

Ranking and Selection with Covariates for Personalized Decision Making

 Haihui Shen,^a L. Jeff Hong,^{b,*} Xiaowei Zhang^c
^aSino-US Global Logistics Institute, Antai College of Economics and Management, Shanghai Jiao Tong University, Shanghai 200030, China;

^bSchool of Management and School of Data Science, Fudan University, Shanghai 200433, China; ^cHKU Business School, The University of Hong Kong, Pokfulam Road, Hong Kong SAR

*Corresponding author

Contact: shenhaihui@sjtu.edu.cn,  <https://orcid.org/0000-0002-4157-1278> (HS); hong_liu@fudan.edu.cn,

 <https://orcid.org/0000-0001-7011-4001> (LJH); xiaoweiz@hku.hk,  <https://orcid.org/0000-0002-5798-646X> (XZ)

Received: August 3, 2018

Revised: October 26, 2019; July 7, 2020; July 13, 2020; July 21, 2020

Accepted: July 22, 2020

Published Online in Articles in Advance: February 12, 2021

<https://doi.org/10.1287/ijoc.2020.1009>
Copyright: © 2021 INFORMS

Abstract. We consider a problem of ranking and selection via simulation in the context of personalized decision making, in which the best alternative is not universal, but varies as a function of some observable covariates. The goal of ranking and selection with covariates (R&S-C) is to use simulation samples to obtain a selection policy that specifies the best alternative with a certain statistical guarantee for subsequent individuals upon observing their covariates. A linear model is proposed to capture the relationship between the mean performance of an alternative and the covariates. Under the indifference-zone formulation, we develop two-stage procedures for both homoscedastic and heteroscedastic simulation errors, respectively, and prove their statistical validity in terms of average probability of correct selection. We also generalize the well-known slippage configuration and prove that the generalized slippage configuration is the least favorable configuration for our procedures. Extensive numerical experiments are conducted to investigate the performance of the proposed procedures, the experimental design issue, and the robustness to the linearity assumption. Finally, we demonstrate the usefulness of R&S-C via a case study of selecting the best treatment regimen in the prevention of esophageal cancer. We find that by leveraging disease-related personal information, R&S-C can substantially improve patients' expected quality-adjusted life years by providing a patient-specific treatment regimen.

History: Accepted by Bruno Tuffin, Area Editor for Simulation.

Funding: H. Shen was sponsored by the National Natural Science Foundation of China [Grants 72001140 and 71931007] and the "Chenguang Program" supported by Shanghai Education Development Foundation and Shanghai Municipal Education Commission [Grant 19CG14]. L. J. Hong was supported by the National Natural Science Foundation of China [Grant 71991473]. X. Zhang was supported by the Hong Kong Research Grants Council [General Research Fund Grants 16211417 and 17201520].

Supplemental Material: The e-companion is available at <https://doi.org/10.1287/ijoc.2020.1009>.

Keywords: ranking and selection • covariates • probability of correct selection • least favorable configuration • experimental design

1. Introduction

Ranking and selection (R&S) is one of the most studied problems in the area of stochastic simulation. It aims to select the one with the best mean performance from a set of alternatives through running simulation experiments; see Kim and Nelson (2006) and Chen et al. (2015) for reviews. In the conventional R&S setting, the mean performance of an alternative is considered as a constant. In context of personalized decision making, however, such a setting may be too rigid. For instance, medical studies show that the effectiveness of a cancer chemotherapy treatment depends on the biometric characteristics of a patient, such as tumor biomarker and gene expression (Yap et al. 2009, Kim et al. 2011). Therefore, for two patients with different characteristics, the best treatment

regimen may be different. Similar examples can also be found in marketing, in which research shows that the effect of an online advertisement depends on customer purchasing preference (Arora et al. 2008), and in training self-driving cars, in which the best driving decision depends on the real-time ambient information collected by all the sensors (Katrakazas et al. 2015). In all these examples, it appears more reasonable to consider the mean performance of an alternative as a function of the *covariates*, which include all of the additional contextual information, such as the biometric characteristics in the cancer treatment example, purchasing preference in the marketing example, and ambient information in the self-driving car example.

One approach to solving the problem is to run conventional R&S procedures once the covariates

are observed. However, this approach may be impractical in many situations for two reasons. First, the decision maker may not have the access or the time to run the simulation model. In the cancer treatment example, the simulation model often involves a complex Markov chain, and one needs to simulate its steady state in order to estimate the treatment effects. It is well known that steady-state simulation is computationally expensive as it may take a long time for the Markov chain to reach the steady state. (The brute force simulation that is based on a simplified Markov chain model and a grid-based interpolation to compute the policy for assigning personalized treatment regimens took about eight days on our desktop computer to finish; see Section 8 for details.) In addition, the doctor may not have access to the simulation model that needs to be run on sophisticated computer systems. In the marketing example, the online retailer has to display the advertisement once the customer is logged in and, thus, has no time to run simulation experiments. In the self-driving car example, the time is more precious, and decisions have to be made in real time. A second reason is about efficiency. Personalized decisions typically need to be made repeatedly for different people upon observing their covariates. Then, running conventional R&S procedures for each person is conceivably less efficient than developing a selection policy, which maps the covariates to the identity of the best alternative, and using it repeatedly for different people.

We note that we are interested in the kind of situations in which the simulation model is expensive to run. Unless the covariates are of a low dimension, it would be computationally prohibitive to discretize the domain space of the covariates into a grid and then simulate the alternatives at each grid point in a brute force manner. Instead, an intelligent design is needed to allocate a computational budget both to the alternatives and over the domain space of the covariates. This is also a critical issue that differentiates this new problem from the conventional R&S problem, which is only concerned with allocating a computational budget to the alternatives.

In this paper, we consider a new R&S setting in which the mean performances of all alternatives are functions of the covariates, and therefore, the identity of the best alternative is also a function of the covariates. One may run simulation experiments to learn the mean performance functions of all alternatives and to use these learned functions to select the best alternative upon observing the covariates. We call this problem *ranking and selection with covariates* (R&S-C). Notice that, under this setting, the time-consuming component is the learning of the mean performance functions of all alternatives, and it requires a significant amount of simulation effort. However, this can

be done off-line. Once the mean performance functions are learned, only these learned functions (which form the selection policy) need to be deployed. Then, the selection of the best upon observing the covariates is basically computing the function values of all alternatives at the values of the covariates, and it can be done online in real time with negligible computation. Notice that such an off-line learning online application approach allows the learned functions to be deployed to many users (e.g., doctors and self-driving cars) and used repeatedly with no additional cost.

1.1. Main Contributions

To tackle the R&S-C problem, we first provide a general frequentist formulation. We generalize important frequentist R&S concepts, such as the indifference zone and the probability of correct selection (PCS), to the R&S-C setting and define the corresponding finite-sample statistical guarantee. We also show that the R&S-C formulation, in general, gives a better outcome than the R&S formulation if one chooses to average the effects of the covariates.

Second, we consider a specific situation of the R&S-C problem, in which the mean performances of all alternatives are linear in the covariates (or linear in certain basis functions of the covariates) with unknown coefficients that may be estimated through linear regression, and show that Stein's lemma (Stein 1945), which is a major cornerstone of the conventional frequentist R&S, may be extended to linear regression contexts. Despite its simplicity, linear models have distinct advantages in terms of their interpretability and robustness to model misspecification and often show good performance in prediction (James et al. 2013).

Third, we propose two-stage procedures to solve R&S-C problems with linear performance functions. These procedures may be viewed as the extensions of the famous Rinott's procedure (Rinott 1978) in the conventional frequentist R&S, and they can handle both homoscedastic and heteroscedastic errors, respectively, in the linear models. Based on the extended Stein's lemma that we develop, we prove that these procedures deliver the desired finite-sample statistical guarantee. We also conduct numerical studies to assess the performances of the procedures and discuss their robustness to the linearity assumption and to the experimental design.

Finally, we consider the personalized prevention regimen for esophageal cancer, in which the effectiveness of the prevention regimens are evaluated using a Markov simulation model developed and calibrated by domain experts in cancer research. We compare the R&S-C formulation with the conventional R&S formulation and show that the R&S-C formulation can significantly improve the expected quality-adjusted life years (QALYs) of patients who

are diagnosed with Barrett's esophagus (BE), a mild precursor to esophageal cancer.

1.2. Related Literature

R&S has been studied extensively in the statistics and stochastic simulation literature. In general, there are two streams of procedures: frequentist and Bayesian. Frequentist procedures typically aim to deliver the PCS under the indifference-zone formulation. There are two-stage procedures (Rinott 1978), sequential procedures (Kim and Nelson 2001, Hong 2006), and procedures designed to handle a large number of alternatives in parallel computing environments (Luo et al. 2015, Ni et al. 2017). These procedures are typically conservative and require more samples than necessary for average cases. Bayesian procedures, on the other hand, often aim to allocate a finite computational budget to different alternatives to either maximize the posterior PCS or minimize the expected opportunity cost. There are a variety of approaches to developing Bayesian procedures, including value of information (Chick and Inoue 2001), knowledge gradient (Frazier et al. 2008), optimal computing budget allocation (Chen et al. 1997), and economics of selection procedures (Chick and Gans 2009, Chick and Frazier 2012). Bayesian procedures often require fewer samples than frequentist ones to achieve the same level of PCS. However, they do not provide a (frequentist) statistical guarantee in general. Frazier (2014) develops a Bayes-inspired procedure that includes many of the Bayesian features while still guaranteeing a frequentist PCS.

The R&S-C problems have been tackled in the Bayesian framework. Hu and Ludkovski (2017) propose to model the performance functions of alternatives as Gaussian random fields and use the expected improvement criteria to develop Bayesian procedures. Pearce and Branke (2017) follow the same framework of Hu and Ludkovski (2017) but focus on how to efficiently estimate the expected improvement over a continuous domain. These Bayesian procedures for R&S-C aim to adaptively allocate a given sampling budget to the alternatives and over the domain of covariates in an efficient way. However, a statistical guarantee on the performance of their solutions is yet to be proved. In contrast to their approaches, we take a frequentist perspective in this paper, for the first time to our knowledge, to model and solve the R&S-C problems, and it aims to achieve a certain finite-sample statistical guarantee on the performance of the solution.

Our research is also related to the literature on the multiarm bandit (MAB) with covariates. MAB is an important class of sequential decision-making problems in the fields of operations research, statistics, and machine learning. It was first proposed by

Robbins (1952) and has been studied extensively since then; see, for instance, Bubeck and Cesa-Bianchi (2012) for a comprehensive review of MAB. In recent years, MAB with covariates (also known as contextual MAB) has drawn considerable attention as a tool for personalized decision making. The mean performances in these problems are often modeled as linear functions of the covariates (Auer 2002, Rusmevichientong and Tsitsiklis 2010). In particular, Goldenshluger and Zeevi (2013) consider a linear model whose coefficients are arm-dependent, which motivates our formulation of R&S-C. Nonparametric models have also been considered in the literature of MAB with covariates (Perchet and Rigollet 2013, Slivkins 2014), and this may be a direction of future study for R&S-C. However, a critical distinction between MAB with covariates and R&S-C lies in the way that the covariates are obtained. The values of the covariates arrive randomly in the former, whereas in the latter, they can be chosen by the decision maker. The additional freedom may conceivably allow one to learn the relationship between the mean performance and the covariates more efficiently.

Linear models have also been considered in R&S problems. For example, Negoescu et al. (2011) adopt a linear model when solving an R&S problem in the context of drug discovery, in which the mean performance of an alternative is a linear combination of attribute contributions. However, the intention of introducing the linear model in Negoescu et al. (2011) is quite different from ours. Specifically, the linear model in their work forms a linear projection from the space of alternatives to the space of attributes, which dramatically reduce the computational complexity. Their final goal is still to select the best alternative as a static decision rather than the kind of decision policy that we seek. Therefore, their R&S problem is still in the conventional sense, which is different from the R&S-C problems considered in this paper.

