This article was downloaded by: [98.253.32.243] On: 04 February 2026, At: 14:54
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Ranking and Contextual Selection

Gregory Keslin, Barry L. Nelson, Bernardo Pagnoncelli, Matthew Plumlee, Hamed Rahimian

To cite this article:
Gregory Keslin, Barry L. Nelson, Bernardo Pagnoncelli, Matthew Plumlee, Hamed Rahimian (2025) Ranking and Contextual Selection. Operations Research 73(5):2695-2707. https://doi.org/10.1287/opre.2023.0378

Full terms and conditions of use: https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2024, INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 73, No. 5, September–October 2025, pp. 2695–2707
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

Ranking and Contextual Selection

Gregory Keslin,[a,*] Barry L. Nelson,[a] Bernardo Pagnoncelli,[b] Matthew Plumlee,[c] Hamed Rahimian[d]

[a] Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208; [b] SKEMA Business School, Université Côte d'Azur, 59777 Lille, France; [c] Amazon, Seattle, Washington 98109; [d] Department of Industrial Engineering, Clemson University, Clemson, South Carolina 29634

*Corresponding author

Contact: gregorykeslin2025@u.northwestern.edu, https://orcid.org/0009-0004-2615-251X (GK); nelsonb@northwestern.edu, https://orcid.org/0000-0002-1325-2624 (BLN); bernardo.pagnoncelli@skema.edu (BP); mplumlee@amazon.com, https://orcid.org/0000-0002-7820-3399 (MP); hrahimi@clemson.edu, https://orcid.org/0000-0003-1385-4120 (HR)

Received: July 11, 2023
Revised: April 17, 2024
Accepted: August 29, 2024
Published Online in Articles in Advance: October 3, 2024

Area of Review: Simulation

https://doi.org/10.1287/opre.2023.0378

Copyright: © 2024 INFORMS

Abstract. This paper proposes a new ranking-and-selection procedure, called ranking and contextual selection, in which covariates provide context for data-driven decisions. Our procedure optimizes over a set of covariate design points off-line and then, given an actual observation of the covariate, makes an online decision based on classification—a distinctly new approach. We prove the existence of an experimental design that yields a pointwise probability of good selection guarantee and derive a postexperiment assessment of our procedure that provides an optimality gap upper bound with guaranteed coverage for decisions with respect to future covariates. We illustrate ranking and contextual selection with an application to assortment optimization using data available from Yahoo!.

Funding: This work was supported by the National Science Foundation [Grant CMMI-2206973].
Supplemental Material: This article includes an online appendix and computer code and data supporting the study's findings at https://doi.org/10.1287/opre.2023.0378.

Keywords: simulation • statistical analysis • experiment design • nonparametric

1. Introduction

Stochastic simulation supports the design of systems that do not yet exist—and, therefore, have no historical record of performance—by generating synthetic outcomes from feasible alternatives. This characterization applies even to the improvement of existing systems for which a simulation can generate outcomes that represent potential, but as of yet not implemented, improvements. Often the goal is to choose the alternative or decision that optimizes some measure or measures of system performance. Traditionally, simulation optimization (SO) targets performance with respect to static (one-time) decisions, such as how many hospital beds to dedicate to mental health patients, what reorder points to set in a sports equipment supply chain, or what schedule to use for campus buses. In the static setting, ranking and selection (R&S) is a workhorse SO tool for choosing the best decision with some guarantee of optimality. R&S applies to problems that have a finite number of feasible decisions, and all of them can be simulated (at least a little bit); therefore, the focus is on efficient allocation of simulation effort and delivery of statistical guarantees.

In contrast to one-time, static decisions, data-driven decision making is a broad term that includes deferring selection of the optimal decision until contemporaneous information, often called the context or covariate, is available to obtain a refined, situation-specific decision. When there is variability in the realized outcome from the decision, this is the difference between optimizing a conditional expected value (context sensitive) rather than an unconditional one (averaging over the contexts).

An application that is pushing SO toward data-driven decisions is digital twin technology, in which a simulation model lives in parallel with the evolving real system it represents. For instance, when responding to airline flight disruptions because of covariates such as bad weather and equipment failure, a digital twin allows the decision maker to explore rerouting, delaying, and cancellation decisions via simulation to minimize their impact on cost for the current context (Rhodes-Leader et al. 2022). When an SO can be completed within the time window between observing the covariate and needing a decision, then real-time SO is the obvious strategy. This paper addresses the corresponding situation in which any simulation experiments must be executed in advance of the decision (off-line) for results to be available for timely use when needed (online), typically because of the computational expense of the SO relative to the decision window. The key, in either case, is the availability of a high-fidelity simulation model that generates performance data. This is different from the setting in which no simulation is available, and therefore, one may have to learn a good decision rule by taking real actions with real costs when trying to minimize regret

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

(e.g., contextual bandits; see, for instance, Bouneffouf et al. 2020). Our methodology requires a simulation.

The most significant advances in SO for data-driven, real-time decisions have come in R&S under the names contextual R&S and R&S with covariates, both of which we abbreviate as R&S+C. See Alban et al. (2021), Cakmak et al. (2021), Ding et al. (2022), Du et al. (2022), Gao et al. (2019), Li et al. (2018, 2020, 2024), Liu et al. (2022), Shen et al. (2021), Zhang et al. (2023), and references therein. The present paper provides a fundamental rethinking of R&S+C: current technology either fails to exploit the decades of advances in R&S procedures or makes substantial (and typically unverifiable) assumptions; instead, we directly exploit R&S technology and make weak assumptions, resulting in simple, scalable algorithms. In fact, our assumptions are so weak that we can employ any sensible R&S procedure although, in practice, there are significant benefits to exploiting efficient, effective ones. The key is treating R&S+C as a classification problem, which is our first major contribution.

To distinguish our approach from earlier efforts, we call it ranking and contextual selection (R&CS). We establish novel supporting theory for R&CS that emphasizes a postexperiment assessment in the form of an empirically generated upper bound on the classification quality of future decisions—a bound that holds with a prespecified confidence. This bound is distribution-free although it typically is advantageous to employ parametric R&S procedures to generate the off-line data. R&CS is also easy to implement by using off-the-shelf R&S procedures and a nearest neighbor classification method. The postexperiment assessment is distinctly different from any other R&S+C procedure and is our second major contribution.

In the R&S+C literature, the most studied application is precision or personalized medicine; see Alban et al. (2021), Du et al. (2022), and Shen et al. (2021). The application we employ is assortment optimization, which is an essential problem in revenue management. In e-commerce, advertisement, and recommender systems, customers are shown a set of items hoping that some of those displayed items are of interest. Whereas companies may choose a static assortment to maximize revenue or engagement averaged over the whole population of potential customers (Rusmevichientong et al. 2010, Sauré and Zeevi 2013, Agrawal et al. 2019), the tendency more recently is to personalize the items displayed to each customer or consumer segment (Lobel 2021, Chen et al. 2022). This so-called assortment personalization leverages the availability of covariates, such as purchasing history, location, age, and gender, among others (Chen et al. 2020, Kallus and Udell 2020, Miao and Chao 2022). Simulation may be needed when the customer choice model upon which the assortment optimization is based is complex, and this is the situation to which our work pertains; see Section 5.

We apply R&CS methodology to a news article recommendation system that selects which assortment to display on a web page given each customer's covariates based on the data set Yahoo! Front Page Today Module User Click Log Data Set, Version 2.0. Clearly this decision is needed almost instantly, rendering a real-time SO impossible. The results show that our upper bound on the achieved optimality gap between the click-through rates of the selected assortment and the best possible assortment is achieved with the desired 95% confidence in an experiment using a real-world data set.

The remainder of the paper is organized as follows. Section 2 provides basic notation and background to contrast R&CS with previous methods for R&S+C. In Section 3, we present a high-level description of the R&CS procedure. Sections 4 and 5 establish the underlying theory and empirical performance of R&CS, respectively. We close the paper with a broader discussion in Section 6. The idea of using classification to solve R&S+C problems is first presented in Keslin et al. (2022), which includes an empirical comparison with Shen et al. (2021) but none of the theory, postexperiment assessment methodology, or realistic example presented here.

2. Background

The scenario in which R&S+C applies is as follows: a high-fidelity simulation model is available, and the simulation experiment is executed off-line in advance of decision making. We assume that there are $p < \infty$ possible decisions; although it is more common in R&S to refer to the alternatives as "systems," we believe "decisions" is appropriate for R&S+C. In Section 5, each decision corresponds to a specific assortment to offer a potential customer. The mean performance of a decision depends on the value of a covariate, denoted generically by $\boldsymbol{X} \in \mathcal{X} \subset \Re^q$ for some finite $q \geq 1$. We treat the covariate as a random variable with a known or estimable distribution $\mathcal{F}$ over $\mathcal{X}$.

