x^2 + b_i x + c_i = 0. \quad (18)$$

To ensure the variance in (17) exists, we need $v_i > 2$, $i = 1, \dots, k$. If we let $\alpha_i^{(0)} > 1$, $i = 1, \dots, k$, in the prior distribution, (17) is well defined for all $s = 1, \dots, T$. To efficiently solve the myopic optimization problem (15) and (17) numerically, we establish the following proposition.

Proposition 4 If $\alpha_i^{(0)} > 1$, $i = 1, \dots, k$, we have, for $i = 1, \dots, k$,

$$\begin{aligned}&E \left[\max \left(\mu_i^{(s)}, \max_{j \neq i} \mu_j^{(s-1)} \right) \mid \mathcal{E}_{s-1}, a_s = i \right] \\ &= \mu_i^{(s-1)} - \left(\mu_i^{(s-1)} - \max_{j \neq i} \mu_j^{(s-1)} \right) \Phi_{v_i}(\zeta_i) + \vartheta_i^{(s-1)} \sqrt{\frac{v_i}{v_i - 2}} \phi_{v_i - 2} \left(\sqrt{\frac{v_i - 2}{v_i}} \zeta_i \right),\end{aligned}$$

where $\phi_{v_i}(\cdot)$ is the probability density function (pdf) of Student's t -distribution with degrees of freedom v_i

$$\phi_{v_i}(t) \triangleq \frac{\Gamma(\frac{v_i+1}{2})}{\sqrt{\pi v_i} \Gamma(\frac{v_i}{2})} \left(1 + \frac{t^2}{v_i} \right)^{-\frac{1+v_i}{2}},$$

and $\Phi_{v_i}(\cdot)$ is the cumulative distribution function (cdf) of Student's t -distribution with degrees of freedom v_i . In addition, if equation (18) has two roots, we have

$$\begin{aligned}
 & E \left[\max \left(\mu_i^{(s)} + \gamma \vartheta_i^{(s)}, \max_{j \neq i} \left\{ \mu_i^{(s-1)} + \gamma \vartheta_i^{(s-1)} \right\} \right) \mid \mathcal{E}_{s-1}, a_s = i \right] \\
 &= c_i'' - c_i (\Phi_{v_i}(\bar{u}_i) - \Phi_{v_i}(\underline{u}_i)) + b_i \sqrt{\frac{v_i}{v_i - 2}} \left(\phi_{v_i-2} \left(\sqrt{\frac{v_i - 2}{v_i}} \bar{u}_i \right) - \phi_{v_i-2} \left(\sqrt{\frac{v_i - 2}{v_i}} \underline{u}_i \right) \right) \\
 &+ d_i \sqrt{\frac{v_i}{v_i - 2}} \left(\bar{u}_i \phi_{v_i-2} \left(\sqrt{\frac{v_i - 2}{v_i}} \bar{u}_i \right) - \underline{u}_i \phi_{v_i-2} \left(\sqrt{\frac{v_i - 2}{v_i}} \underline{u}_i \right) \right) \\
 &+ d_i \frac{v_i}{v_i - 2} \left[1 - \left(\Phi_{v_i-2} \left(\sqrt{\frac{v_i - 2}{v_i}} \bar{u}_i \right) - \Phi_{v_i-2} \left(\sqrt{\frac{v_i - 2}{v_i}} \underline{u}_i \right) \right) \right].
 \end{aligned} \tag{19}$$

otherwise, (19) = $c_i'' + d_i \frac{v_i}{v_i - 2}$.

Therefore, with Proposition 4, by calculating the pdf and cdf of t -distribution, gamma function, and some basic functions, we can efficiently solve the myopic optimization problem (15) and (17) numerically.

5 NUMERICAL EXAMPLE

In this section, numerical results are offered for comparison of various selection and allocation policies experimentally. We implement two allocation policies: the first-order KG policy denoted as KG₁ and the second-order KG policy denoted as KG₂. To combine the allocation policy with the selection policy, we denote KG₁(1) as the SADS policy comprising of allocation policy KG₁ and the selection policy \mathcal{D}^{pm} , KG₁(2) with \mathcal{D}^* . Same notation system is applied to the SADS policy with the second-order KG policy. In all experiments, the simulation replications of one design are i.i.d. normally distributed, replications between alternative design are independent, and the IPCSs, which are accurate up to three decimal points, are obtained by 10^5 independent experiments. The estimated IPCSs are graphed as a function of simulation budget.