A preliminary version of this paper (Shen et al. 2017) was published in the Proceedings of the 2017 Winter Simulation Conference. This paper extends Shen et al. (2017) significantly by providing all the proofs for statistical validity and the least favorable configuration, discussing the experimental design and robustness to linearity assumptions, adding the discussions on nonnormal simulation errors, and applying the proposed formulation and procedure to a case study on personalized medicine.

The remainder of the paper is organized as follows. In Sections 2 and 3, we formulate the R&S-C problem and introduce the linear models. In Section 4, we develop procedures for homoscedastic, heteroscedastic, and nonnormal simulation errors. The least favorable configuration is discussed in Section 5, followed by numerical experiments in Section 6 and a robustness

study in Section 7. We demonstrate the practical value of R&S-C in the context of personalized medicine in Section 8 and conclude in Section 9. Technical proofs are included in the e-companion of this paper.

2. Problem Formulation

Suppose there are k alternatives whose mean performances, denoted as $\mu_1(\mathbf{X}), \dots, \mu_k(\mathbf{X})$, are functions of $\mathbf{X} = (X_1, \dots, X_d)^\top$, which is the vector of the observable random covariates with support $\Theta \subset \mathbb{R}^d$. Our goal is to develop a policy that selects the alternative with the largest mean performance upon observing the values of the covariates, that is, identifying $i^*(x) := \arg \max_{1 \leq i \leq k} \{\mu_i(\mathbf{X}) | \mathbf{X} = x\}$ for any $x = (x_1, \dots, x_d) \in \Theta$. In the cancer treatment example considered in Section 1, for instance, the alternatives are the different treatments, the covariates are the biometric characteristics of a patient, the mean performances are the expected QALYs of the patient under different treatments, and the goal is to identify a policy that selects the best treatment for the patient once the biometric characteristics of the patient are observed.

In this paper, we suppose that there are simulation models that allow us to estimate $\mu_1(\mathbf{X}), \dots, \mu_k(\mathbf{X})$ once the values of \mathbf{X} are given. The critical issue here is how to design off-line simulation experiments to learn $\mu_1(x), \dots, \mu_k(x)$ accurately so that they may be used to select the best alternative in real time with a pre-determined level of precision upon observing the values of \mathbf{X} (e.g., PCS in a frequentist sense). We call this problem R&S-C to emphasize that the decision is conditional on the covariates.

Remark 1. Throughout this paper, we assume that the value of \mathbf{X} is observable before making the selection decision. This assumption is reasonable in many practical situations, including the three examples introduced in Section 1. Specifically, in the cancer treatment example, patients’ characteristics, such as tumor biomarkers and gene expressions, can be identified through medical tests; in the marketing example, customer preference can be inferred from the demographic and behavioral information as well as the purchasing history (if available); and in the self-driving car example, the ambient information is collected directly by the sensors.

2.1. Value of Covariates

In the conventional R&S problem, the goal may be viewed as selecting the unconditional best, that is, to identify $i^\dagger := \arg \max_{1 \leq i \leq k} \mu_i$, where $\mu_i := \mathbb{E}[\mu_i(\mathbf{X})]$, $i = 1, \dots, k$, and the expectation is taken with respect to the distribution of \mathbf{X} . In the cancer treatment example, for instance, the conventional R&S selects the best treatment for the entire population instead of

the best for an individual patient. Notice that $\mu_i(\mathbf{X})$ is a random variable. Then, by Jensen’s inequality,

$$\begin{aligned} \mathbb{E}[\mu_{i^*(\mathbf{X})}(\mathbf{X})] &= \mathbb{E}\left[\max_{1 \leq i \leq k} \mu_i(\mathbf{X})\right] \geq \max_{1 \leq i \leq k} \mathbb{E}[\mu_i(\mathbf{X})] \\ &= \mathbb{E}[\mu_{i^\dagger}(\mathbf{X})]. \end{aligned} \tag{1}$$


Therefore, the R&S-C formulation typically outperforms the conventional R&S formulation if the covariates are observable before the selection decision is made. In the cancer treatment example, for instance, Equation (1) implies that the personalized-best treatment typically outperforms the population-best treatment. This point is also demonstrated in the cancer prevention example considered in Section 8.

Remark 2. The distribution of \mathbf{X} is assumed to be known in this paper. This is a common assumption in conventional R&S, in which the distribution of \mathbf{X} needs to be known to evaluate $\mathbb{E}[\mu_i(\mathbf{X})]$ for all $i = 1, \dots, k$. Here, the distribution can be discrete, continuous, or even a mixture of them. Moreover, the elements of \mathbf{X} can be dependent on each other. In practice, the distribution of \mathbf{X} is typically estimated through the input modeling process (Law and Kelton 2000).

2.2. Indifference Zone

The concept of an indifference zone (IZ) plays a key role in conventional frequentist R&S (Bechhofer 1954). It defines the smallest difference δ that the decision maker considers worth detecting. Therefore, alternatives whose mean performances are within δ to the best are in the IZ and are considered “indifferent” from the best. In a frequentist setting, R&S procedures need to deliver the PCS under any configuration of the means. Without the IZ, the best and the other alternatives may be arbitrarily close in their means. One would then need infinitely many samples to identify the true best even with PCS less than one. In the presence of the IZ, the goal is to select one of the

Figure 1. (Color online) Example of Mean Configuration with $d = 1$ and $k = 3$



alternatives in the IZ, and therefore, a finite number of samples are needed to ensure that the selected alternative is in the IZ with PCS less than one.

In the setting of R&S-C, the configurations of the means depend on the values of the covariates. They may be arbitrarily close if the mean surfaces $\mu_1(x), \dots, \mu_k(x)$ intersect at some values of $x \in \Theta$ (see, for instance, Figure 1). Therefore, we also need IZ. Given an IZ parameter, we define the event of correct selection (CS) given $X = x$ as

$$CS(x) := \left\{ \mu_{i^*(x)}(x) - \mu_{\widehat{i^*(x)}}(x) < \delta \right\},$$

where $\widehat{i^*(x)}$ denotes the selected best, and a CS event implies that the selected best is in the IZ. Notice that our definition of a CS event is also known as a good selection event in some of the R&S literature (see, for instance, Ni et al. 2017), where a CS event may be defined more restrictively as $\{i^*(x) = \widehat{i^*(x)}\}$ given that there are no alternatives in the IZ other than the best. However, this more restrictive definition of CS often does not make sense in the context of R&S-C. For instance, when the support of the covariates cover the intersection points of the mean surfaces (see point x^* in Figure 1), no matter how small the IZ parameter δ is, for certain values of the covariates that are in the neighborhood of the intersection points, there are always alternatives in the IZ other than the best, which makes the more restrictive definition of CS inapplicable.

Remark 3. Recently, Fan et al. (2016) show that IZ may be unnecessary for the conventional frequentist R&S if sequential procedures are used. However, in order for their procedures to stop in finite time (with finite samples), the means of all alternatives have to be different. In the R&S-C context, however, the mean values of some alternatives are the same at the intersection points (see Figure 1). Therefore, it is not clear how the procedures of Fan et al. (2016) may be applied in the R&S-C context.

2.3. Probability of Correct Selection

We are now ready to define the PCS, which is the statistical guarantee that frequentist R&S procedures typically deliver. Let $\widehat{i^*(x)}$ denote the selection policy produced by an R&S-C procedure. Notice that the presence of the covariates complicates the definition of PCS because one has to answer whether the PCS is defined for an individual or the population. To address the issue, we first define the *conditional* PCS, given $X = x$, as

$$PCS(x) := \mathbb{P}\left\{ \mu_{i^*(x)}(X) - \mu_{\widehat{i^*(x)}}(X) < \delta \mid X = x \right\}, \quad (2)$$

where the probability is taken with respect to the distribution of simulated samples that are used to estimate the

mean functions $\mu_1(x), \dots, \mu_k(x)$ and to derive the selection policy $i^*(x)$ for all $x \in \Theta$.

Notice that $PCS(x)$ may be viewed as the PCS for an individual whose covariates take the value x . However, the covariates are random variables, and therefore, $PCS(X)$ is also a random variable. To use it as the statistical guarantee for R&S-C procedures, one way is to consider some summary statistics of $PCS(X)$. To that end, we define the average PCS, denoted by PCS_E , as

$$PCS_E := \mathbb{E}[PCS(X)]. \quad (3)$$

Notice that

$$\begin{aligned} PCS_E &= \mathbb{E}\left[\mathbb{P}\left\{ \mu_{i^*(X)}(X) - \mu_{\widehat{i^*(X)}}(X) < \delta \mid X \right\} \right] \\ &= \mathbb{P}\left\{ \mu_{i^*(X)}(X) - \mu_{\widehat{i^*(X)}}(X) < \delta \right\}. \end{aligned}$$

Therefore, PCS_E is the *unconditional* PCS, and it is for the entire population. If we set $PCS_E \geq 1 - \alpha$, we are $(1 - \alpha)$ confident that a random individual from the population will select the individual's personalized best decision or a decision that is within the IZ. We want to point out that other summary statistics of $PCS(X)$ may also be used as decision criteria, for instance, one may define it to be a certain quantile of the random variable $PCS(X)$ or even $\min_{x \in \Theta} PCS(x)$ to be more risk averse. The $\min_{x \in \Theta} PCS(x)$ statistic focuses on the worst $PCS(x)$ over Θ , a much more conservative criterion compared with PCS_E , whereas quantile of $PCS(X)$ is a more flexible statistic, whose conservativeness can be adjusted by the quantile parameter.

Now, we summarize our problem. We want to design a selection procedure that samples each alternative off-line to estimate the mean functions $\mu_1(x), \dots, \mu_k(x)$ and then produce the selection policy $i^*(x)$ for all $x \in \Theta$. This policy is used to select the best alternative in real time upon observing the values of X , and it should reach the prespecified target of PCS_E , say, $1 - \alpha$.

Remark 4. We focus on PCS_E in the main text of this paper. Selection procedures can be developed and analyzed analogously if the prespecified target is $PCS_{\min} := \min_{x \in \Theta} PCS(x) \geq 1 - \alpha$. Detailed discussion including numerical experiments is provided in Section EC.6 of the e-companion.

3. Linear Models and the Extended Stein's Lemma

Notice that the general formulation of R&S-C problems, presented in Section 2, allows the mean performance functions $\mu_1(x), \dots, \mu_k(x)$ to take any forms. To solve the problems, however, one needs to decide how to estimate these functions. There are two approaches:

parametric and nonparametric. Both have pros and cons, and both are widely used in function estimation. In this paper, we take a parametric approach and assume that $\mu_1(x), \dots, \mu_k(x)$ are linear functions of the covariates x with unknown coefficients that need to be estimated through simulation experiments. Let $Y_i(x)$ denote the random performance of alternative i at the covariates x for all $i = 1, \dots, k$ and $x \in \Theta$. We make the following assumption on the forms of $\mu_1(x), \dots, \mu_k(x)$ and distributions of $Y_1(x), \dots, Y_k(x)$.

Assumption 1. For all $i = 1, \dots, k$,

$$\begin{aligned}\mu_i(x) &= x^\top \beta_i, \\ Y_i(x) &= \mu_i(x) + \epsilon_i,\end{aligned}$$

where $\beta_i = (\beta_{i1}, \dots, \beta_{id})^\top \in \mathbb{R}^d$ is a vector of unknown coefficients and the simulation error ϵ_i follows a normal distribution with mean zero and variance $\sigma_i^2 < \infty$. In addition, the simulation errors are independent among different alternatives, different covariates, and different replications.

The linear model in Assumption 1 is crucial for analyzing the performance in terms of PCS_E of the proposed selection procedures. In particular, without the linear structure, it would be challenging to obtain a finite-sample performance guarantee even with the normality assumption on the simulation errors. Under a nonparametric model, such as kernel regression and tree-based methods, selection procedures as well as the analysis of their statistical performance would be drastically different from what is done in this paper. We leave the investigation to future research.