Let the mean performance of the jth decision, given $\boldsymbol{X} = \boldsymbol{x}$, be denoted by $\mu_j(\boldsymbol{x}), j = 1, 2, \ldots, p$. The ideal R&S+C procedure selects $j^*(\boldsymbol{x}) = \arg\max_{j=1,2,\ldots,p} \mu_j(\boldsymbol{x})$, but because the problem is stochastic, some compromise is accepted. By stochastic, we mean that, whereas the function $\mu_j(\boldsymbol{x})$ is not available under context $\boldsymbol{x}$, we can simulate independent and identically distributed (i.i.d.) replications $Y_{1j}(\boldsymbol{x}), Y_{2j}(\boldsymbol{x}), \ldots$ of the jth decision's performance with mean $\mu_j(\boldsymbol{x}) = \mathrm{E}(Y_{ij}(\boldsymbol{x}))$.

When the real-world covariate is realized, $\boldsymbol{X} = \boldsymbol{x}$, a decision $j(\boldsymbol{x})$ is chosen online based on the information available from the simulation experiment but without further simulation. R&S+C encompasses both the design of the off-line simulation experiment and the online procedure for selecting a decision; obviously, these two aspects need to work together.

The dominant paradigm for R&S+C is to employ an off-line experiment to estimate p metamodels,

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

$\hat{\mu}_j(\cdot)$, $j = 1, 2, \ldots, p$, and then to make the online decision by choosing $j(\boldsymbol{x}) = \arg\max_{j=1,2,\ldots,p} \hat{\mu}_j(\boldsymbol{x})$. See the references in Section 1 on R&S+C. In Keslin et al. (2022), we refer to this approach as "weight then optimize" because many metamodels can be viewed as weighted averages of the simulation data, and the optimization takes place after the covariate is realized. Gaussian process metamodels are typical, and the goal of the experiment design is to expend a finite budget of simulation effort in a way that produces metamodels yielding good online decisions without any particular guarantee on the achieved optimality gap. We believe a statement about the optimality gap is essential. By optimality gap, we mean

$$\Delta(\boldsymbol{x}) = \mu_{j^*(\boldsymbol{x})}(\boldsymbol{x}) - \mu_{j(\boldsymbol{x})}(\boldsymbol{x}), \tag{1}$$

that is, the shortfall of the mean of the selected decision below the mean of the best decision. When the values of the means in (1) are sufficiently close, that is, when $\Delta(\boldsymbol{x}) \le \delta$ for some acceptable $\delta > 0$, we say $j(\boldsymbol{x})$ is a good decision.

An exception to the lack of an optimality gap guarantee is Shen et al. (2021), which is why it is the only reference we review in detail. Taking as input $\delta > 0$, a confidence level $1 - \alpha$, and assuming that $\boldsymbol{X} \sim \mathcal{F}$ with $\mathcal{F}$ known, they create metamodels $\hat{\mu}_j(\cdot)$ that provide an expected probability of good selection (PGS) guarantee

$$\Pr\{\Delta(\boldsymbol{X}) \le \delta\} \ge 1 - \alpha. \tag{2}$$

PGS is one of the many types of guarantees found in the R&S literature (Eckman and Henderson 2018); Shen et al. (2021) provide this strong guarantee in expectation over the simulation experiment and covariate space, which we refer to as a predesign guarantee. To achieve this, they assume that $\mu_j(\boldsymbol{x}) = \boldsymbol{\beta}_j^\top \boldsymbol{x}$ with $\boldsymbol{\beta}_j$ unknown. That is, they assume that the relationship between the covariate and the true mean of the jth decision is linear.

For R&CS, we do not build metamodels of $\mu_j(\boldsymbol{x})$. Instead, for each covariate value in an experiment design $\mathcal{D}_m = \{\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_m\}$, we execute an off-the-shelf R&S procedure; this R&S procedure could be as simple as simulating all decisions equally, but there are advantages to employing an efficient one with correct selection or good selection guarantees. In any event, the R&S procedure yields a corresponding collection of selected decisions $\mathcal{R}_m = \{R(\boldsymbol{X}_1), R(\boldsymbol{X}_2), \ldots, R(\boldsymbol{X}_m)\}$, where $R(\boldsymbol{X}_i) \in \{1, 2, \ldots, p\}$. The design points $\mathcal{D}_m$ may be chosen systematically or randomly. We call $(\mathcal{D}_m, \mathcal{R}_m)$ the database. When a covariate $\boldsymbol{X}$ is realized, we classify it as one of the decisions in $\mathcal{R}_m$ based on the nearness of $\boldsymbol{X}$ to the covariates in $\mathcal{D}_m$. In Keslin et al. (2022), we refer to the classification approach as optimize then weight because we first optimize via R&S at each covariate in the design set and later select a decision for a new covariate by classification; classification can often be viewed as a weighted score with respect to the database.

The potential advantages of R&CS are obvious: (i) it exploits decades of advancements in R&S procedures, (ii) classification is a well-developed machine learning methodology, and (iii) implementation is simple. These benefits lead to computational and practical gains, but what can be proven about this approach?

In Keslin et al. (2022), we show that, under very weak conditions, a k-nearest neighbors (kNN) classifier is asymptotically consistent in the sense that the PGS converges to one pointwise when the design size $m \to \infty$ implies that the number of nearest neighbors $k \to \infty$ and $k/m \to 0$. In this paper, we demonstrate the following more useful results:

i. When the $\mu_j(\cdot)$'s are Lipschitz continuous, there exists a finite-m experimental design $\mathcal{D}_m$ that achieves a pointwise PGS guarantee.

ii. Given a finite-m design beyond a sufficient size (that we specify) and no assumptions on the $\mu_j(\cdot)$'s, the $1 - \alpha$ quantile of the achieved optimality gap of the design can be obtained. This is a postexperiment assessment of the R&CS procedure for classification of future covariates with a predesign confidence guarantee.

iii. A computationally efficient and asymptotically valid empirical method for approximating the quantile in item (ii) above is available under very weak assumptions.

iv. R&CS has the desired performance in a realistic problem.

In the next section, we present the R&CS procedure and set up the questions that we answer in theory and in practice in Sections 4 and 5.

3. R&CS Classification Procedure

As a prelude to the technical details in Section 4, we present our R&CS procedure here. There are possible variations, but we outline the one that we think is most likely to be applied in practice.

R&CS consists of an off-line setup procedure and an online classification procedure; see Figure 1. A database $(\mathcal{D}_m, \mathcal{R}_m)$ is constructed off-line. When presented with a new covariate value $\boldsymbol{X}$, the online classification procedure uses the database to return a decision $J_m(\boldsymbol{X}) \in \{1, 2, \ldots, p\}$. Associated with the database is an empirical assessment of its quality, specifically a quantity $\hat{\Delta}_m^{1-\alpha}$ such that

$$\Pr\{\mu_{j^*(\boldsymbol{X})}(\boldsymbol{X}) - \mu_{J_m(\boldsymbol{X})}(\boldsymbol{X}) \le \hat{\Delta}_m^{1-\alpha}\} \ge 1 - \alpha. \tag{3}$$

In other words, the achieved optimality gap is no larger than $\hat{\Delta}_m^{1-\alpha}$ with high confidence. This is a key innovation in our work: it is analogous to the expected PGS (EPGS) guarantee of Shen et al. (2021) but without any unverifiable assumptions about the output distributions, functional relationship between $\mu_j(\boldsymbol{x})$ and $\boldsymbol{x}$, or R&S procedure deployed to build the database. We do share the assumption of known or estimated covariate distribution $\mathcal{F}$.

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

The postexperiment off-line assessment leading to $\hat{\Delta}_m^{1-\alpha}$ is based on a leave-one-out estimate of the optimality gap for each covariate in the database $\mathcal{D}_m$. A confidence level $1-\alpha$ is required, and it may be different from any confidence level used by the R&S procedure if it has one.

As an additional input, the user provides an R&S procedure. All R&S procedures produce point estimators as a by-product of selection. For simplicity, we assume that these are sample means of i.i.d. replications although that is not essential. Let $\overline{Y}_j(\boldsymbol{x})$ denote the sample mean of the outputs generated for decision j at covariate value $\boldsymbol{x}$ by the R&S procedure. Notice that, for many procedures, $R(\boldsymbol{x}) = \arg\max_{j=1,2,\ldots,p} \overline{Y}_j(\boldsymbol{x})$. Because the true means are not known, we also use these point estimators as plug-ins for our optimality gap assessment; whenever plug-ins are used, we replace hats with bars in the notation as in $\overline{\Delta}_m^{1-\alpha}$ rather than $\hat{\Delta}_m^{1-\alpha}$.