There are 3 designs in this example, which illustrates the enhancement by implementing the optimal selection policy, and the improvement of the KG policy by adding posterior variance information. The prior distribution is the normal-Gamma conjugate prior with parameters given as: $\mu_i^{(0)} = 0$, $\kappa_i^{(0)} = 10^3$, $\alpha_i^{(0)} = 3$, $i = 1, 2, 3$, $\beta_1^{(0)} = 10^4$, and $\beta_2^{(0)} = \beta_3^{(0)} = 10^3$. From (8) and (9), we know the expectation of true mean $E[\mu_i] = 0$, $i = 1, 2, 3$, and the expectation of true variance $E[\sigma_1^2] = 5000$, and $E[\sigma_2^2] = E[\sigma_3^2] = 500$. In this example, the expected value of the simulation variance is significantly large, and the expected variance of the first design is 10 times of that other two.

From Figure 1, we can see the KG₂ outperforms KG₁ whether the selection policy is \mathcal{D}^* and \mathcal{D}^{pm} , and the IPCS of the KG₂(1) is even slightly higher than that of KG₁(2). Since the simulation budget is relatively small with the expected variance so high, the numerical results in this example can be understood as the small sample performance of various SADS policies.

6 CONCLUSION

This paper rigorously formulates the statistical selection problem as an optimal sampling allocation and design selection policy. We show that the optimal selection policy is not necessarily choosing the design with the largest sample mean, but can also depend on the correlation structure induced by the posterior variance. We derive a second-order KG policy to better utilize the information in posterior variance. The numerical experiments shows that our optimal selection policy and the myopic allocation policy can improve the IPCS especially in the small sample case. Ongoing work adds the dynamic structure for finding an optimal allocation policy.

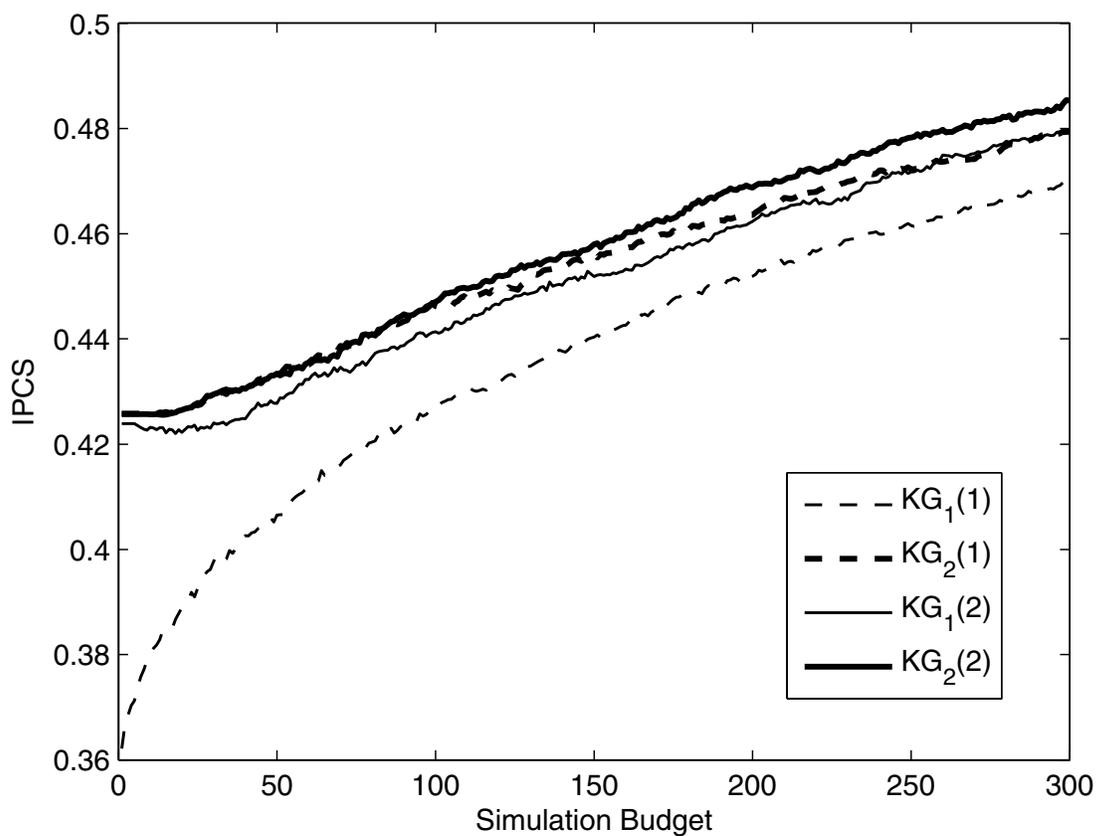


Figure 1: The prior distribution is the normal-Gamma conjugate prior, with parameters $\mu_i^{(0)} = 0$, $\kappa_i^{(0)} = 10^3$, $\alpha_i^{(0)} = 3$, $i = 1, 2, 3$, and $\beta_1^{(0)} = 10^4$, and $\beta_2^{(0)} = \beta_3^{(0)} = 10^3$. The IPCSs and their standard error are estimated by 10^5 independent experiments.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation (NSF) under Grants CMMI-0856256, CMMI-1233376, and EECs-0901543, by the Air Force Office of Scientific Research (AFOSR) under Grant FA9550-10-1-0340, by the Department of Energy under Grant DE-SC0002223, by the National Institutes of Health (NIH) under Grant 1R21DK088368-01, by the National Natural Science Foundation of China under Grants 71071040, 71028001, 70832002, 71061160506, by the National Science Council of Taiwan under Grant NSC-100-2218-E-002-027-MY3, by the Program for Professor of Special Appointment (Eastern Scholar) at Shanghai Institution of Higher Learning, and by the Programme for Cultivating Innovative Students in Key Disciplines of Fudan University.