Despite the simplicity, linear models usually have the advantages of high interpretability and robustness to model misspecification (James et al. 2013). We study in Section 7 the robustness of the procedures developed in this paper when the linearity assumption does not hold.

Remark 5. The linear model in Assumption 1 can be generalized to capture the nonlinearity of $\mu_i(x)$ in x by the use of basis functions. That is, we may postulate $\mu_i(x) = f(x)^\top \beta_i$, where $f(x)$ is a vector of basis functions (e.g., polynomials or radial basis functions) that one selects carefully. (Nevertheless, selecting a good set of basis functions is a nontrivial task, and it is beyond the scope of this paper.) Note that, if we view $f(x)$ as a new set of covariates through a change of variables, Assumption 1 and the analysis in the sequel still hold. Moreover, we often set $X_1 \equiv 1$ to allow an intercept term in the linear model. We may also include categorical variables in x by the use of dummy variables.

Notice that Assumption 1 basically requires all $Y_i(x)$ to follow the standard linear regression assumption (James et al. 2013) so that the unknown coefficient vectors β_i may be estimated using a standard ordinary least squares (OLS) method. Furthermore, Assumption 1 is a natural extension of the normality assumption commonly used in the R&S literature. For instance, both Rinott (1978) and Kim and Nelson (2001) assume that $Y_i = \mu_i + \epsilon_i$ for all $i = 1, \dots, k$. We basically extend the mean μ_i to a linear function $\mu_i(x) = x^\top \beta_i$ to add the effect of the covariates. Moreover, we see later in this section that the OLS estimators of the unknown parameters β_i under Assumption 1 resemble the sample-mean estimators of the unknown means under the normality assumption. This resemblance gives us great convenience to develop statistically valid R&S-C procedures.

3.1. Fixed Design

Based on Assumption 1, a critical issue in solving an R&S-C problem is to obtain estimates of β_1, \dots, β_k that are accurate enough. Therefore, we need to decide when to run simulation experiments (i.e., the *design points*) and how many observations to run (i.e., the *sample sizes*). Here, we want to emphasize again that estimation of β_i is conducted off-line based on simulation experiments at the chosen design points instead of real experiments at randomly observed values of the covariates. Hence, we are free in choosing the number of design points, their locations, and the number of samples to be taken at each design point. This naturally becomes an experimental design problem, which could be formulated as an optimization problem with the objective of minimizing certain metrics of the error in estimating the policy $\hat{i}^*(\cdot) = \arg \max_{1 \leq i \leq k} \mu_i(\cdot)$. However, this problem is much more challenging to solve than the experimental design problem for linear regression (Silvey 1980), primarily because the $\arg \max$ operation is nonlinear in the unknown surfaces. It is beyond the scope of this paper to find both the optimal design and the optimal sample size that jointly provide a guarantee on the PCS of the estimated policy $\hat{i}^*(\cdot)$.

In this paper, we choose to use a fixed set of design points to estimate β_i for all $i = 1, \dots, k$. In particular, we select a set of m design points, denoted as $x_1, \dots, x_m \in \Theta$, with $m \geq d$, and conduct simulation experiments only at these design points for all alternatives. Notice that the use of fixed design points eliminates the randomness in choosing design points. It simplifies the analysis and makes statistically valid R&S-C procedures significantly easier to develop. When adopting a fixed design, the placement of the design points is an important issue. We discuss it in

Section 7 under the premise that the number of design points is given. As for now, we simply consider the situation in which m design points are chosen, and they satisfy that $\mathcal{X}^\top \mathcal{X}$ is a nonsingular matrix, in which $\mathcal{X} = (x_1, \dots, x_m)^\top \in \mathbb{R}^{m \times d}$. Notice that the nonsingularity of $\mathcal{X}^\top \mathcal{X}$ is a standard condition in linear regression (James et al. 2013). It ensures that all β_i may be estimated properly.

3.2. Extended Stein’s Lemma

Conventional R&S procedures often have a first stage to estimate the means and variances of all alternatives and use them to determine the remaining sample sizes (or sampling policy). For instance, the two-stage procedures of Dudewicz and Dalal (1975) and Rinott (1978) and the sequential procedures of Kim and Nelson (2001) and Hong (2006) use the sample variances, and the optimal computing budget allocation procedure of Chen et al. (1997) uses both the sample means and variances. However, this may create a statistical issue because the overall sample size of an alternative now depends on its first-stage samples. Then, what is the distribution of the overall sample mean?

Stein’s lemma (Stein 1945) critically answers this question. The lemma shows that, if Y_1, Y_2, \dots are independent and identically distributed (i.i.d.) normal random variables and N depends on the first-stage samples only through the sample variance, then the overall sample mean $(Y_1 + \dots + Y_N)/N$, conditionally on the first-stage sample variance, still has a normal distribution. Consequently, this lemma became a cornerstone of the conventional frequentist R&S procedures in proving finite-sample statistical guarantees with unknown variances; see Dudewicz and Dalal (1975) and Rinott (1978) for early use of this lemma in designing two-stage R&S procedures, and theorem 2 of Kim and Nelson (2006) for a rephrased version of the lemma.

In R&S-C, we also face the problem of unknown variances, that is, $\sigma_1^2, \dots, \sigma_k^2$ are unknown in the linear models. Moreover, we have to deal with the OLS estimators $\hat{\beta}_i$ instead of only the sample means as in the conventional R&S. Suppose that we have $m \geq d$ design points with the design matrix \mathcal{X} and each sample includes an observation from every design point. Then, we have the following extended Stein’s lemma. We provide its proof in Section EC.1 of the e-companion, in which a more general version is stated and proved, but we remark here that the assumption of the linear models is crucial.

Lemma 1 (Extended Stein’s Lemma). *Let $Y = \mathcal{X}\beta + \epsilon$, where $\beta \in \mathbb{R}^d$, $\mathcal{X} \in \mathbb{R}^{m \times d}$, and $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathcal{I})$ with $\mathbf{0}$ denoting the zero vector in \mathbb{R}^m and \mathcal{I} the identity matrix in $\mathbb{R}^{m \times m}$. Assume that $\mathcal{X}^\top \mathcal{X}$ is nonsingular. Let T be a random variable independent of $\sum_{\ell=1}^n Y_\ell$ and of $\{Y_\ell : \ell \geq n + 1\}$,*

where Y_1, Y_2, \dots are independent samples of Y . Suppose that $N \geq n$ is an integer-valued function of T and no other random variables. Let $\hat{\beta} = N^{-1}(\mathcal{X}^\top \mathcal{X})^{-1} \mathcal{X}^\top \sum_{\ell=1}^N Y_\ell$. Then, for any $x \in \mathbb{R}^d$,

- i. $x^\top \hat{\beta} | T \sim \mathcal{N}(x^\top \beta, \frac{\sigma^2}{N} x^\top (\mathcal{X}^\top \mathcal{X})^{-1} x)$.
- ii. $\frac{\sqrt{N}(x^\top \hat{\beta} - x^\top \beta)}{\sigma \sqrt{x^\top (\mathcal{X}^\top \mathcal{X})^{-1} x}}$ is independent of T and has the standard normal distribution.

Remark 6. If we set $m = d = 1$ and $\mathcal{X} = 1$, Y becomes a scalar, and it follows $\mathcal{N}(\beta_1, \sigma^2)$. Then, Lemma 1 becomes Stein’s lemma (Stein 1945). In this sense, Lemma 1 is an extension of Stein’s lemma to the linear regression context.

Remark 7. In Lemma 1, if we let T denote the OLS estimator of the variance σ^2 computed using samples Y_1, \dots, Y_n , then, by Rencher and Schaalje (2008, theorem 7.6b), $(nm - d)T/\sigma^2$ follows a chi-square distribution with $(nm - d)$ degrees of freedom, and it is independent of $\sum_{\ell=1}^n Y_\ell$ and of $\{Y_\ell : \ell \geq n + 1\}$. Therefore, similar to the conventional frequentist R&S, we may let the sample sizes of all alternatives depend on their first-stage OLS variance estimators and still keep the desired statistical properties.

4. Two-Stage Procedures

For conventional frequentist R&S procedures, there are two-stage and sequential procedures. Even though both types of procedure are designed based on the least favorable configuration of means, sequential procedures, such as those of Kim and Nelson (2001) and Hong (2006), take advantage of the information on means and allow the procedures to terminate earlier if the differences between the best and rest of the alternatives are significantly larger than the IZ parameter. On the other hand, two-stage procedures, such as those of Dudewicz and Dalal (1975) and Rinott (1978), do not take the mean information into consideration and are, thus, often more conservative.

In R&S-C problems, however, the configurations of the means depend on the realizations of the covariates. For some realizations of the covariates, the differences may be larger than the IZ parameter, and for other realizations, the differences may be much smaller, even close to zero (see, for instance, Figure 1). The procedures that we intend to design need to deliver a selection policy $\hat{r}^*(x)$ for all $x \in \Theta$ before the covariates are realized, and the policy may be used repeatedly for many realizations of the covariates. Therefore, it is not clear whether sequential procedures may still be advantageous in the R&S-C context. In this paper, we focus on designing two-stage procedures that deliver the desired finite-sample statistical guarantee.

4.1. The Procedure and Statistical Validity

We develop a two-stage procedure, called procedure TS, for R&S-C problems under Assumption 1. In the first stage, the procedure takes a small number of samples from all design points to estimate the total sample size required to deliver the desired statistical guarantee, and in the second stage, it takes the additional samples and produces a selection policy based on all samples. The structure of the procedure resembles many of the conventional two-stage R&S procedures, including those of Dudewicz and Dalal (1975) and Rinott (1978).

Procedure TS

Setup: Specify the target $\text{PCSE} 1 - \alpha$, the IZ parameter $\delta > 0$, the first-stage sample size $n_0 \geq 2$, the number of design points $m \geq d$, and the design matrix \mathcal{X} with a nonsingular $\mathcal{X}^\top \mathcal{X}$. Let h satisfy the following equation:

$$\mathbb{E} \left\{ \int_0^\infty \left[\int_0^\infty \Phi \left(\frac{h}{\sqrt{(n_0 m - d)(t^{-1} + s^{-1}) \mathbf{X}^\top (\mathcal{X}^\top \mathcal{X})^{-1} \mathbf{X}}} \right) \times \eta(s) ds \right]^{k-1} \eta(t) dt \right\} = 1 - \alpha, \quad (4)$$

where $\Phi(\cdot)$ is the cumulative distribution function (cdf) of the standard normal distribution, $\eta(\cdot)$ is the probability density function (pdf) of the chi-square distribution with $(n_0 m - d)$ degrees of freedom, and the expectation is taken with respect to the distribution of \mathbf{X} .

First-stage sampling: Take n_0 independent samples from each alternative i at each design point \mathbf{x}_j through simulation and denote them by $\mathbf{Y}_{i\ell} = (Y_{i\ell}(\mathbf{x}_1), \dots, Y_{i\ell}(\mathbf{x}_m))^\top$, $i = 1, \dots, k$, and $\ell = 1, \dots, n_0$. For each $i = 1, \dots, k$, let

$$\widehat{\boldsymbol{\beta}}_{i0} = \frac{1}{n_0} (\mathcal{X}^\top \mathcal{X})^{-1} \mathcal{X}^\top \sum_{\ell=1}^{n_0} \mathbf{Y}_{i\ell},$$

$$S_i^2 = \frac{1}{n_0 m - d} \sum_{\ell=1}^{n_0} (\mathbf{Y}_{i\ell} - \mathcal{X} \widehat{\boldsymbol{\beta}}_{i0})^\top (\mathbf{Y}_{i\ell} - \mathcal{X} \widehat{\boldsymbol{\beta}}_{i0}).$$

Second-stage sampling: Compute the total sample size $N_i = \max\{[h^2 S_i^2 / \delta^2], n_0\}$ for each i , where $[a]$ denotes the smallest integer no less than a . Take $N_i - n_0$ additional independent samples from alternative i at all design points through simulation, $\mathbf{Y}_{i, n_0+1}, \dots, \mathbf{Y}_{i, N_i}$, $i = 1, \dots, k$. For each alternative i , let

$$\widehat{\boldsymbol{\beta}}_i = \frac{1}{N_i} (\mathcal{X}^\top \mathcal{X})^{-1} \mathcal{X}^\top \sum_{\ell=1}^{N_i} \mathbf{Y}_{i\ell}.$$

Selection: Return $\widehat{i}^*(\mathbf{x}) = \arg \max_{1 \leq i \leq k} \{\mathbf{x}^\top \widehat{\boldsymbol{\beta}}_i\}$ as the selection policy.