The online classification procedure we employ is the nearest neighbor (NN); that is, we predict that the best decision for covariate $\boldsymbol{X}$ is the decision associated with the nearest covariate value in the database $\mathcal{D}_m$ to $\boldsymbol{X}$. Thus, a distance metric $d(\cdot,\cdot)$ in $\Re^q$ is also required.

Figure 2 illustrates R&CS when there are $p = 7$ possible decisions: the dots are $m = 9$ covariate values $\boldsymbol{X}_i$ that have been sampled off-line from the covariate space $\mathcal{X}$ to create the database along with the decisions $R(\boldsymbol{X}_i) \in \{1,2,3,4,5,6,7\}$ assigned to each of them by the R&S procedure. Notice that decision 6 was never selected. In the postexperiment assessment of the optimality gap, $\boldsymbol{X}_8$ is the NN of $\boldsymbol{X}_7$, and therefore, its decision "1" is assigned to $\boldsymbol{X}_7$ in the leave-one-out assessment; see the thick arrow. The same occurs for each design point. When a new covariate is observed online (circle), the NN classification using the database assigns the decision "5" to it (thin arrow).

Clearly, the R&CS procedure described in Figure 1 is easy to apply, requiring no special software beyond an R&S procedure. In Sections 4 and 5, we address questions raised by the procedure:

- The database is constructed by sampling values from the covariate space $\mathcal{X}$ according to $\mathcal{F}$. Why, and

Figure 1. R&CS Off-line Setup and Online Classification Procedures

R&CS Offline Set-Up Procedure

1. Database:

(a) Generate the design set $\mathcal{D}_m = \{\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_m\}$ by generating m i.i.d. values from $\mathcal{F}$ for some $m+1 \geq \lceil 2/\alpha \rceil$.

(b) Solve m independent R&S problems yielding the corresponding selections $\mathcal{R}_m = \{R(\boldsymbol{X}_1), R(\boldsymbol{X}_2), \ldots, R(\boldsymbol{X}_m)\}$.

(c) Return the database $(\mathcal{D}_m, \mathcal{R}_m)$.

2. Post-experiment Assessment:

(a) For $i \in \{1,2,\ldots,m\}$ do

i. Let $I_m^{-i}(\boldsymbol{X}_i) = \arg\min_{h=1,2,\ldots,m;h\neq i} d(\boldsymbol{X}_h, \boldsymbol{X}_i)$, the NN to $\boldsymbol{X}_i$ in the database.

ii. Let $J_m^{-i}(\boldsymbol{X}_i) = R(\boldsymbol{X}_{I_m^{-i}(\boldsymbol{X}_i)})$, the decision associated with the NN to $\boldsymbol{X}_i$.

iii. Let

$$\bar{\Delta}_m^{-i}(\boldsymbol{X}_i) = \bar{Y}_{R(\boldsymbol{X}_i)}(\boldsymbol{X}_i) - \bar{Y}_{J_m^{-i}(\boldsymbol{X}_i)}(\boldsymbol{X}_i)$$

the estimated gap between the optimal decision at $\boldsymbol{X}_i$ and the NN decision at $\boldsymbol{X}_i$.

(b) Return the assessment $\bar{\Delta}_m^{1-\alpha} = \bar{\Delta}_{m,(i^\star)}$, the $i^\star$ order statistic of the $\bar{\Delta}_m^{-i}(\boldsymbol{X}_i)$, where $i^\star$ is chosen as prescribed by Theorem 2.

R&CS Online Classification Procedure

1. Observe a new covariate value $\boldsymbol{X} \sim \mathcal{F}$.
2. Find the NN $I_m(\boldsymbol{X}) = \arg\min_{i=1,2,\ldots,m} d(\boldsymbol{X}_i, \boldsymbol{X})$.
3. Return the decision $J_m(\boldsymbol{X}) = R(\boldsymbol{X}_{I_m(\boldsymbol{X})})$.

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

Figure 2. Illustration of R&CS with an Off-line Database of $m = 9$ Design Points (Dots) and Their Associated R&S Decisions Along with an Online Covariate Value $\boldsymbol{X}$ (Circle) to Assign a Decision

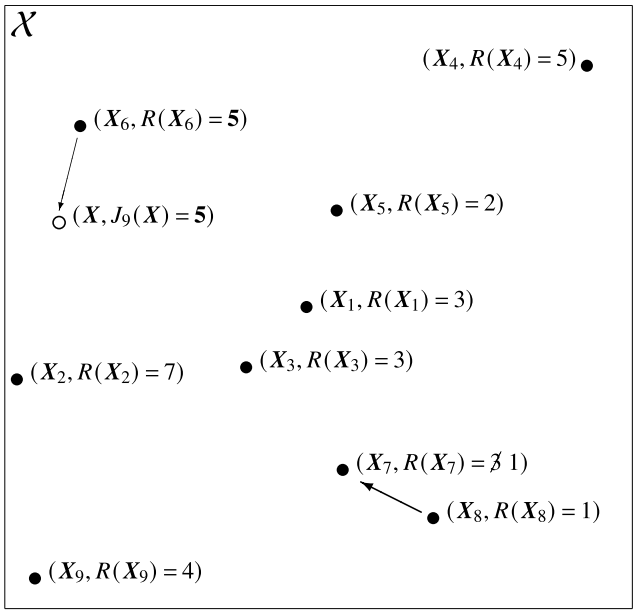

Note. The thick arrow is the decision assigned to design point $\boldsymbol{X}_7$ in the leave-one-out postexperiment assessment; the thin arrow is the decision assigned to online covariate $\boldsymbol{X}$ using the NN from the database.

are there potential advantages to selecting the design points systematically? Theorem 1 shows that very strong guarantees can be obtained by carefully selecting design points, but only under correspondingly strong assumptions and by employing a very large number of design points.

- A postexperiment assessment of the achieved optimality gap of the database is provided rather than the specific size-δ guaranteed optimality gap provided by the multistage procedure of Shen et al. (2021). However, our assessment holds quite generally. How is this proven? Theorems 2 and 3 provide finite-sample and asymptotic answers, respectively. Probabilistically exchangeable design points and the NN classifier are key.
- How tight is the optimality gap estimate, and how does performance compare with Shen et al. (2021), which provides an EPGS guarantee under much stronger assumptions? Section 5 contains empirical evidence bolstered by others in Keslin et al. (2022).

The reader will no doubt notice that we employ the very simple kNN classifier with $k = 1$, raising the questions of why not a more sophisticated classifier and why not $k > 1$?

Probabilistically exchangeable design points and the kNN classifier are key to proving Theorems 2 and 3, which deliver the postexperiment assessment. We strongly suspect that they cannot be proven for other classifiers, at least not with the same machinery. However, other classifiers could be used, and we expect they would have good performance but without a way to assess it.

The premise behind kNN is that design points closer to the new covariate $\boldsymbol{X}$ are more likely to be correct than ones farther away. In machine learning applications, kNN often employs $k > 1$ neighbors to dampen noise, which is particularly important in kNN regression. However, because we use R&S procedures with small, user-controlled probability of error α_R, there is very little noise at our design points; in fact, when they are misclassified (the best is not chosen by R&S) the class that is chosen tends to be close in expected value (i.e., the good selection guarantee). If we use $k > 1$, then design points farther away come into play, and the majority class has to be correct for $\boldsymbol{X}$ to be classified correctly. However, unlike typical machine learning applications, we will not have a densely filled design space; in fact, we show that $m = 39$ design points is all we need to get 95% optimality gap confidence bounds. Therefore, if we employ (say) $k = 7$ neighbors, then design points quite distant from $\boldsymbol{X}$ come into play, which is problematic. As noted, we are doing classification, not regression, so the majority vote of the neighbors is what matters. Extending Theorems 2 and 3 to $k > 1$ also forces the sufficient design size m to be k times larger. Therefore, we only present the NN versions of the theorems even though they can be adapted to more neighbors.

4. Supporting Theory

In this section, we focus on properties of the optimality gap produced by R&CS. We first establish the existence of an experimental design in the covariate space $\mathcal{X}$ for which a PGS guarantee even stronger than Shen et al. (2021) is obtained under similar assumptions on $\mu_j(\cdot)$ and a certain assumption on the R&S procedure employed to build the database. This establishes that a classification approach can achieve the same guarantee as a linear metamodel when a linear model is correct. We then show the more important result that, in the absence of these assumptions, the predesign coverage guarantee for the postexperiment estimated optimality gap quantile $\hat{\Delta}_m^{1-\alpha}$ still holds for R&CS—a guarantee that is entirely new for R&S+C.

Specifically, in Section 4.2, we prove that, if all of the mean response functions $\mu_j(\cdot)$ are Lipschitz continuous with the largest constant known and the R&S procedure provides a marginal PGS guarantee, then we can construct an experimental design for R&CS with a prespecified pointwise PGS guarantee. However, this design would typically demand many more design points than needed for acceptable performance.