REFERENCES

- Bechhofer, R. E., T. J. Santner, and D. M. Goldsman. 1995. *Design and Analysis for Statistical Selection, Screening, and Multiple Comparisons*. New York: John Wiley and Sons.
- Berger, J. O. 2004. *Statistical Decision Theory and Bayesian Analysis*, Volume 2. Springer.
- Branke, J., S. E. Chick, and C. Schmidt. 2007. "Selecting a selection procedure". *Management Science* 53 (12): 1916–1932.
- Chen, C.-H., D. He, and M. Fu. 2006. "Efficient dynamic simulation allocation in ordinal optimization". *IEEE Transactions on Automatic Control* 51 (12): 2005–2009.
- Chen, C.-H., J. Lin, E. Yücesan, and S. E. Chick. 2000. "Simulation budget allocation for further enhancing the efficiency of ordinal optimization". *Discrete Event Dynamic Systems* 10 (3): 251–270.
- Frazier, P. I., W. B. Powell, and S. Dayanik. 2008. "A knowledge-gradient policy for sequential information collection". *SIAM Journal on Control and Optimization* 47 (5): 2410–2439.
- Fu, M. C., J.-Q. Hu, C.-H. Chen, and X. Xiong. 2007. "Simulation allocation for determining the best design in the presence of correlated sampling". *INFORMS Journal on Computing* 19 (1): 101–111.
- Kelton, W. D., and A. M. Law. 2000. *Simulation Modeling and Analysis*. McGraw Hill Boston, MA.
- Peng, Y., C. Chen, M. Fu, and J. Hu. 2013. "Efficient simulation resource sharing and allocation for selecting the Best". *IEEE Transactions on Automatic Control* 58 (4): 1017–1023.
- Rose, C., and M. D. Smith. 2002. "mathStatICA: Mathematical Statistics with Mathematica". In *Compstat*, 437–442. Springer.
- Van der Vaart, A. W. 2000. *Asymptotic Statistics*, Volume 3. Cambridge university press.

AUTHOR BIOGRAPHIES

YI-JIE PENG is a Ph.D. student in the Management Science Department of School of Management at Fudan University. His research interests lie in simulation methodology, modeling, analysis, and optimization. His email address is 10110690016@fudan.edu.cn.

CHUN-HUNG CHEN is a Professor of Systems Engineering and Operations Research at George Mason University and is also affiliated with National Taiwan University. Dr. Chen has led research projects in stochastic simulation and optimization, sponsored by the NSF, NSC, FAA, Air Force, and NASA. He served as Co-Editor of the *Proceedings of the 2002 Winter Simulation Conference* and Program Co-Chair for 2007 Inform Simulation Society Workshop. He has served on the editorial boards of *IEEE Transactions on Automatic Control*, *IEEE Transactions on Automation Science and Engineering*, *IIE Transactions*, *Journal of Simulation Modeling Practice and Theory*, and *International Journal of Simulation and Process Modeling*. He received his Ph.D. degree from Harvard University in 1994. His email address is cchen9@gmu.edu.

MICHAEL C. FU is Ralph J. Tyser Professor of Management Science in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research and affiliate faculty appointment in the Department of Electrical and Computer Engineering, all at the University of Maryland. His re-

search interests include simulation optimization and applied probability, with applications in supply chain management and financial engineering. He has a Ph.D. in applied math from Harvard and degrees in math and EECS from MIT. He served as WSC2011 Program Chair, NSF Operations Research Program Director, *Management Science* Stochastic Models and Simulation Department Editor, and *Operations Research* Simulation Area Editor. He is a Fellow of **INFORMS** and **IEEE**. His email address is mfu@umd.edu.

JIAN-QIANG HU is a Professor with the Department of Management Science, School of Management, Fudan University. He received his B.S. degree in applied mathematics from Fudan University, China, and M.S. and Ph.D. degrees in applied mathematics from Harvard University. He was an Associate Professor with the Department of Mechanical Engineering and the Division of Systems Engineering at Boston University before joining Fudan University. His research interests include discrete-event stochastic systems, simulation, queueing network theory, stochastic control theory, with applications towards supply chain management, risk management in financial markets and derivatives, and communication networks. His e-mail addresses is hujq@fudan.edu.cn.