Remark 8. Similar to typical two-stage R&S procedures, the first-stage sample size n_0 is chosen heuristically here. If n_0 is too small, then h calculated from (4) tends to be large, leading to excessive second-stage samples to compensate for the inaccurate variance estimator S_i^2 in the first stage. Taking $n_0 \geq 10$ is a common recommendation (Kim and Nelson 2006).

Remark 9. The constant h , defined in (4), is computed numerically. In our numerical experiments, the integrations and expectations are computed by the MATLAB built-in numerical integration function `integral`, and h is solved by the MATLAB built-in root-finding function `fzero`. However, the numerical integration may suffer from the curse of dimensionality if the dimension of \mathbf{X} is large. In such situations, one may use the Monte Carlo method to approximate the expectation or apply the stochastic approximation method (Robbins and Monro 1951) to find the root h . See more discussion in Section EC.2 of the e-companion.

The following theorem states that procedure TS is statistically valid under Assumption 1. We include the proof in Section EC.3 of the e-companion, but remark here that the proof relies critically on the extended Stein's lemma (Lemma 1).

Theorem 1. *Suppose that procedure TS is used to solve the R&S-C problem and Assumption 1 is satisfied. Then, $\text{PCSE} \geq 1 - \alpha$.*

4.2. Handling Heteroscedastic Errors

In Assumption 1, we assume that the variance of simulated samples of an alternative does not change with respect to the values of the covariates. This implies that the linear models all have homoscedastic simulation errors. However, this assumption may not always hold. In many practical situations, such as queueing and financial applications, simulation errors are often heteroscedastic. In this section, we present a two-stage R&S-C procedure to take care of the heteroscedasticity in the linear models.

We first extend Assumption 1 to the following to allow heteroscedastic errors.

Assumption 2. *For all $i = 1, \dots, k$,*

$$\mu_i(\mathbf{x}) = \mathbf{x}^\top \boldsymbol{\beta}_i,$$

$$Y_i(\mathbf{x}) = \mu_i(\mathbf{x}) + \epsilon_i(\mathbf{x}),$$

where $\boldsymbol{\beta}_i = (\beta_{i1}, \dots, \beta_{id})^\top \in \mathbb{R}^d$ is a vector of unknown coefficients and the simulation error $\epsilon_i(\mathbf{x})$ follows a normal distribution with mean zero and variance $\sigma_i^2(\mathbf{x}) < \infty$. In addition, the simulation errors are independent among different alternatives, different covariates, and different replications.

When linear models have heteroscedastic errors, OLS estimators of $\boldsymbol{\beta}_i$ are still consistent estimators.

However, subtle controls are needed to deliver the required PCS_E . Because our experiments are controlled simulation experiments, we may run multiple simulation runs at each design point to calculate the sample variance at the design point. We then use these sample variances to determine the (different) total sample sizes at different design points. We call this new two-stage procedure TS^+ . One distinct feature of procedure TS^+ is that it allows different design points to have different total sample sizes to handle heteroscedastic errors, and procedure TS always assigns the same total sample size to all design points. The following is the procedure. For simplicity, we use χ_v^2 to denote the chi-square distribution with v degrees of freedom.

Procedure TS^+

Setup: Specify the target PCS_E $1 - \alpha$, the IZ parameter $\delta > 0$, the first-stage sample size $n_0 \geq 2$, the number of design points $m \geq d$, and the design matrix \mathcal{X} with a nonsingular $\mathcal{X}^T \mathcal{X}$. Let h_{Het} satisfy the following equation:

$$\mathbb{E} \left\{ \int_0^\infty \left[\int_0^\infty \Phi \left(\frac{h_{Het}}{\sqrt{(n_0 - 1)(t^{-1} + s^{-1})} \mathbf{X}^T (\mathcal{X}^T \mathcal{X})^{-1} \mathbf{X}} \right) \times \gamma_{(1)}(s) ds \right]^{k-1} \gamma_{(1)}(t) dt \right\} = 1 - \alpha, \tag{5}$$

where $\gamma_{(1)}(\cdot)$ is the pdf of the smallest order statistic of m i.i.d. $\chi_{n_0-1}^2$ random variables, that is,

$$\gamma_{(1)}(t) = m\gamma(t)(1 - \Gamma(t))^{m-1},$$

with $\gamma(\cdot)$ and $\Gamma(\cdot)$ denoting the pdf and cdf of the $\chi_{n_0-1}^2$ distribution, respectively, and the expectation is taken with respect to the distribution of \mathbf{X} .

First-stage sampling: Take n_0 independent samples of each alternative i from each design point x_j through simulation and denote them by $Y_{i\ell}(x_1), \dots, Y_{i\ell}(x_m)$, $i = 1, \dots, k$, and $\ell = 1, \dots, n_0$. For each i and j , let

$$\bar{Y}_{ij} = \frac{1}{n_0} \sum_{\ell=1}^{n_0} Y_{i\ell}(x_j) \quad \text{and} \\ S_{ij}^2 = \frac{1}{n_0 - 1} \sum_{\ell=1}^{n_0} (Y_{i\ell}(x_j) - \bar{Y}_{ij})^2.$$

Second-stage sampling: Compute the total sample size $N_{ij} = \max\{\lceil h_{Het}^2 S_{ij}^2 / \delta^2 \rceil, n_0\}$ for each i and j . Take $N_{ij} - n_0$ additional independent samples from alternative i at design point x_j through simulation Y_{i,n_0+1}

$(x_j), \dots, Y_{iN_{ij}}(x_j)$, $j = 1, \dots, m$, and $i = 1, \dots, k$. For each alternative i , let

$$\hat{\beta}_i = (\mathcal{X}^T \mathcal{X})^{-1} \mathcal{X}^T \hat{\mathbf{Y}}_i,$$

where $\hat{\mathbf{Y}}_i = (\hat{Y}_{i1}, \dots, \hat{Y}_{im})^T$ and

$$\hat{Y}_{ij} = \frac{1}{N_{ij}} \sum_{\ell=1}^{N_{ij}} Y_{i\ell}(x_j).$$

Selection: Return $\hat{i}^*(\mathbf{x}) = \arg \max_{1 \leq i \leq k} \{\mathbf{x}^T \hat{\beta}_i\}$ as the selection policy.

Remark 10. The smallest order statistics in Equation (5) are introduced to make the computation of h_{Het} feasible. Without it, the equation for computing the constant h_{Het} would involve $(2m)$ -dimensional numerical integration, which becomes prohibitively difficult to solve for $m \geq 3$. The price of using the smallest order statistic is that h_{Het} is (slightly) larger than necessary, which introduces some conservativeness in the procedure. See Remark EC.2 in the e-companion for more details.

The following theorem states that procedure TS^+ is statistically valid under Assumption 2. Its proof, which is included in Section EC.4 of the e-companion, is similar to that of Theorem 1 but technically more involved. We remark here that the proof relies critically on a more generalized extension of Stein’s lemma (Stein 1945), which is stated and proved as Lemma EC.1 in the e-companion.

Theorem 2. Suppose that procedure TS^+ is used to solve the R&S-C problem and Assumption 2 is satisfied. Then, $PCS_E \geq 1 - \alpha$.

Remark 11. Let N_{TS} and N_{TS^+} denote the expected total sample sizes of procedure TS of procedure TS^+ , respectively. It can be shown that $N_{TS} = \mathcal{O}(k^{1+\frac{2}{n_0 m-d}})$ and $N_{TS^+} = \mathcal{O}(k^{1+\frac{2}{n_0-1}})$ as $k \rightarrow \infty$; meanwhile, $N_{TS} = \mathcal{O}(\alpha^{-\frac{2}{n_0 m-d}})$ and $N_{TS^+} = \mathcal{O}(\alpha^{-\frac{2}{n_0-1}})$ as $\alpha \rightarrow 0$. The proofs are provided in Section EC.7 of the e-companion. It turns out that two classical selection procedures for conventional R&S problems—the two-stage procedure in Rinott (1978) and the sequential procedure in Kim and Nelson (2001)—have upper bounds on their sample sizes that are similar to those of procedure TS and procedure TS^+ (Zhong and Hong 2021, lemma 4). In this sense, procedure TS and procedure TS^+ can solve the R&S-C problem, an extension of the conventional R&S problem, without substantially increasing the sample size complexity. Nevertheless, it must be stressed that this property relies heavily on the linear model assumption.

4.3. Comparison Between Procedure TS and Procedure TS⁺

Clearly, the assumption of homoscedasticity yields more analytical and computational tractability than the assumption of heteroscedasticity. However, if procedure TS is used in the presence of heteroscedastic errors, it may fail to deliver the desired PCS_E guarantee. An intuitive explanation is that using a single variance estimate for all the design points may underestimate the variance at some design points, leading to insufficient sampling effort at those design points.

On the other hand, procedure TS⁺ may behave in an overly conservative manner when used in the case of homoscedastic errors. This is because procedure TS⁺ requires estimation of the variances at all design points, which amounts to estimating the common variance repeatedly in the homoscedasticity setting, resulting in an excessive sampling effort. To be more specific, let us consider the estimators of the common variance σ_i^2 in procedure TS and procedure TS⁺, which are S_i^2 and S_{ij}^2 , respectively. It is easy to see that they are both unbiased estimators of σ_i^2 with the former having variance $2\sigma_i^4/(n_0m - d)$ and the latter $2\sigma_i^4/(n_0 - 1)$. Because $n_0m - d \geq n_0d - d \geq n_0 - 1$, S_{ij}^2 has a larger variance. This is not surprising as S_{ij}^2 uses just n_0 samples to estimate the variance, whereas S_i^2 uses n_0m samples. Hence, procedure TS⁺ requires more second-stage samples to compensate for the less accurate variance estimator. Furthermore, the use of the order statistic in procedure TS⁺ further loosens the lower bound of the PCS_E and results in more excessive sample sizes. These behaviors are revealed clearly through the numerical experiments in Section 6.

This discussion provides us a rule of thumb for choosing the procedures in practice. Procedure TS may be preferred if either the problem has approximately homoscedastic errors or the decision maker can tolerate some underachievement relative to the desired PCS_E. On the other hand, procedure TS⁺ may be a better choice if the errors are notably heteroscedastic or if the decision maker is stringent on delivering the PCS_E guarantee.

4.4. Handling Nonnormal Errors

We now discuss nonnormal simulation errors. We first consider the case of homoscedasticity and relax Assumption 1 to the following.

Assumption 3. For all $i = 1, \dots, k$, $\mu_i(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\beta}_i$ and $Y_i(\mathbf{x}) = \mu_i(\mathbf{x}) + \epsilon_i$, where $\boldsymbol{\beta}_i = (\beta_{i1}, \dots, \beta_{id})^T \in \mathbb{R}^d$ is a vector of unknown coefficients and the simulation error ϵ_i has mean zero and variance $\sigma_i^2 < \infty$. In addition, the simulation errors are independent among different alternatives, different covariates, and different replications.

In the absence of normality, the extended Stein's lemma (Lemma 1) does not hold. As a consequence, procedure TS does not provide finite-sample statistical validity in terms of PCS_E under Assumption 3. Instead, we establish its statistical validity in an asymptotic sense. In particular, we adopt the "small δ " regime, that is, $\delta \rightarrow 0$. This asymptotic regime is often used in R&S literature; see, for example, Kim and Nelson (2006) and Luo et al. (2015).