Remarkably, our second set of results in Section 4.3 establishes a related and easily interpreted performance

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

guarantee with only a modest sized experiment design, assuming little. To understand this guarantee, notice that, for the classifier constructed from the database, $J_m(\cdot)$, there exists an unknown optimality gap for which it achieves an EPGS of $1-\alpha$. The quality assessment $\hat{\Delta}_m^{1-\alpha}$, is an estimate of this value. Section 4.3.1 proves its finite-sample coverage guarantees, whereas Section 4.3.2 establishes asymptotic coverage guarantees for $\overline{\Delta}_m^{1-\alpha}$, the plug-in version of $\hat{\Delta}_m^{1-\alpha}$. To deliver these results, no assumptions on the $\mu_j(\cdot)$, the distributions of the simulation outputs $Y_{\ell j}(\cdot)$, or the R&S procedure are required.

Before stating and proving results, we outline the spectrum of optimality gap guarantees that one might consider.

4.1. A Spectrum of Guarantees

In R&CS, we append a subscript m to quantities to indicate the dependence on the design $\mathcal{D}_m$ of size m. The inference in R&CS is the consequence of a process that can be broken into four sequential steps:

i. The (possibly random) experimental design, $\mathcal{D}_m$, is realized as $\mathcal{D}_m = d_m$.

ii. Given $\mathcal{D}_m = d_m$, an R&S procedure is executed conditioned on each $\boldsymbol{x}_i \in d_m$. The selected decision of each R&S procedure is realized as $R(\boldsymbol{x}_i) = r_i$. The vector $\rho_m = (r_1, r_2, \ldots, r_m)$ is paired with d_m to construct the realized database used for classification, $(\mathcal{D}_m, \mathcal{R}_m) = (d_m, \rho_m)$.

iii. An NN classifier, $j_m : \mathcal{X} \to \{1, 2, \ldots, p\}$, is constructed from the database (d_m, ρ_m). Specifically, let $d : \mathcal{X} \times \mathcal{X} \to \mathbb{R}^+$ be a metric and $I_m(\boldsymbol{x}) \in \operatorname{argmin}_{i=1,2,\ldots,m} d(\boldsymbol{x}_i, \boldsymbol{x})$, where ties are broken by choosing the covariate in d_m with the smallest index. Then, $j_m(\boldsymbol{x}) = r_{I_m(\boldsymbol{x})}$. In this paper, the metric is the Euclidean distance.

iv. The real-world covariate is realized as $\boldsymbol{X} = \boldsymbol{x}$, and the decision $j_m(\boldsymbol{x})$ is returned with (unknown) optimality gap $\hat{\Delta}_m(\boldsymbol{x})$.

We now identify three steps at which one could define PGS guarantees:

- Predesign: Prior to the realization of $\mathcal{D}_m$, the design $\mathcal{D}_m$, classifier $J_m(\cdot)$, and real-world covariate $\boldsymbol{X}$ are all random. The EPGS guarantee

$$\Pr\{\hat{\Delta}_m(\boldsymbol{X}) \le \delta\} \ge 1-\alpha$$

is a predesign guarantee because this probability is an expectation with respect to the random variables $\mathcal{D}_m$, $J_m(\cdot)$, and $\boldsymbol{X}$.

- Postdesign: Conditional on $\mathcal{D}_m = d_m$ but prior to the realization of $\mathcal{R}_m$ (and, hence, J_m), a similar postdesign EPGS can be defined as

$$\Pr\{\hat{\Delta}_m(\boldsymbol{X}) \le \delta | \mathcal{D}_m = d_m\}.$$

In the postdesign EPGS, the expectation is with respect to the random variables $J_m(\cdot)$ and $\boldsymbol{X}$, whereas $\mathcal{D}_m$ is fixed as d_m.

- Postsimulation: Conditional on $J_m = j_m$ and $\mathcal{D}_m = d_m$, a postsimulation EPGS can be defined as

$$\Pr\{\hat{\Delta}_m(\boldsymbol{X}) \le \delta | J_m = j_m, \mathcal{D}_m = d_m\}.$$

In the postsimulation EPGS, the expectation is with respect to the real-world covariate, $\boldsymbol{X}$, because the classifier is fixed as j_m.

In Theorem 1, we demonstrate the existence of an experimental design that implies a postdesign EPGS guarantee, whereas Theorems 2 and 3 establish predesign guarantees analogous to Shen et al. (2021).

4.2. Finite Design with PGS Guarantee

Theorem 1 shows that, for a confidence level $1-\alpha \in (1/p, 1)$ and optimality gap $\delta > 0$, there exists a design d_m that delivers the postdesign pointwise PGS guarantee; clearly, a pointwise guarantee implies an EPGS guarantee. The intuition behind Theorem 1 is that, with Lipschitz bounded changes in the mean functions, one can execute PGS-guaranteeing R&S procedures employing a smaller-than-target optimality gap δ_R on a dense enough design set to control the selection error for any new covariate and its nearest design-point neighbor.

We first establish sufficient conditions for d_m to attain the postdesign pointwise PGS guarantee in Lemma 1 and then exhibit such a design in Theorem 1. We assume the following.

Assumption 1 (Classifier). *The classifier $J_n(\cdot)$ is NN with associated distance metric $d(\cdot,\cdot)$. The realized classifier, $j_m(\cdot)$, breaks ties by choosing the design covariate with the smallest index in d_m.*

Assumption 2 (Lipschitz Continuity). *There exists $\lambda > 0$ such that, for every decision j, and any $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{X}$, we have $|\mu_j(\boldsymbol{x}) - \mu_j(\boldsymbol{x}')| \le \lambda\, d(\boldsymbol{x}, \boldsymbol{x}')$.*

Assumption 3 (Good Selection Guarantee). *For all $\boldsymbol{x} \in \mathcal{X}$, any fixed $\delta_R < \delta$, and any fixed $1-\alpha_R \ge 1-\alpha$, the R&S procedure satisfies $\Pr\{\mu_{j^*(\boldsymbol{x})}(\boldsymbol{x}) - \mu_{R(\boldsymbol{x})}(\boldsymbol{x}) \le \delta_R\} \ge 1-\alpha_R$.*

Assumption 4 (Design Density). *For the realized design, d_m, and all $\boldsymbol{x} \in \mathcal{X}$, there exists $\boldsymbol{x}_i \in d_m$ such that $d(\boldsymbol{x}_i, \boldsymbol{x}) \le (\delta - \delta_R)/(2\lambda)$.*

Lemma 1. *Under Assumptions 1–4, $\Pr\{\hat{\Delta}_m(\boldsymbol{x}) \le \delta | \mathcal{D}_m = d_m\} \ge 1-\alpha$ for all $\boldsymbol{x} \in \mathcal{X}$.*

The proof is found in the Electronic Companion EC.1. We now demonstrate the existence of a d_m that satisfies Assumption 4.

Theorem 1. *Under Assumptions 1–3, if $\mathcal{X}$ is a compact subset of $\Re^q$ for the topology induced by $d(\cdot,\cdot)$, then there exists a positive integer, m, and a design $d_m \in \mathcal{X}^m$ such that, for all $\boldsymbol{x} \in \mathcal{X}$, $\Pr\{\hat{\Delta}_m(\boldsymbol{x}) \le \delta | \mathcal{D}_m = d_m\} \ge 1-\alpha$.*

Proof. Let $\delta_{\text{diff}} = \delta - \delta_R$. For every $\boldsymbol{x} \in \mathcal{X}$, denote the ball centered at $\boldsymbol{x}$ with radius $\delta_{\text{diff}}/(2\lambda)$ by $B_{\boldsymbol{x}} = B(\boldsymbol{x}, \delta_{\text{diff}}/$

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

(2λ)). From the compactness of $\mathcal{X}$, there exists a finite enumeration of these balls, $B_{x_1}, B_{x_2}, \ldots, B_{x_m}$, such that $\mathcal{X} \subseteq \cup_{i=1}^{m} B_{x_i}$. Let d_m be the centers of these balls: $d_m = \{x_1, x_2, \ldots, x_m\}$.

Fix an arbitrary $x \in \mathcal{X}$. Let $\tilde{x} \in \arg\min_{x_i \in d_m} d(x_i, x)$ be the closest design covariate in d_m to x with ties broken by selecting the design point with the smallest index. Because $\mathcal{X} \subseteq \cup_{i=1}^{m} B_{x_i}$ and the radius of every B_{x_i} is identical, $x \in B_{\tilde{x}}$. Therefore, $d(\tilde{x}, x) \leq \delta_{\text{diff}}/(2\lambda)$. Thus, d_m satisfies Assumption 4, and therefore, $\Pr\{\hat{\Delta}_m(x) \leq \delta | \mathcal{D}_m = d_m\} \geq 1 - \alpha$ by Lemma 1. □

Of course, finding such a design d_m in practice requires knowledge of an upper bound on the Lipschitz constant, which is similar to the requirement that each $\mu_j(\cdot)$ be a linear function for Shen et al. (2021). This theorem establishes a connection between the classification and metamodeling approaches, but the novel results in the next section are more practically applicable.

4.3. Postdesign Assessment

Rather than imposing strong conditions on $\mu_j(\cdot)$, R&CS supports assessment of the performance of the postsimulation classifier, $j_m(\cdot)$, by providing a predesign $1-\alpha$ confidence guarantee on the optimality gap $\hat{\Delta}_m(\boldsymbol{X})$. Specifically, we estimate the optimality gap quantile $\Delta_m^{1-\alpha} | \{d_m, j_m\}$, defined as $\inf\{c \in \mathbb{R} : \Pr\{\hat{\Delta}_m(\boldsymbol{X}) \leq c | d_m, j_m\} \geq 1 - \alpha\}$, where we let $\{d_m, j_m\}$ denote the conditioning event $\{\mathcal{D}_m = d_m, J_m = j_m\}$ for brevity. Notice that, if some specific optimality gap δ is greater than $\Delta_m^{1-\alpha} | \{d_m, j_m\}$, then the postsimulation EPGS is larger than $1-\alpha$. Thus, by estimating $\Delta_m^{1-\alpha} | \{d_m, j_m\}$, we can infer whether $j_m(\cdot)$ achieves any given optimality gap δ, and in any event, we can rely on the estimated gap with high confidence. What remains to show is how to realize $\Delta_m^{1-\alpha}$, which is a predesign random variable. By definition, we have that $\Pr\{\hat{\Delta}_m(\boldsymbol{X}) \leq \Delta_m^{1-\alpha}\} \geq 1 - \alpha$.

The guarantee borrows from an extensive statistical literature that leverages the property of exchangeability of random variables to produce predictive inference with minimal assumptions. Procedures of this type are often referred to as *conformal inference* (Vovk et al. 2005). In the next section, we establish the predesign coverage guarantee for $\hat{\Delta}_m^{1-\alpha}$ under the simplifying assumption that the NN classifier's leave-i-out optimality gap $\hat{\Delta}_m^{-i}(\boldsymbol{X}_i) = \mu_{j^*(\boldsymbol{X}_i)}(\boldsymbol{X}_i) - \mu_{J_m^{-i}(\boldsymbol{X}_i)}(\boldsymbol{X}_i)$ can be evaluated without error for any $\boldsymbol{X}_i$ in the database. This proof helps illuminate the theory and illustrate the most important aspects of our approach. Section 4.3.2 addresses the practical case of unknown but estimated $\hat{\Delta}_m^{-i}(\boldsymbol{X}_i)$ to obtain $\overline{\Delta}_m^{1-\alpha}$ as shown in the R&CS algorithm in Section 3.

4.3.1. Finite-Sample Coverage.

The fundamental concept underlying our approach is the exchangeability property of the randomly sampled off-line covariates in our database and a random future covariate. To enhance understanding, envision a simplified scenario in which we are provided with a classifier $\tilde{J}(\cdot)$ that is entirely independent of our database of covariates $\mathcal{D}_m = \{\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_m\}$ and implies an optimality gap function denoted by $\tilde{\Delta}(x) = \mu_{j^*(x)}(x) - \mu_{\tilde{J}(x)}(x)$.

Let $\boldsymbol{X}$ be a yet-to-be-observed covariate that is sampled independently of $\mathcal{D}_m$ from $\mathcal{F}$. Therefore, $\{\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_m, \boldsymbol{X}\}$ is a vector of exchangeable random variables, and so is $\{\tilde{\Delta}(\boldsymbol{X}_1), \tilde{\Delta}(\boldsymbol{X}_2), \ldots, \tilde{\Delta}(\boldsymbol{X}_m), \tilde{\Delta}(\boldsymbol{X})\}$. Let $\tilde{\Delta}_{(i)}$ be the ith order statistic of $\tilde{\Delta}(\boldsymbol{X}_1), \tilde{\Delta}(\boldsymbol{X}_2), \ldots, \tilde{\Delta}(\boldsymbol{X}_m)$. Clearly, $\Pr\{\tilde{\Delta}(\boldsymbol{X}) \leq \tilde{\Delta}_{(i)}\} \geq i/(m+1)$. Thus, given a classifier and an observed collection of optimality gaps, we can select the smallest order statistic larger than $(1-\alpha)(m+1)$ to produce a valid quality assessment of $\tilde{\Delta}(\boldsymbol{X})$ with minimal assumptions.

The technical difficulty for us is that R&CS uses $\mathcal{D}_m$ to both fit the classifier $J_m(\cdot)$ and provide the quality assessment. This places us outside the simplified scenario above as $\hat{\Delta}_{m+1}(\boldsymbol{X}_1), \hat{\Delta}_{m+1}(\boldsymbol{X}_2), \ldots, \hat{\Delta}_{m+1}(\boldsymbol{X}_{m+1})$ and $\hat{\Delta}_{m+1}(\boldsymbol{X})$ are no longer exchangeable because adding $\boldsymbol{X}$ to the database changes the NN classifier. We resolve this problem by using leave-one-out optimality gaps and treating $\boldsymbol{X}$ as the unobserved $(m+1)$st covariate from a design of size $m+1$. In brief, the tedious proof in Electronic Companion EC.2 shows that, by using a larger order statistic paired with a sufficient design size m, we can account for the possibility of a consequential change in the classifier from an $(m+1)$st design point.

We need two assumptions.

Assumption 5 (Design Independence). *Each $\boldsymbol{X}_i \in \mathcal{D}_{m+1}$ is i.i.d. with $\boldsymbol{X}_i \sim \mathcal{F}$.*

Assumption 6 (Classifier). *$J_m(\cdot)$ is an NN classifier with associated distance metric $d(\cdot,\cdot)$ that breaks ties by choosing the design covariate with the smallest index in $\mathcal{D}_m$.*

Theorem 2. *Suppose $m+1 \geq \lceil 2/\alpha \rceil$. Let $\xi = (\sqrt{\alpha(m+1)} - 1)^2$. Set,*

$$i^\star = \begin{cases} m, & \text{if } m+1 < 4/\alpha \\ \lceil m+1-\xi \rceil, & \text{otherwise.} \end{cases}$$

Then, under Assumptions 5 and 6, $\Pr\{\hat{\Delta}_m(\boldsymbol{X}) \leq \hat{\Delta}_{m,(i^\star)}\} \geq 1-\alpha$, where $\hat{\Delta}_{m,(i^\star)}$ is the $i^\star$ order statistic of $\hat{\Delta}_m^{-i}(\boldsymbol{X}_i), i = 1, 2, \ldots, m$.

The proof can be found in Electronic Companion EC.2.

For the $i^\star$ given by Theorem 2, this result shows that $\hat{\Delta}_{m,(i^\star)}$ delivers a predesign coverage guarantee of $1-\alpha$ for any covariate distribution and mean functions, $\mu_j(\cdot)$. Furthermore, no assumptions are made on the R&S procedure used to classify each design covariate. In fact, any method to classify each design covariate could be employed, including ones that do not satisfy a $1-\alpha_R$ PGS guarantee at each design covariate as in Section 4.2.

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

By way of intuition, notice that a naive method to produce an order statistic to satisfy a $1-\alpha$ predesign coverage guarantee of $\Delta_m^{1-\alpha}$ is to use the $i = \lceil m(1-\alpha) \rceil$ order statistic. For the $i^\star$ recommended by Theorem 2, we can show that, as the design size m increases, $i^\star/m$ converges to $1-\alpha$ as one would hope.

However, $\hat{\Delta}_{m,(i)}$ relies on knowing the optimality gap of the leave-one-out NN classifier at each design covariate, which we will not know in practice. Section 4.3.2 resolves this final issue.

4.3.2. Asymptotic Coverage. The estimator of the optimality gap quantile provided in Section 4.3.1 relies on knowledge of the leave-one-out value of the optimality gap at each design covariate $\boldsymbol{X}_i$, denoted by $\hat{\Delta}_{m_i}^{-i}(\boldsymbol{X}_i)$. Here, we use a plug-in estimate, denoted by $\overline{\Delta}_m^{-i}(\boldsymbol{X}_i)$, instead. Recall that the R&CS procedure in Figure 1 employs plug-ins.