Note that, as $\delta \rightarrow 0$, the smallest difference that the decision maker deems worth detecting vanishes, and the R&S-C problem becomes increasingly difficult, requiring infinitely many samples eventually. Meanwhile, the central limit theorem suggests that the estimates of the linear coefficients are asymptotically normal, which would lead to the asymptotic validity of procedure TS. The proof of the following theorem is given in Section EC.5 of the e-companion.

Theorem 3. Suppose that procedure TS is used to solve the R&S-C problem and Assumption 3 is satisfied. Then, $\liminf_{\delta \rightarrow 0} \text{PCS}_E \geq 1 - \alpha$.

Furthermore, we relax Assumption 2 likewise and show that procedure TS⁺ is statistically valid asymptotically as $\delta \rightarrow 0$ as well. The proof is similar to that of Theorem 3, so we omit the details.

Assumption 4. For all $i = 1, \dots, k$, $\mu_i(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\beta}_i$ and $Y_i(\mathbf{x}) = \mu_i(\mathbf{x}) + \epsilon_i(\mathbf{x})$, where $\boldsymbol{\beta}_i = (\beta_{i1}, \dots, \beta_{id})^T \in \mathbb{R}^d$ is a vector of unknown coefficients and the simulation error $\epsilon_i(\mathbf{x})$ has mean zero and variance $\sigma_i^2(\mathbf{x}) < \infty$. In addition, the simulation errors are independent among different alternatives, different covariates, and different replications.

Theorem 4. Suppose that procedure TS⁺ is used to solve the R&S-C problem and Assumption 4 is satisfied. Then, $\liminf_{\delta \rightarrow 0} \text{PCS}_E \geq 1 - \alpha$.

5. Least Favorable Configuration

For conventional R&S problems, the so-called *least favorable configuration* (LFC) is an important concept because it defines the most difficult configuration of the means of the alternatives for the selection procedures (Bechhofer 1954). Indeed, many selection procedures are designed by analyzing the LFC. If a selection procedure can meet the target PCS under its LFC, it can certainly meet the same target for all mean configurations. It is well known that under the IZ formulation, the LFC for R&S problems is the *slippage configuration* (SC) for many procedures (Gupta and Miescke 1982). The SC is a configuration in which there exists a unique best alternative, and all other alternatives have equal means that differ from the best by exactly the IZ parameter.

To better understand our proposed procedures for R&S-C, it is important to investigate their LFCs, formally defined as follows. For given k , distribution of X , and $\sigma_i^2(x)$, $i = 1, \dots, k$, the LFC for a R&S-C procedure is the value of $\beta = (\beta_i : 1 \leq i \leq k)$ that minimizes the PCS_E of that procedure. That is,

$$LFC := \arg \min_{\beta=(\beta_i; 1 \leq i \leq k)} PCS_E(\beta),$$

where $PCS_E(\beta)$ denotes the PCS_E of the procedure under the configuration β . Note that this definition of LFC generalizes the same notion for the conventional R&S problem. If a selection procedure can meet the target PCS_E under its LFC, it will meet the same target for any other configurations.

We first generalize the SC in conventional R&S problems to the R&S-C setting and define the *generalized slippage configuration* (GSC) as follows:


$$\mu_i(x) - \mu_j(x) = \delta, \quad \text{for all } x \in \Theta \text{ and all } i = 2, \dots, k. \quad (6)$$

Under the linearity assumption (Assumption 1 or 2), the GSC becomes

$$x^T \beta_1 - x^T \beta_i = \delta, \quad \text{for all } x \in \Theta \text{ and all } i = 2, \dots, k. \quad (7)$$

Hence, under the GSC, the best alternative is the same for all $x \in \Theta$, and all other alternatives have equal mean performances. It is worth mentioning that the GSC of linear mean surfaces implies the existence of an intercept term (i.e., $X_1 \equiv 1$). Geometrically, the GSC means that the hyperplanes formed by the mean performances of the inferior alternatives are identical and parallel to the hyperplane of the best alternative, and the vertical distance between the two hyperplanes (i.e., the difference between the intercepts) is exactly δ ; see Figure 2 for an illustration for $d = 3$.

Figure 2. (Color online) Geometrical Illustration of the GSC of Linear Mean Surfaces for $d = 3$



Note. Coordinate x_1 is omitted because it is always one.

It turns out that the GSC defined in (7) is the LFC for both procedures TS and TS⁺ under the IZ formulation. We summarize this result in the following theorem. A slightly more general result is provided and proved in Section EC.8 of the e-companion.

Theorem 5. *The GSC is the LFC for procedures TS and TS⁺.*

Remark 12. Theorem 5 not only deepens our understanding of our procedures, but also helps us design numerical experiments to serve as a stress test for the proposed procedures.

6. Numerical Experiments

In this section, we investigate numerically the statistical validity of the two proposed procedures. We create a number of problem instances to test the procedures. For each problem instance, we need to specify the number of alternatives k , the dimension of the covariates d , the design matrix \mathcal{X} , the mean configuration parameters β_i , the variance configuration $\sigma_i^2(\cdot)$, and the distribution of X . Instead of specifying these aspects in a combinatorial fashion, which would result in an excessively large number of problem instances, we first create a benchmark problem and then investigate the effect of a factor by varying it while keeping the others unchanged. All the numerical studies, including the numerical experiments in this section and Section 7.2 and the case study in Section 8, are implemented in MATLAB on a desktop computer with Windows 10 OS, 3.60 GHz CPU, and 16 GB RAM. The source code is available at <https://github.com/shenhaihui/rsc>.

The benchmark problem is formulated as follows. Let $d = 4$ and $k = 5$. Suppose that $X = (1, X_2, \dots, X_d)^T$ and X_2, \dots, X_d are i.i.d. uniform[0,1] random variables. Here, the first covariate is always one, which is used to include the intercept terms for linear models. We set each except the first entry of a d -dimensional design point to be 0 or 0.5, so there are $m = 2^{d-1}$ design points in total. We set the configuration of the means to be the GSC, that is, $\beta_{11} - \delta = \beta_{i1} = 0$, $\beta_{1w} = \beta_{iw} = 1$ for $i = 2, \dots, k$ and $w = 2, \dots, d$, and set the simulation errors to be homoscedastic, particularly $\sigma_i(x) \equiv \sigma_i = 10$ for $i = 1, \dots, k$.

We then create nine test problems by varying one factor of the benchmark problem at a time while keeping other factors the same:

1. $k = 2$.
2. $k = 8$.
3. Randomly generated components of β_i from uniform[0,5], $i = 1, \dots, 5$.
4. Increasing variances (IV) configuration: $\sigma_1 = 5$, $\sigma_2 = 7.5$, $\sigma_3 = 10$, $\sigma_4 = 12.5$, and $\sigma_5 = 15$.
5. Decreasing variances (DV) configuration: $\sigma_1 = 15$, $\sigma_2 = 12.5$, $\sigma_3 = 10$, $\sigma_4 = 7.5$, and $\sigma_5 = 5$.

6. Heteroscedastic simulation errors: $\sigma_i(\mathbf{x}) = 10\mathbf{x}^\top \boldsymbol{\beta}_i$, $i = 1, \dots, 5$.
7. $d = 2$.
8. $d = 6$.
9. $X_i \sim \mathcal{N}(0.5, 1)$ truncated on $[0, 1]$ and $\text{Cov}(X_i, X_j) = 0.5$ for $i, j = 2, 3, 4$ and $i \neq j$.

Compared with the benchmark problem, problems 1 and 2 change the number of alternatives, problem 3 changes the configuration of the means so it is no longer the GSC, problems 4 and 5 change the configuration of the variances while retaining homoscedasticity, problem 6 considers heteroscedasticity, problems 7 and 8 change the dimensionality of the covariates, and problem 9 changes the distribution of the covariates.

We further create three large-scale problems:

10. $k = 100$.
11. $d = 50$.
12. $k = 100, d = 50$.

Note that problem 12 changes both k and d relative to the benchmark problem. For these three problems, instead of taking 2^{d-1} design points as before, we use the Latin hypercube sampling with $m = 2d$ design points.

In all the problem instances, we set $\alpha = 0.05$, $\delta = 1$, and $n_0 = 50$. We conduct $R = 10^4$ macroreplications for each problem–procedure combination. In each macroreplication $r = 1, \dots, R$, we apply procedures TS and TS⁺, respectively, to a problem to obtain a selection policy $\hat{\mu}_{i_r^*}(\mathbf{x})$ and then apply it to select the best alternative for each \mathbf{x}_t , a realization of \mathbf{X} that is

randomly generated from its distribution for $t = 1, \dots, T$ with $T = 10^5$. We calculate the achieved PCS_E as

$$\widehat{\text{PCS}}_E := \frac{1}{R} \sum_{r=1}^R \frac{1}{T} \sum_{t=1}^T \mathbb{I} \left\{ \mu_{i_r^*(\mathbf{x}_t)}(\mathbf{x}_t) - \mu_{\hat{i}_r^*(\mathbf{x}_t)}(\mathbf{x}_t) < \delta \right\}, \quad (8)$$

where $\mathbb{I}\{\cdot\}$ denotes the indicator function. We also report the average total sample size used by each procedure for producing the selection policy.

The numerical results are shown in Table 1, from which we have the following observations. First, as expected, both procedures can deliver the target PCS_E in their respective domains. Procedure TS can deliver the designed PCS_E if the simulation errors are homoscedastic, and procedure TS⁺ can deliver the designed PCS_E even when the simulation errors are heteroscedastic. Moreover, the achieved PCS_E is higher than the target in general; see, for example, the column “ $\widehat{\text{PCS}}_E$ ” under “Procedure TS” of Table 1, except the entry for problem 6. This is especially the case if the configuration of the means is not the GSC, that is, problem 3. Overshooting the target PCS_E suggests that the total sample size is larger than necessary for meeting the target PCS_E. Such conservativeness is a well-known issue for R&S procedures under the IZ formulation; see Fan et al. (2016) for an exposition on the issue.

Second, if procedure TS is applied to the instance of heteroscedasticity, (i.e., problem 6), the target PCS_E cannot be met. By contrast, if procedure TS⁺ is applied to the instances of homoscedasticity, (i.e., all problem instances except 6), it becomes overly conservative

Table 1. Results When the Target Is PCS_E ≥ 95%

Problem	Procedure TS			Procedure TS ⁺		
	h	Sample size	$\widehat{\text{PCS}}_E$	h_{Het}	Sample size	$\widehat{\text{PCS}}_E$
0 Benchmark	3.423	46,865	0.9610	4.034	65,138	0.9801
1 $k = 2$	2.363	8,947	0.9501	2.781	12,380	0.9702
2 $k = 8$	3.822	93,542	0.9650	4.510	130,200	0.9842
3 Non-GSC	3.423	46,865	0.9987	4.034	65,138	0.9994
4 IV	3.423	52,698	0.9618	4.034	73,265	0.9807
5 DV	3.423	52,720	0.9614	4.034	73,246	0.9806
6 Het	3.423	58,626	0.9232	4.034	81,555	0.9846
7 $d = 2$	4.612	21,288	0.9593	4.924	24,266	0.9662
8 $d = 6$	2.141	73,428	0.9656	2.710	117,626	0.9895
9 Normal Dist	3.447	47,529	0.9626	4.063	66,061	0.9821
10 $k = 100$	4.346	1,133,384	0.9758	5.117	1,570,911	0.9918
11 $d = 50$	3.222	508,977	0.9583	4.312	911,326	0.9926
12 $k = 100, d = 50$	4.886	23,400,677	0.9765	6.702	44,024,486	0.9991

Notes. (i) In the presence of heteroscedasticity, the boxed number suggests that procedure TS fails to deliver the target PCS_E, whereas the bold number suggests that procedure TS⁺ succeeds in doing so. (ii) In these experiments, the sampling effort is negligible because it only involves generating normally distributed errors. Thus, the run time for producing the selection policy by each procedure reflects the computational overhead of each procedure. It is found that, even for the relatively large-scale problem 12, the run time is shorter than one second, which indicates negligible overhead.

compared with procedure TS. This is reflected by the achieved PCS_E being substantially higher than the target and the sample size being substantially larger than that of procedure TS.