Our experience is that reusing the simulation output data employed to construct the database $(\mathcal{D}_m, \mathcal{R}_m)$ to produce plug-in estimators works well. Indeed, the empirical results of Section 5 and controlled studies not reported here demonstrate this. In this section, we establish asymptotic $1-\alpha$ predesign coverage of $\overline{\Delta}_m^{1-\alpha}$ as the minimal sample size used in the R&S procedures increases.

Let the simulation output data from which the database is constructed be denoted by

$$\mathcal{O}_m = ((\boldsymbol{X}_i, (Y_{\ell j}(\boldsymbol{X}_i))_{\ell=1}^{N(\boldsymbol{X}_i,j)})_{i=1}^m)_{j=1}^p,$$

where $N(\boldsymbol{X}_i, j)$ is the number of replications obtained for decision j conditional on $\boldsymbol{X}_i$. From $\mathcal{O}_m$ the plug-in estimates $\overline{\Delta}_m^{-i}(\boldsymbol{x}_i)$ are derived. Recall the definition of the leave-one-out classifier function $J_m^{-i}(\boldsymbol{X}_i)$ employing the nearest neighbor classification of $\boldsymbol{X}_i$ for design $\mathcal{D}_m \setminus \boldsymbol{X}_i$ and classification vector $\mathcal{R}_m \setminus R(\boldsymbol{X}_i)$. For each $\boldsymbol{X}_i \in \mathcal{D}_m$, suppose that we estimate $\hat{\Delta}_m^{-i}(\boldsymbol{X}_i)$ via the sample means from $\mathcal{O}_m$ as described below.

Let $\overline{Y}_j(\boldsymbol{X}_i) = \sum_{\ell=1}^{N(\boldsymbol{X}_i,j)} Y_{\ell j}(\boldsymbol{X}_i)/N(\boldsymbol{X}_i, j)$, the sample mean of the outputs in $\mathcal{O}_m$ from decision j conditional on design covariate $\boldsymbol{X}_i$. Define

$$\overline{\Delta}_m^{-i}(\boldsymbol{X}_i) = \max_{j \in \{1,2,\ldots,p\}} \overline{Y}_j(\boldsymbol{X}_i) - \overline{Y}_{J_m^{-i}(\boldsymbol{X}_i)}(\boldsymbol{X}_i).$$

Notice that this definition is free of any particular R&S procedure and simply assumes we employ the largest sample mean as the point estimator of the optimal value at design point $\boldsymbol{X}_i$. The vector of the leave-one-out estimated optimality gaps is $\{\overline{\Delta}_m^{-1}(\boldsymbol{X}_1), \overline{\Delta}_m^{-2}(\boldsymbol{X}_2), \ldots, \overline{\Delta}_m^{-m}(\boldsymbol{X}_m)\}$; let $\overline{\Delta}_{m,(i)}$ denote the ith order statistic of this set and let $\overline{\Delta}_m^{1-\alpha} = \overline{\Delta}_{m,(i^\star)}$ for $i^\star$ chosen to approximately satisfy the $1-\alpha$ predesign coverage guarantee.

Because it does not incorporate the estimation error of $\overline{\Delta}_m^{-i}(\boldsymbol{X}_i)$, the $i^\star$ recommended by Theorem 2 is an aggressive choice if the sample size used to estimate $\hat{\Delta}_m^{-i}(\boldsymbol{X}_i)$ is small. However, as the sample size increases, the estimation error of $\overline{\Delta}_m^{-i}(\boldsymbol{X}_i)$ decreases, and the $i^\star$ from Theorem 2 is satisfactory. This intuition is supported by Theorem 3.

Theorem 3 demonstrates that $\overline{\Delta}_m^{1-\alpha}$ satisfies the predesign coverage guarantee as the minimum number of replications generated for each decision increases to infinity; notice that the number of design points does not increase. Depending on the R&S procedure, there are multiple mechanisms to drive the minimum number of replications to infinity. In the case of fixed-precision procedures such as the KN (Kim and Nelson 2001) and the two-stage NM procedure (Nelson and Matejcik 1995), as $\delta_R \to 0$ or $\alpha_R \to 0$, every decision is sampled infinitely often. In the case of many fixed-budget procedures, every decision is sampled infinitely often as the budget increases to infinity; see, for instance, Ryzhov (2016). Theorem 3 simply assumes that this happens, and selection is based on the largest sample mean.

Theorem 3. *Let* $N^\star = \min_{j=1,2,\ldots,p, i=1,2,\ldots,m} N(\boldsymbol{X}_i, j)$. *Let* $i^\star$ *be defined as in Theorem 2. Under Assumptions 5 and 6, and provided* $m+1 \geq \lceil 2/\alpha \rceil$, *then for any* $\epsilon > 0$,

$$\liminf_{N^\star \to \infty} \Pr\{\hat{\Delta}_m(\boldsymbol{X}) \leq \overline{\Delta}_m^{1-\alpha} + \epsilon\} \geq 1-\alpha.$$

The proof can be found in Electronic Companion EC.3.

As a practical matter, what is the computational cost of the plug-in procedure? First note that a sufficient design size m (number of R&S problems to solve) does not depend on the number of decisions p: $m \geq \lceil 2/\alpha \rceil - 1$, which for 95% confidence is 39. As long as this condition is satisfied, the user obtains inference on the quality of the decisions via Theorem 2.

To actually apply R&CS, we have to solve m R&S problems, each with p possible decisions. The efficiency of R&S procedures is well-studied, but a very pessimistic number of replications needed for decision j at design point $\boldsymbol{x}_i$ is $n_j(\boldsymbol{x}_i) \approx (h\sigma(\boldsymbol{x}_i)/\delta_R)^2$, where δ_R is the desired indifference-zone/good-selection bound, $\sigma^2(\boldsymbol{x}_i)$ is the output variance, and h is a constant that grows slowly with the number of decisions p. If, for simplicity, we assume all variances are the same, then the computational burden is $m \times p \times n$. Notice that users typically set δ_R so that differences in performance of $< \delta_R$ are meaningless, implying that mean performance is well-estimated. However, in controlled computational studies not reported here, we have found that setting δ_R too large for the problem at hand does not impact the optimality gap coverage; it just tends to make the optimality gap quantile larger to compensate for noisier estimates.

This analysis, as in Theorem 3, reaches its conclusion by forcing estimates of $\mu_j(\boldsymbol{x}_i)$ for all decisions j and design points $\boldsymbol{x}_i$ to be close to the truth, which is more demanding than required for good performance. For instance, the modified KN procedure we employ for the assortment-optimization problem in Section 5 ensures

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

that, at each design point, the mean performance of the decision chosen by the R&S procedure and the mean performance of the decision chosen by its nearest neighbor are precisely estimated—not the means of all decisions, which is a substantial reduction.

5. Empirical Evaluation and Illustration

In Keslin et al. (2022), we reported extensive experimental results that compare a version of R&CS to Shen et al. (2021) on the test suite of problems that they employ; their test suite contains synthetic problems for which the true relationship between the response and covariate is linear, and they vary factors such as the dimension of the covariate and response variability. We also applied both procedures to a supply chain reorder point problem. The results show that R&CS performs as well or better in terms of required number of replications and delivered the desired EPGS relative to a fixed δ even though R&CS does not promise to control the latter.

In this section, we present results for a realistic personalized assortment problem in which a firm offers a set of products to each arriving customer tailored to their covariate information such as age, gender, purchasing behavior, etc.

5.1. Description of the Assortment Problem

Consider a firm that sells a collection of N products, denoted by $\mathcal{N} = \{1, 2, \ldots, N\}$; we let zero denote the no-purchase choice. Because of display capacity considerations, the firm can offer $K \leq N$ products, called an assortment, to an arriving customer whose observed covariate vector is $\boldsymbol{X}$. The customer has a utility associated with each product, so the customer either purchases one of the products in the assortment or leaves with no purchase. The firm needs to determine the assortment of products to offer to each customer to maximize the probability of choosing a product among the ones available, which is sometimes called the click-through rate.

To formulate the problem, let $\mathcal{S} = \{\mathcal{A} \subset \mathcal{N} : |\mathcal{A}| = K\}$ denote the collection of all assortments $\mathcal{A}$ with cardinality K. Moreover, let $U_{\mathcal{A}}(i|\boldsymbol{X})$ denote the utility of product $i \in \mathcal{A} \cup \{0\}$ for a customer with covariate vector $\boldsymbol{X}$, provided they are offered assortment $\mathcal{A}$. Given assortment $\mathcal{A}$, the probability of choosing a product for a customer with covariate vector $\boldsymbol{X}$ is

$$\mu_{\mathcal{A}}(\boldsymbol{X}) = \Pr\left\{\max_{i\in\mathcal{A}} U_{\mathcal{A}}(i|\boldsymbol{X}) > U_{\mathcal{A}}(0|\boldsymbol{X})\right\}.$$

The personalized assortment optimization problem for a customer with covariate vector $\boldsymbol{X}$ is

$$\mathcal{A}^*(\boldsymbol{X}) \in \arg\max_{\mathcal{A}\in\mathcal{S}} \mu_{\mathcal{A}}(\boldsymbol{X}). \tag{4}$$

We apply R&CS to (4). Notice that contextual multi-armed bandits are often used for solving such problems when there is no simulation model to exploit, so a good assortment rule must be learned, ideally by minimizing the regret incurred in making online assortment decisions for actual arriving customers (Chen et al. 2020, Kallus and Udell 2020, Miao and Chao 2022). We describe the simulation model next; recall that a simulation is essential to R&CS.