Third, as the number of alternative k increases, which corresponds to problems 1, 0, and 2, the sample size allocated to each alternative on average (measured by the ratio of the total sample size to k) increases as well. This is caused by the increase in the constant h as k increases. Notice that the sample size required for alternative i on one design point in procedure TS is $N_i = \max\{[h^2 S_i^2 / \delta^2], n_0\}$. Thus, a larger h means a larger N_i . A similar argument holds for procedure TS^+ as well. This suggests that, as k increases, each alternative must be estimated more accurately in order to differentiate them.

Fourth, the numerical results of problems 4 and 5 are almost identical. In particular, the value of h is identical for both problems because the equations that determine h (Equation (4)) and h_{Het} (Equation (5)) do not depend on the configuration of the variances. Then, as the sum of the variances is the same for both problems, the total sample size that is approximately proportional to h^2 times the sum of the variances is almost the same for both problems.

Finally, the results for problems 10–12 are also as expected. They show that both procedure TS and procedure TS^+ can be used to handle relatively large-scale problems. Note that h and h_{Het} , the key quantities for determining the second-stage sample size of the two procedures, respectively, can be computed via the Monte Carlo method or the stochastic approximation method; see further discussion in Section EC.2 of the e-companion. Therefore, the computational requirement for determining the two quantities and, thus, the total sample size as well as the sample allocation is negligible relative to the expenses of running the simulation model.

7. Experimental Design and Robustness to Linearity Assumptions

We have assumed so far that the design points are given with the design matrix \mathcal{X} satisfying that $\mathcal{X}^T \mathcal{X}$ is nonsingular. In this section, we discuss how to select the design points. We show that the extreme design, that is, locating the design points at the corners of the design region Θ , is typically a good strategy under the linearity assumptions (e.g., Assumptions 1–4). In practice, however, linearity assumptions are often satisfied only approximately. Then, the selection of design points is critically related to the robustness of the linearity assumptions. We show through numerical experiments that the extreme design may perform poorly when the linearity assumptions are violated mildly, but distributing the design points

evenly in the design region Θ appears to be quite robust to the linearity assumptions.

7.1. Optimal Design Under Linearity Assumptions

Experimental design is a classical problem in statistics. In classical design for linear regression, the objective is often to choose a design that optimizes a certain criterion given a fixed total sample size. Popularly used criteria include D-optimal design that minimizes the determinant of the covariance matrix of the OLS estimator of β , G-optimal design that minimizes the maximal variance of the fitted response over the design region, and many others; see Silvey (1980, chapter 2) for more details on the subject. Some of the optimal designs are equivalent under certain conditions. For instance, Kiefer and Wolfowitz (1960) prove that the D- and G-optimal designs are equivalent in the continuous case (also called the approximate case) in which the integer constraint on the sample size at each design point is relaxed; see Silvey (1980, chapter 3) for a more careful and complete discussion on the general equivalence theory.

However, the optimal design in the R&S-C context is different from the classical ones. In our procedures, an optimal design is the design that minimizes the total sample size required by the procedures to deliver the predetermined PCS_E . Using procedure TS as an example, the total sample size is $\sum_{i=1}^k N_i m$, where N_i is approximately $h^2 S_i^2 / \delta^2$. Because the design matrix $\mathcal{X} = (\mathbf{x}_1, \dots, \mathbf{x}_m)^T$, we may formulate the optimal design problem as the following optimization problem:

$$\begin{aligned} & \min_{m, \mathbf{x}_1, \dots, \mathbf{x}_m} h^2 m \\ & \text{s.t. } \mathbb{E} \left\{ \int_0^\infty \left[\int_0^\infty \Phi \left(\frac{h}{\sqrt{(n_0 m - d)(t^{-1} + s^{-1}) \mathbf{X}^T (\mathcal{X}^T \mathcal{X})^{-1} \mathbf{X}}} \right)} \right. \right. \\ & \quad \left. \left. \times \eta(s) ds \right]^{k-1} \eta(t) dt \right\} = 1 - \alpha, \\ & \text{rank}(\mathcal{X}^T \mathcal{X}) = d, \\ & m \geq d, \text{ integer}, \\ & \mathbf{x}_1, \dots, \mathbf{x}_m \in \Theta, \end{aligned}$$

where the first constraint is exactly (4) and the second constraint ensures the nonsingularity of $\mathcal{X}^T \mathcal{X}$. The problem is, in general, a nonconvex integer programming problem, and it is difficult to solve. Moreover, even without concerning the integer constraint, the optimal design is a function of the distribution of \mathbf{X} and is often difficult to characterize. In this section, we derive an optimal design, which is invariant to the distribution of \mathbf{X} , for a simplified case in which m is fixed and an additional constraint is imposed.

Specifically, we assume that

$$\Theta = \{1\} \times [l_2, u_2] \times \cdots \times [l_d, u_d],$$


$$d \geq 2 \text{ and } l_w < u_w \text{ for all } w = 2, \dots, d. \quad (9)$$

Here, the first covariate is always one, which is used to take care of the intercept terms in the linear models, and all other $d - 1$ covariates are in an interval $[l_w, u_w]$ for $w = 2, \dots, d$. Notice that all the following analysis can also apply to the linear models without intercept terms with similar arguments. Suppose that we want to allocate $m = b(2^{d-1})$ design points in Θ , where $b \geq 1$ is a fixed integer. We denote these design points by $\mathbf{a}_i^j \in \Theta$, $i = 1, \dots, 2^{d-1}$, and $j = 1, \dots, b$. Moreover, for any j , we let $\mathbf{a}_1^j, \dots, \mathbf{a}_{2^{d-1}}^j$ be symmetric with respect to the Cartesian coordinate system located at the center of Θ . See Figure 3 for an illustration with $d = 3$ and $b = 2$. Let $\mathcal{S}^j := \{\mathbf{a}_1^j, \dots, \mathbf{a}_{2^{d-1}}^j\}$ for $j = 1, \dots, b$. Then, $\{\mathcal{S}^1, \dots, \mathcal{S}^b\}$ denotes a set of $b(2^{d-1})$ design points (in which duplicates are allowed), and we call it a *symmetric design*. Notice that the symmetric design ensures that $\text{rank}(\mathcal{X}^T \mathcal{X}) = d$. The reason that we only consider symmetric designs is, without considering the distribution of the covariates \mathbf{X} , symmetric designs are natural choices because of the symmetric nature of the design region Θ .

Let \mathcal{S}^0 denote the set of corner points of Θ . It is easy to see that \mathcal{S}^0 has 2^{d-1} elements. A simple design is to use all the points in \mathcal{S}^0 for b times, that is, $\mathcal{S}^1 = \cdots = \mathcal{S}^b = \mathcal{S}^0$, and we call it the *extreme design*. Notice that the extreme design aims to spread out all the design points so that the OLS estimators of β_i can have small variances.

The extreme design is also a symmetric design. In the following theorem, we show that the extreme design is the best symmetric design regardless of

Figure 3. (Color online) Geometrical Illustration of the Symmetric Design for $d = 3$ and $b = 2$



Note. Coordinate x_1 is omitted because it is always one.

the distribution of \mathbf{X} . The proof is included in Section EC.9 of the e-companion.

Theorem 6. Suppose that Assumption 1 or 3 holds, procedure TS is used to solve an R&S-C problem, and $m = b(2^{d-1})$ design points are allocated in Θ as assumed in (9). Then, among all symmetric designs, the extreme design $\mathcal{S}^1 = \cdots = \mathcal{S}^b = \mathcal{S}^0$ minimizes the expected total sample size.

There is an interesting link between the optimal design in R&S-C with that in the classic linear regression setting. That is, the extreme design is also both the D-optimal and G-optimal designs in linear regression when the total sample size is $b(2^{d-1})$ among all feasible designs (without the symmetry constraint). This result is formally stated in Theorem 7, and its proof is included in Section EC.10 of the e-companion, in which the formal definitions of D- and G-optimality are also given. We want to emphasize that Theorem 7 further justifies the consideration of the extreme designs for R&S-C problems.

Theorem 7. Consider the linear regression problem $Y(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\beta} + \epsilon$, where $\boldsymbol{\beta}, \mathbf{x} \in \mathbb{R}^d$ and ϵ is random error with mean zero and variance σ^2 . Let $d \geq 2$ and $x_1 \equiv 1$ so that the intercept term is included. Suppose that $m = b(2^{d-1})$ design points are allocated in Θ as assumed in (9). Then, among all feasible designs, the extreme design $\mathcal{S}^1 = \cdots = \mathcal{S}^b = \mathcal{S}^0$ is both D- and G-optimal.

Remark 13. The total sample size of procedure TS⁺ depends on the variances of the design points, which are not known a priori. Therefore, we can only prove that, among all symmetric designs, the extreme design minimizes the constant h_{Het} defined in Equation (5).

7.2. Robustness to Linearity Assumptions

In practice, the linearity assumptions (i.e., Assumptions 1–4) often hold only approximately. Notice that the linear models can be generalized to capture nonlinearity in $\mu_i(\mathbf{x})$ by the use of basis functions; see Remark 5. Here, by saying that the linearity assumption does not hold, we actually mean that $\mu_i(\mathbf{x})$ is not linear in \mathbf{x} , and we do not have a proper set of basis functions to perform a change of variables.

It is argued in James et al. (2013) that linear models are often robust to nonlinear behaviors and lead to good predictions. However, the extreme design is, in general, not robust to nonlinearity because it allocates no design points in the interior of the design region and leaves the fitted models depending completely on the corner points. To improve the robustness of the experimental design, one can allocate design points evenly in the design region. One such design that is widely used is the so-called minimax design, which, roughly speaking, ensures that all points in the design region are not too far from the nearest design points. It is shown by Johnson et al. (1990) that the minimax

design is asymptotically equivalent to the Bayesian G-optimal design for general Gaussian processes, which mimics the classical G-optimal design but is defined in a Bayesian framework. In the rest of this section, we conduct numerical studies to compare the extreme design and the minimax design and to understand their behaviors under different scenarios.

We consider the case in which $\Theta = \{1\} \times [0, 1]^{d-1}$ and generate the true surfaces randomly from a $(d - 1)$ -dimensional second-order stationary Gaussian random field with mean $\mathbf{0}$ and isotropic covariance $\text{Cov}(z, z') = \exp\{-\lambda\|z - z'\|^2\}$ for $z, z' \in \mathbb{R}^{d-1}$, where $\|\cdot\|$ represents the Euclidean norm. Notice that parameter λ controls the scale of the random field, and larger λ often leads to a higher level of violation of the linearity assumption. To keep the linearity assumptions approximately true, we discretize the surfaces with a step size 0.01 for each coordinate, calculate the R^2 of the discretized observations, and only keep the surfaces whose R^2 is above 0.8. We first obtain 50 such approximately linear random surfaces. Then, we randomly create 100 R&S-C problems, each with five surfaces that are randomly drawn from those 50 surfaces. We consider only the homoscedastic errors and add normal noises with $\sigma_1^2 = \dots = \sigma_5^2 = 1$.