5.2. Description of Data Set and Data Preparation

We simulate a news article recommendation system inspired by the Yahoo! Front Page Today Module User Click Log Data Set, Version 2.0. We assume that, when a user with covariate vector $\boldsymbol{X}$ visits the front page of Yahoo!'s website, an assortment $\mathcal{A}$ of articles is displayed, and either the customer clicks one of the offered articles or leaves (no-click option).

The data set contains users' covariate information and their clicking reactions to the news articles. Each data point in the data set is composed of (i) a time stamp, (ii) the user's covariate encoded as 136 binary variables, (iii) a single displayed article out of a set of 36 possible articles that can be offered, and (iv) the user's choice encoded as a binary variable (zero indicating no-click and one indicating click). We treat this data set as a recommendation system with the display capacity $K = 1$. In other words, we have access to $\{(\boldsymbol{x}_k, \hat{\mathcal{A}}_k, \hat{i}_k)\}$, where $\boldsymbol{x}_k$ is the covariate vector of customer k, $\hat{\mathcal{A}}_k$ is the offered assortment to customer k (with cardinality one), and $\hat{i}_k$ is the kth customer's choice.

This data set consists of millions of data points that represent user interactions with Yahoo!'s platform in an interval of 10 days in the month of May 2009. Because we are not investigating second-by-second differences in user behavior, we selected the first 50,000 data points for our experiments, which were sufficient to accurately estimate the problem's parameters. To construct the distribution of the covariate vector, we assumed independence among covariates and calculated the empirical probabilities of each covariate. To reduce the dimension of the covariate vector, we performed a principal components analysis (PCA) on the 136 binary covariate variables, and the q most informative covariates were chosen. To construct an observed covariate vector $\boldsymbol{X} = \boldsymbol{x}$ with dimension q, we generated a 136-dimensional vector of zeros and ones, independently from the empirical distributions, and trimmed this vector with the q PCA covariates.

Moreover, to model the utility functions, we did the following: for all $\mathcal{A} \in \mathcal{S}$ and $i \in \mathcal{A} \cup \{0\}$, we assumed

$$U_{\mathcal{A}}(i|\boldsymbol{X}) = \beta_i^0 + \beta_i^\top \boldsymbol{X} + \varepsilon_i,$$

where $\varepsilon_i \sim \text{Gumbel}(0,1)$. Therefore, given covariate vector $\boldsymbol{X}$, the best assortment to offer for Problem (4) is

$$\mathcal{A}^*(\boldsymbol{X}) \in \arg\max_{\mathcal{A}\in\mathcal{S}} \frac{\sum_{i\in\mathcal{A}} e^{(\beta_i^0-\beta_0^0)+(\beta_i-\beta_0)^\top \boldsymbol{X}}}{1+\sum_{i\in\mathcal{A}} e^{(\beta_i^0-\beta_0^0)+(\beta_i-\beta_0)^\top \boldsymbol{X}}},$$

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

and the corresponding probability of choosing a product in the assortment is

$$\mu_{\mathcal{A}^*(\boldsymbol{X})}(\boldsymbol{X}) = \frac{\sum_{i\in\mathcal{A}^*(\boldsymbol{X})} e^{(\beta_i^0-\beta_0^0)+(\beta_i-\beta_0)^\top \boldsymbol{X}}}{1+\sum_{i\in\mathcal{A}^*(\boldsymbol{X})} e^{(\beta_i^0-\beta_0^0)+(\beta_i-\beta_0)^\top \boldsymbol{X}}}.$$

We note that $(\beta_i^0-\beta_0^0)+(\beta_i-\beta_0)\boldsymbol{X}$ is the mean difference in utility of choosing to click article i over the no-click choice, for $i\in\mathcal{N}$, which induces the probabilities. To estimate the model, we first filtered the 50,000 data points to include only those data points for which articles in $\mathcal{N}$ were offered to customers. Then, we applied logistic regression to these data points. In this way, we created a realistic assortment simulation, but one for which we know the true click-through rates so that we can evaluate R&CS. In a real problem, the only difference would be a more complex and, therefore, less tractable, utility model that required simulation.

5.3. Experiment Design and Evaluation

We present results for R&CS and for TS$^+$ of Shen et al. (2021), which is the version of their procedure that accommodates the unequal output variances that are present in this problem.

Both TS$^+$ and R&CS provide a decision when presented with a new covariate value $\boldsymbol{X}$. The primary promise of TS$^+$ is to attain an EPGS of at least $1-\alpha$ when averaged over the covariate space, for which a good selection has true mean within a user-specified optimality gap δ. The strong assumption of response linearity in the covariates supports this guarantee, so attained EPGS is the key measure of statistical validity. Notice that the relationship is not linear in our personalized assortment problem. The simulation effort required to attain the guarantee (or not) is a measure of computational efficiency.

The primary promise of R&CS is to provide an estimate of the $1-\alpha$ quantile of the attained optimality gap of the database; therefore, the actual coverage probability is the key measure of statistical validity. The simulation effort required to build the database is a measure of computational efficiency. Although R&CS makes no EPGS promise with respect to a prespecified δ, we nevertheless evaluate its attained EPGS relative to the optimality gap δ that drives TS$^+$.

R&CS requires an R&S procedure, and we adopt KN of Kim and Nelson (2001). KN is a fully sequential indifference zone procedure that eliminates systems from further simulation until only one remains, and it exploits common random numbers. KN has a probability of correct selection (PCS) guarantee, so we set the PCS and indifference zone parameter for KN to the same $1-\alpha$ and δ, respectively, as in TS$^+$. Recall that the coverage of our optimality gap quantile estimator does not depend on the parameters of the R&S procedure. We made a small modification to KN that is described in Electronic Companion EC.4; its purpose is to ensure that KN leaves behind good plug-in optimality gap estimates. Specifically, the modification guarantees that the difference between the means of the decision at a covariate design point and the decision recommended by its leave-one-out NN is estimated to within $\pm\delta$ with $1-\alpha$ confidence. The NN is determined by Euclidean distance.

The TS$^+$ experiments do not employ common random numbers as the validity of the procedure is unproven and the impact of common random numbers on it appears to be negative. Stated differently, we design the TS$^+$ experiment so that the assumptions underpinning it are as closely satisfied as possible.

The specific version of the assortment optimization problem we consider has $N=11$ articles from which we test assortments of $K=9$ or 2; thus, in each setting, there are $p=55$ possible assortment decisions. When $K=9$, the assortments are very similar, making the best one hard to distinguish. We test covariates with dimensions $q=2,4,8$; a larger dimension implies more information about the customer. The minimal covariate design size for R&CS is $m=\lceil 2/\alpha\rceil-1=39$; we also test $m=79$ to evaluate the impact of approximately doubling the design size. Thus, in total, there are 12 cases. The same designs are employed for TS$^+$. In all cases, $1-\alpha=0.95$, which corresponds to the nominal EPGS for TS$^+$, the nominal PCS for KN, and the desired optimality gap coverage for R&CS. The indifference zone parameter δ for TS$^+$ and KN is tailored for each (K,q) combination to be the 90th percentile of the difference in click-through rates between the best and second best assortment across the covariate space, which was determined via extensive simulation; these values are listed in Electronic Companion EC.5. By setting δ in this way, we ensure that the optimality gap is equally difficult to attain in each experiment. Of course, in practice, this value would be driven by economic considerations.

The results reported below are based on 500 macroreplications. Within each macroreplication, an additional 5,000 simulations are used to estimate quantities such as the true coverage probability and EPGS. In summary, we find that, across all cases considered, R&CS attains at least its desired optimality gap coverage—its promised guarantee—whereas TS$^+$ does not always attain the target EPGS—its promised guarantee. In most cases, R&CS requires less—often dramatically less—simulation effort. The EPGS attained by R&CS is often close to but sometimes below the target set for TS$^+$ even though R&CS makes no EPGS promise for the given δ. Larger covariate designs improve the EPGS of R&CS as one would hope.

5.4. Empirical Results Summary

The optimality gap coverage results for R&CS are reported in Table 1, the computational effort comparisons are reported in Table 2, and the EPGS comparisons are reported in Table 3.