We consider $\lambda = 0.5$ and $\lambda = 3$ to capture small and large violations of the linearity assumption. We also consider $d = 2$ and $d = 3$. Figure 4 shows the typical shapes of these randomly generated surfaces. For each R&S-C problem, we let X_2, \dots, X_d be i.i.d. uniform $[0, 1]$ random variables and set $\alpha = 0.05$, $\delta = 0.2$,

$n_0 = 50$, $R = 10^3$, and $T = 10^4$. We compare the extreme design and minimax design with $2(2^{d-1})$ points, that is, four design points when $d = 2$ and eight design points when $d = 3$. The design matrices are listed in Table 2. We report the means and standard deviations (SD) of the average total sample size and the achieved PCSE (i.e., $\overline{\text{PCSE}}$) over 100 problems in Table 3. We also calculate the average *regret* (also called the opportunity cost in Bayesian R&S literature), which is defined as $\frac{1}{R} \sum_{r=1}^R \frac{1}{T} \sum_{t=1}^T \{\mu_{i^*(x_t)}(x_t) - \mu_{\hat{i}^*(x_t)}(x_t)\}$. The means and SD of the average regrets are also reported in Table 3.

From Table 3, we see that the extreme designs lead to significantly smaller total sample sizes than the minimax designs if the linearity assumption is more or less satisfied (e.g., $\lambda = 0.5$), but their achieved PCSE and regrets are significantly poorer than those of the minimax designs if the linearity assumption is more violated (e.g., $\lambda = 3$). Based on these observations, we have the following conclusions on experimental design and robustness on linearity assumptions.

- The proposed procedures perform well when the surfaces are approximately linear though the statistical guarantee may not always hold.
- When the true surfaces are exactly linear or only slightly nonlinear, the extreme design is preferred because it requires fewer samples to deliver the required PCSE .
- When the true surfaces are relatively nonlinear, even designs, such as the minimax design, are preferred because they are more robust to nonlinearity.

Figure 4. (Color online) Randomly Generated Surfaces with $R^2 \geq 0.8$



Note. In the general case, nonlinear surfaces do not necessarily mean that the linearity assumption is violated.

Table 2. Extreme Designs and Minimax Designs for $d = 2, 3$

$d = 2$		$d = 3$	
Extreme design	Minimax design	Extreme design	Minimax design ^a
$\begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 0 \\ 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 1/8 \\ 1 & 3/8 \\ 1 & 5/8 \\ 1 & 7/8 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0.1557 & 0.2086 \\ 1 & 0.1557 & 0.7914 \\ 1 & 0.8443 & 0.2086 \\ 1 & 0.8443 & 0.7914 \\ 1 & 0.2468 & 0.5000 \\ 1 & 0.7532 & 0.5000 \\ 1 & 0.5000 & 0.1794 \\ 1 & 0.5000 & 0.8206 \end{pmatrix}$

^aSee Melissen and Schuur (1996).

• This intuition suggests that, when the design region Θ is of a general shape other than a hyper-rectangle, it is better to allocate the design points far from each other if the linearity is strong and more evenly in the region if the linearity is weak.

8. A Case Study: Personalized Treatment for Cancer Prevention

Esophageal cancer (see Figure 5) is the seventh-leading cause of cancer death among males (making up 4%) in the United States, according to Cancer Facts & Figures 2016 (American Cancer Society 2016). Esophageal adenocarcinoma (EAC) is a main subtype of esophageal cancer, and its incidence has increased 500% over the past 40 years (Hur et al. 2013, Choi et al. 2014). Thus, the management of BE, a precursor to EAC, is an active topic in cancer research. A common strategy for BE management is endoscopic surveillance, which attempts to prevent EAC through dysplasia treatment or to identify EAC before it becomes invasive. Recently, chemoprevention has received substantial attention as a method to lower the progression of BE to EAC, and aspirin and statins are two particular drugs that are demonstrated to be effective (Kastelein et al. 2011). For each BE patient, the progression rate to cancer depends on a variety of factors, including

age; weight; lifestyle habits, such as smoking and alcohol use; the grade of dysplasia, etc. In addition, each patient may have a different response to drugs depending on drug resistance and tolerance. Hence, it is conceivable that the best treatment regimen for BE is patient-specific.


We formulate the problem of selecting the best treatment regimen for each BE patient as an R&S-C problem. There are three alternatives: endoscopic surveillance only ($i = 1$), aspirin chemoprevention with endoscopic surveillance ($i = 2$), and statin chemoprevention with endoscopic surveillance ($i = 3$). For simplicity, we consider only the starting age of a treatment regimen, risk (i.e., the annual progression rate of BE to EAC), and drug effects (i.e., the progression reduction effect of a drug) as patient characteristics that determine the effectiveness of a treatment regimen. More specifically, the vector of covariates is $X = (1, X_1, X_2, X_3, X_4)^T$, where X_1 is the starting age, X_2 is the risk, and X_3 and X_4 are the drug effects of aspirin and statin, respectively. We use the expected QALYs as the performance measure to compare different alternatives.

To solve this R&S-C problem, we need a model to simulate the QALYs of the treatment regimens for different patients. Fortunately, a discrete-time Markov

Table 3. Means and SD (in Parentheses) over 100 Problems

Case	R^2	Extreme design			Minimax design		
		Sample size	\widehat{PCS}_E	Regret	Sample size	\widehat{PCS}_E	Regret
$d = 2, \lambda = 0.5$	0.965 (0.002)	1,730 (2)	0.9948 (0.0086)	0.007 (0.008)	2,869 (5)	0.9978 (0.0037)	0.005 (0.005)
$d = 2, \lambda = 3$	0.921 (0.003)	1,730 (2)	0.8558 (0.1394)	0.100 (0.109)	2,941 (50)	0.9799 (0.0297)	0.013 (0.016)
$d = 3, \lambda = 0.5$	0.917 (0.003)	2,282 (70)	0.9528 (0.0586)	0.024 (0.027)	4,659 (16)	0.9876 (0.0118)	0.008 (0.007)
$d = 3, \lambda = 3$	0.863 (0.002)	2,425 (120)	0.7358 (0.1306)	0.204 (0.139)	4,904 (96)	0.9133 (0.0502)	0.047 (0.030)

Figure 5. (Color online) Diagram of Esophageal Cancer




chain model developed by Hur et al. (2004) and Choi et al. (2014) may be used. The model simulates the transitions among different health states of a BE patient until death, and the transition diagram of the model is shown in Figure 6. The transition probability matrices are well calibrated so that the simulation outputs match the published results. We adopt this model to simulate individual patients with specific characteristics, which are defined by the covariates X and assumed to be observable. This Markov chain model is, of course, a highly simplified model compared with those having more detailed biological mechanisms (see, for example, the multistage clonal expansion EAC model of Curtius et al. 2015). However, as an illustrative purpose, we adopt this simple model because of its accessibility and relatively short running time because we need to run the model in a brute force way to obtain the mean performance surfaces of all alternative, that is, $\mu_i(x)$ for $i = 1, 2, 3$, and use them as the true values to evaluate the performance of the proposed procedures.

In this case study, we assume that the distributions of the covariates are known because there are often ample historical data to calibrate these distributions in practice. Furthermore, we specify these distributions as follows: We assume $X_1 \in [55, 80]$ as it is documented by Naef and Savary (1972) that there is a BE incidence peak for individuals with ages within this range. We assume $X_2 \in [0, 0.1]$ following the specification in Hur et al. (2004) and set $X_3 \in [0, 1]$ and $X_4 \in [0, 1]$ by definition. Moreover, we assume $\mathbb{E}[X_3] = 0.53$ and $\mathbb{E}[X_4] = 0.54$ following the study by Kastelein et al. (2011). Nevertheless, because of a lack of detailed data, we do not know the distribution of covariates exactly among the entire population of BE patients. Instead, we suppose that X_1, \dots, X_4 are independent, and their distributions are specified in Table 4. The design points are specified as follows. We take X_1 from $\{61, 74\}$, X_2 from $\{0.1/4, 0.3/4\}$, X_3 from $\{1/4, 3/4\}$, and X_4 from $\{1/4, 3/4\}$ and then combine them in a full factorial way. Therefore, it is a relatively even design with 16 design points.

Before carrying out the R&S-C procedures, we conduct several trial runs of the simulation model, and we find that the linearity assumptions hold approximately and the simulation errors are clearly heteroscedastic. Therefore, procedure TS^+ is used. Notice that, to calculate the achieved PCS_E (i.e., \widehat{PCS}_E) of our procedure, we need the true response surfaces $\mu_i(x)$ for all $x \in \Theta$ and $i = 1, 2, 3$ to identify the true best selection policy $i^*(x)$. To that end, we use extensive simulation to approximate the true response surfaces. We discretize X_2 with a step size of 0.01 and discretize X_3 and X_4 with a step size of 0.1. At each discretization point, we run the simulation model for 10^6 replications so that the estimation error is negligible (e.g., the half width of the 95% confidence interval is less than 0.02 QALYs). The response at any other x is approximated via a linear interpolation. To compute \widehat{PCS}_E , we conduct $R = 300$ macroreplications. For each macroreplication r , we apply procedure TS^+ to obtain the selection policy $\widehat{i}_r^*(x)$ and then apply it to

Figure 6. (Color online) Transition Diagram of the Markov Simulation Model



Notes. (1) A person in each state may die from age-related, all-cause mortality. (These transitions are omitted in the diagram.) (2) The time duration between state transitions is one month. (3) The details of the state transitions inside the dotted box depends on whether aspirin chemoprevention or statin chemoprevention is used.


Table 4. Distributions of the Covariates

Covariate	Distribution	Support	Mean
X_1	Discrete (Figure 7)	$\{55, \dots, 80\}$	64.78
X_2	Uniform (0, 0.1)	$[0, 0.1]$	0.05
X_3	Triangular (0, 0.59, 1)	$[0, 1]$	0.53
X_4	Triangular (0, 0.62, 1)	$[0, 1]$	0.54

select the best treatment regimen for $T = 10^5$ simulated BE patients whose characteristics are randomly generated from the distribution of \mathbf{X} . Other parameters of procedure TS^+ are specified as follows: $\alpha = 0.05$, $\delta = 1/6$ (i.e., two months), and $n_0 = 100$. Our case study shows that the $\widehat{PCS}_E = 99.5\%$, which is substantially higher than the target level $1 - \alpha = 95\%$. This is because the configuration of the means of this problem is much more favorable than the GSC, and thus, the selection procedure behaves in an overly conservative manner in this situation. (Recall that problem 3 in Section 6 has a similar behavior.)

Remark 14. In principle, one could compute the “true” response surfaces of this simulation model that correspond to the three treatment regimens in a brute force way subject to a discretization scheme and then identify the best alternative for each individual patient. However, this would be very time-consuming even for a coarse discretization scheme and a moderate level of estimation accuracy as specified. (It takes about eight days on a desktop computer with Windows 10 OS, 3.60 GHz CPU, and 16 GB RAM to complete the simulation implemented in MATLAB.) By contrast, it takes less than one minute for procedure TS^+ to obtain a selection policy. This demonstrates the practical value of our model and selection procedures in real-world applications.

Figure 7. (Color online) Probability Mass Function of X_1 (Truncated)




Note. Data source: U.S. 2016 population data, U.S. Census Bureau.

To demonstrate the usefulness of R&S-C as a decision-making framework, we compare the personalized approach with a more traditional approach, which selects the treatment regimen that is the best for the entire population, that is, $i^+ = \arg \max_{1 \leq i \leq 3} \mathbb{E}[\mu_i(\mathbf{X})]$. The latter corresponds to a conventional R&S approach. In this problem, we find $i^+ = 3$, which indicates that alternative 3 is better than the others based on the population average. Notice that choosing the population best, that is, always selecting alternative 3, can also be regarded as a selection policy. Based on our numerical study, we find that this policy corresponds to a \widehat{PCS}_E of 75.8%; that is, alternative 3 is indeed the best or within the IZ for 75.8% of the population. In contrast, the personalized approach that we report earlier has a \widehat{PCS}_E of 99.5%. The 23.7% difference in \widehat{PCS}_E demonstrates clearly the advantage of the personalized approach.

In addition to \widehat{PCS}_E , we consider QALYs regret as another criterion to compare the two approaches. More specifically, we consider the expected QALYs regret, which is the expected difference between the QALYs under the true optimal treatment regimen and the selected one by each approach. Conditionally on $\mathbf{X} = \mathbf{x}$, the expected regret is $\mu_{i^*(\mathbf{x})}(\mathbf{x}) - \mu_3(\mathbf{x})$ for the traditional approach and $\mu_{i^*(\mathbf{x})}(\mathbf{x}) - \mu_{\widehat{i}^*(\mathbf{x})}(\mathbf{x})$ for the personalized approach, where $\widehat{i}^*(\mathbf{x})$ comes from one macroreplication of procedure TS^+ . The results are plotted in Figure 8, in which the left panel shows the distribution of regret for the entire BE population (i.e., $\mathbf{X} \in \Theta$), and the right panel shows the distribution of regret for a specific group of patients (i.e., $X_3 = 0.9$ and $X_4 = 0.2$).

From these results, we see that using the personalized approach (i.e., the R&S-C approach), BE patients have much lower expected QALYs regret than using the traditional approach (i.e., the conventional R&S approach). Among the entire BE population (left panel of Figure 8), when the personalized approach is used, more than 99% of the patients have either no regret or a regret that is less than or equal to two months (i.e., the IZ parameter). However, when the traditional approach is used, close to a quarter (i.e., 24%) of the patients have a regret that is more than two months, and 2% of them have a regret that is more than 12 months.

Figure 8. (Color online) Bar Charts of $\mathbb{E}[\text{QALYs}|X]$ Regret Under the Selected Treatment Regimen

Notes. Left: Entire population, $X \in \Theta$. Right: Specific population, $X = (1, X_1, X_2, 0.9, 0.2)^T$.

If we look at the specific group of patients, for example, the group as considered in the right panel of Figure 8, we see that the reduction of the regret using the personalized approach is even more substantial, which demonstrates the key point of personalized medicine, that is, a universal treatment, even it seems fairly good for the entire population, may perform quite poorly for certain groups of patients, and we can do much better with the help of personalized medicine.

9. Conclusions

Ranking and selection is a long-standing research problem in simulation literature. The emerging popularity of personalized decision making leads us to consider this classic problem in a new environment in which the performance of an alternative depends on some observable random covariates. A critical feature in the new setting is that the goal is not to seek a single alternative having a superior performance, but a selection policy as a function of the covariates. Albeit computed off-line via simulation model, the selection policy can be applied online to specify the best alternative for the subsequent individuals after observing their covariates. Therefore, R&S-C reflects a shift in perspective regarding the role of simulation: a tool for system control instead of system design. In particular, we demonstrate the practical value of R&S-C via a case study of personalized medicine for selecting the best treatment regimen in prevention of esophageal cancer.

This paper uses a linear model to capture the relationship between the response of an alternative and the covariates and develops two-stage selection procedures accordingly under the IZ formulation. However, the presence of covariates complicates the concept of PCS because the best alternative varies as a function of the covariates. We define the statistical validity of a procedure in terms of average PCS although other forms of unconditional PCS are also

possible. This paper is a first step toward understanding R&S-C problems under a frequentist perspective. There are many potential directions for future work, such as nonparametric models and sequential selection procedures.

Acknowledgments

The authors thank the area editor, Professor Bruno Tuffin, the associate editor, and anonymous referees for their insightful and detailed comments that have significantly improved this paper.

References

- American Cancer Society (2016) Cancer facts & figures. Report, American Cancer Society, Atlanta.
- Arora N, Dreze X, Ghose A, Hess JD, Iyengar R, Jing B, Joshi YV, et al (2008) Putting one-to-one marketing to work: Personalization, customization, and choice. *Marketing Lett.* 19(3):305–321.
- Auer P (2002) Using confidence bounds for exploitation-exploration trade-offs. *J. Machine Learn. Res.* 3:397–422.
- Bechhofer RE (1954) A single-sample multiple decision procedure for ranking means of normal populations with known variances. *Ann. Math. Statist.* 25(1):16–39.
- Bubeck S, Cesa-Bianchi N (2012) *Regret Analysis of Stochastic and Nonsmooth Multi-armed Bandit Problems*. Foundations and Trends in Machine Learning, vol. 5, no. 1 (Now Publishers Inc., Boston).
- Chen HC, Chen CH, Dai L, Yücesan E (1997) New development of optimal computing budget allocation for discrete event simulation. Andradóttir S, Healy KJ, Withers DH, Nelson BL, eds. *Proc. 1997 Winter Simulation Conf.* (IEEE Computer Society, Washington, DC), 334–341.
- Chen CH, Chick SE, Lee LH, Pujowidianto NA (2015) Ranking and selection: Efficient simulation budget allocation. Fu MC, ed., *Handbook of Simulation Optimization* (Springer, New York), 45–80.
- Chick SE, Frazier PI (2012) Sequential sampling with economics of selection procedures. *Management Sci.* 58(3):550–569.
- Chick SE, Gans N (2009) Economic analysis of simulation selection problems. *Management Sci.* 55(3):421–437.
- Chick SE, Inoue K (2001) New two-stage and sequential procedures for selecting the best simulated system. *Oper. Res.* 49(5): 732–743.
- Choi SE, Perzan KE, Tramontano AC, Kong CY, Hur C (2014) Statins and aspirin for chemoprevention in Barrett's esophagus: Results of a cost-effectiveness analysis. *Cancer Prevention Res.* 7(3): 341–350.

- Curtius K, Hazelton WD, Jeon J, Luebeck EG (2015) A multiscale model evaluates screening for neoplasia in Barrett's esophagus. *PLOS Comput. Biol.* 11(5):e1004272.
- Dudewicz EJ, Dalal SR (1975) Allocation of observations in ranking and selection with unequal variances. *Sankhyā: Indian J. Statist., Ser. B* 37(1):28–78.
- Fan W, Hong LJ, Nelson BL (2016) Indifference-zone-free selection of the best. *Oper. Res.* 64(6):1499–1514.
- Frazier PI (2014) A fully sequential elimination procedure for indifference-zone ranking and selection with tight bounds on probability of correct selection. *Oper. Res.* 62(4):926–942.
- Frazier PI, Powell WB, Dayanik S (2008) A knowledge-gradient policy for sequential information collection. *SIAM J. Control Optim.* 47(5):2410–2439.
- Goldenshluger A, Zeevi A (2013) A linear response bandit problem. *Stochastic Systems* 3(1):230–261.
- Gupta SS, Miescke KJ (1982) On the least favorable configurations in certain two-stage selection procedures. Kallianpur G, Krishnaiah PR, Ghosh JK, eds. *Statistics and Probability: Essays in Honor of C. R. Rao* (North Holland, Amsterdam), 295–305.
- Hong LJ (2006) Fully sequential indifference-zone selection procedures with variance-dependent sampling. *Naval Res. Logist.* 53(5):464–476.
- Hu R, Ludkovski M (2017) Sequential design for ranking response surfaces. *SIAM/ASA J. Uncertainty Quantification* 5(1):212–239.
- Hur C, Nishioka NS, Gazelle GS (2004) Cost-effectiveness of aspirin chemoprevention for Barrett's esophagus. *J. National Cancer Inst.* 96(4):316–325.
- Hur C, Miller M, Kong CY, Dowling EC, Nattinger KJ, Dunn M, Feuer EJ (2013) Trends in esophageal adenocarcinoma incidence and mortality. *Cancer* 119(6):1149–1158.
- James G, Witten D, Hastie T, Tibshirani R (2013) *An Introduction to Statistical Learning*, vol. 112 (Springer, New York).
- Johnson ME, Moore LM, Ylvisaker D (1990) Minimax and maximin distance designs. *J. Statist. Planning Inference* 26(2):131–148.
- Kastelein F, Spaander MCW, Biermann K, Steyerberg EW, Kuipers EJ, Bruno MJ, Probar-Study Group (2011) Nonsteroidal anti-inflammatory drugs and statins have chemopreventative effects in patients with Barrett's esophagus. *Gastroenterology* 141(6):2000–2008.
- Katrakazas C, Quddus M, Chen WH, Deka L (2015) Real-time motion planning methods for autonomous on-road driving: State-of-the-art and future research directions. *Transportation Res. C: Emerging Tech.* 60:416–442.
- Kiefer J, Wolfowitz J (1960) The equivalence of two extremum problems. *Canadian J. Math.* 12:363–366.
- Kim SH, Nelson BL (2001) A fully sequential procedure for indifference-zone selection in simulation. *ACM Trans. Model. Comput. Simulation* 11(3):251–273.
- Kim SH, Nelson BL (2006) Selecting the best system. Henderson SG, Nelson BL, eds. *Handbooks in Operations Research and Management Science*, vol. 13 (Elsevier, Amsterdam), 501–534.
- Kim ES, Herbst RS, Wistuba II, Lee JJ, Blumenschein GR Jr, Tsao A, Stewart DJ, et al (2011) The BATTLE trial: Personalizing therapy for lung cancer. *Cancer Discovery* 1(1):44–53.
- Law AM, Kelton WD (2000) *Simulation Modeling and Analysis*, 3rd ed. (McGraw-Hill, New York).
- Luo J, Hong LJ, Nelson BL, Wu Y (2015) Fully sequential procedures for large-scale ranking-and-selection problems in parallel computing environments. *Oper. Res.* 63(5):1177–1194.
- Melissen JBM, Schuur PC (1996) Improved coverings of a square with six and eight equal circles. *Electronic J. Combin.* 3(1):1–10.
- Naef A, Savary M (1972) Conservative operations for peptic esophagitis with stenosis in columnar-lined lower esophagus. *Ann. Thoracic Surgery* 13(6):543–551.
- Negoescu DM, Frazier PI, Powell WB (2011) The knowledge-gradient algorithm for sequencing experiments in drug discovery. *INFORMS J. Comput.* 23(3):346–363.
- Ni EC, Ciocan DF, Henderson SG, Hunter SR (2017) Efficient ranking and selection in parallel computing environments. *Oper. Res.* 65(3):821–836.
- Pearce M, Branke J (2017) Efficient expected improvement estimation for continuous multiple ranking and selection. Chan WK, D'Ambrogio A, Zacharewicz G, Mustafee N, Wainer G, Page EH, eds. *Proc. 2017 Winter Simulation Conf.* (IEEE Press, Piscataway, NJ), 2161–2172.
- Perchet V, Rigollet P (2013) The multi-armed bandit problem with covariates. *Ann. Statist.* 41(2):693–721.
- Rencher AC, Schaalje GB (2008) *Linear Models in Statistics*, 2nd ed. (John Wiley & Sons, Inc., Hoboken, NJ).
- Rinott Y (1978) On two-stage selection procedures and related probability-inequalities. *Comm. Statist. Theory Methods* 7(8):799–811.
- Robbins H (1952) Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc. (N.S.)* 58(5):527–535.
- Robbins H, Monro S (1951) A stochastic approximation method. *Ann. Math. Statist.* 22(3):400–407.
- Rusmevichientong P, Tsitsiklis JN (2010) Linearly parameterized bandits. *Math. Oper. Res.* 35(2):395–411.
- Shen H, Hong LJ, Zhang X (2017) Ranking and selection with covariates. *Proc. 2017 Winter Simulation Conf.*, 2137–2148.
- Silvey SD (1980) *Optimal Design: An Introduction to the Theory for Parameter Estimation* (Chapman and Hall, London).
- Slivkins A (2014) Contextual bandits with similarity information. *J. Machine Learn. Res.* 15(1):2533–2568.
- Stein C (1945) A two-sample test for a linear hypothesis whose power is independent of the variance. *Ann. Math. Statist.* 16(3):243–258.
- Yap TA, Carden CP, Kaye SB (2009) Beyond chemotherapy: Targeted therapies in ovarian cancer. *National Rev. Cancer* 9(3):167–181.
- Zhong Y, Hong LJ (2021) Knockout-tournament procedures for large-scale ranking and selection in parallel computing environments. *Oper. Res.* Forthcoming.