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

Table 1. Estimated Optimality Gap Coverage for R&CS

Case	q	K	m	Coverage	Standard error
1	2	9	39	0.983	0.001
2	4	9	39	0.978	0.001
3	8	9	39	0.977	0.001
4	2	2	39	0.985	0.001
5	4	2	39	0.978	0.001
6	8	2	39	0.975	0.001
7	2	9	79	0.985	0.001
8	4	9	79	0.978	0.001
9	8	9	79	0.977	0.001
10	2	2	79	0.986	0.001
11	4	2	79	0.978	0.001
12	8	2	79	0.974	0.001

Table 1 reveals that the optimality gap coverage is slightly conservative with all cases above the nominal 0.95. This is not surprising given the inequalities used to establish Theorem 2. More importantly, they show that there is no loss of coverage from using plug-in estimators, an approach supported by Theorem 3. The estimated and actual optimality gap quantiles are provided in Online Section EC.5.

TS$^+$ is a two-stage procedure that prescribes sufficient replications for each covariate value and decision pair in the experiment design to deliver a prespecified EPGS relative to a given good selection bound δ. Table 2 compares the number of replications consumed by R&CS and TS$^+$ for the same experimental design. Notice that TS$^+$ may require as many as two orders of magnitude more simulation effort than R&CS, whereas only occasionally does R&CS require more and then not substantially more. This is a consequence of their differing goals and approaches: TS$^+$ must control the optimality gap error across all $p = 55$ metamodels simultaneously to deliver a guaranteed EPGS, whereas R&CS attacks the R&S problem at each design point efficiently, exploiting common random numbers, and then estimates the achieved optimality gap for the database. This demonstrates the benefit of being able to exploit R&S procedures directly.

Table 3 compares the achieved EPGS of the two methods. Because the relationship between the click-through rate and the covariates is not linear, TS$^+$ falls short of the promised EPGS in cases 3, 6, and 9 but certainly achieves a practically valuable EPGS in all cases. This is in contrast to R&CS, which achieves its promised optimality gap coverage in all cases as shown in Table 1. In addition, R&CS often achieves the nominal 0.95 EPGS despite not directly targeting it, is sometimes just below it, and is substantially low in case 6. Comparing cases 1–6 to 7–12 shows that a larger database leads to a larger achieved EPGS for R&CS; cases 1–6 employ the minimum allowable design size for R&CS, and going larger is feasible because of the efficiency of R&S.

6. Discussion

R&CS uses simulation to create an off-line database of covariate values and good decisions for those values and then solves the online decision problem by assigning a class (decision) to an observed covariate using the database. R&S procedures are a natural choice for database construction, and an advantage of the R&CS approach is that it can directly exploit the vast literature on efficient R&S procedures and any new procedures created in the future.

In this paper, we consider frequentist R&S procedures that provide an indifference zone correct selection or good selection guarantee. When one employs an R&S procedure with a marginal PGS guarantee and the response means are Lipschitz in the covariate, Theorem 1 establishes the existence of a finite database delivering a postdesign pointwise PGS guarantee across the covariate space. Several such R&S procedures are available; see, for instance, Nelson and Matejcik (1995), Eckman and Henderson (2018), and Zhong and Hong (2018).

However, the postexperiment optimality gap inference associated with R&CS has no such requirements. Theorem 2 shows that this inference is distribution- and R&S procedure–free when the optimality gaps associated with the database are known, which is a remarkable

Table 2. Average Number of Replications for R&CS and TS$^+$

Case	q	K	m	R&CS replications	Standard error	TS$^+$ replications	Standard error
1	2	9	39	2.4E+06	5.2E+03	3.6E+07	3.6E+05
2	4	9	39	2.2E+06	4.4E+03	6.4E+07	5.8E+05
3	8	9	39	1.8E+06	2.3E+03	1.5E+08	1.4E+06
4	2	2	39	2.2E+06	2.8E+03	3.6E+06	3.4E+04
5	4	2	39	1.5E+06	1.6E+03	3.8E+06	3.6E+04
6	8	2	39	1.1E+06	2.0E+02	2.6E+06	2.1E+04
7	2	9	79	4.9E+06	7.2E+03	3.5E+07	2.3E+05
8	4	9	79	4.5E+06	6.5E+03	6.0E+07	3.7E+05
9	8	9	79	3.6E+06	3.3E+03	1.3E+08	7.0E+05
10	2	2	79	4.4E+06	3.7E+03	3.6E+06	2.5E+04
11	4	2	79	3.0E+06	2.2E+03	3.5E+06	2.0E+04
12	8	2	79	2.2E+06	3.1E+02	2.2E+06	1.1E+04

Table 3. Estimated EPGS for R&CS and TS$^+$

Case	q	K	m	R&CS EPGS	Standard error	TS$^+$ EPGS	Standard error
1	2	9	39	0.986	0.001	0.997	0.000
2	4	9	39	0.945	0.001	0.976	0.001
3	8	9	39	0.873	0.001	0.886	0.003
4	2	2	39	0.978	0.001	0.991	0.001
5	4	2	39	0.895	0.001	0.964	0.001
6	8	2	39	0.798	0.001	0.928	0.001
7	2	9	79	0.992	0.000	0.998	0.000
8	4	9	79	0.962	0.000	0.982	0.001
9	8	9	79	0.905	0.001	0.917	0.002
10	2	2	79	0.990	0.000	0.995	0.000
11	4	2	79	0.932	0.001	0.978	0.001
12	8	2	79	0.842	0.001	0.951	0.001

result. In practice, one needs to employ plug-in estimates of these gaps, so Theorem 3 proves the asymptotic validity of using them.

Thus, in practice, one needs R&S procedures that leave behind good point estimates of the performance of each decision and the decision corresponding to the NN for each covariate in the database. Procedure NM in Nelson and Matejcik (1995) is an example. Nevertheless, in Section 5, we chose to use a highly efficient eliminating procedure, KN (Kim and Nelson 2001) and then to enhance the relevant point estimates, and this worked very well.

Fixed-budget Bayesian and Bayesian-inspired procedures are also compatible with our postexperiment assessment, and because they typically do not eliminate any decisions, they yield good point estimates if the budget is large enough relative to the number of decisions, p. See, for instance, Chen and Lee (2011) and Ryzhov (2016). In fact, as the budget increases, they conform perfectly with the asymptotic result in Theorem 3.

Exploiting efficient R&S procedures is a computational advantage for R&CS; a second computational advantage is that R&CS can easily exploit parallel computing. Parallelizing is not easily accomplished for the metamodeling approaches to R&S+C with the exception of Shen et al. (2021). In R&CS, the application of R&S to each of the m covariate design points requires no coordination among them. Therefore, if there are at least m parallel processors available, then all m R&S optimizations to construct the database can be executed simultaneously. When the number of parallel processors is large, then groups of processors can be assigned to each design point, and parallel R&S procedures can be exploited within each group (Hunter and Nelson 2017). Only the postexperiment optimality gap assessment requires coordination.

Our postexperiment inference on the optimality gap is a random variable and, therefore, is not known in advance. We expect it to decrease as the size of the experiment design m increases as shown in our experiments. Heuristically, one could increase m until the bound is below a desired gap, δ, but at present, we cannot provide statistical guarantees on such a sequential design. However, it seems clear that adding one design point at a time until the bound is small enough is likely to lead to undercoverage; therefore, a significant number (say, 25% more design points) should be added on each round. Sequential design in this setting is a problem for future work.

Pulling back from the specific contributions of this paper, the widespread availability of data, generated by social media, transactional systems, sensors, or other means, has transformed the industrial landscape and the way of thinking in many disciplines. Most of those applications belong to the predictive realm, but an increasing number address the integration of data and decisions. Terms such as "end-to-end learning" (Wilder et al. 2019, Bergman et al. 2022) or "smart, predict then optimize" (Elmachtoub and Grigas 2022) refer to integrated approaches that connect data to decisions by combining predictive and prescriptive methods. We view R&CS as an example of how simulation can contribute to this trend.

Acknowledgments

The authors thank the associate editor and referees for their insightful and careful reviews.

References

Agrawal S, Avadhanula V, Goyal V, Zeevi A (2019) MNL-bandit: A dynamic learning approach to assortment selection. *Oper. Res.* 67(5):1453–1485.

Alban A, Chick SE, Zoumpoulis SI (2021) Expected value of information methods for contextual ranking and selection: Clinical trials and simulation optimization. Kim S, Feng B, Smith K, Masoud S, Zheng Z, eds. *Proc. 2021 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 1–12.

Bergman D, Huang T, Brooks P, Lodi A, Raghunathan AU (2022) Janos: An integrated predictive and prescriptive modeling framework. *INFORMS J. Comput.* 34(2):807–816.

Bouneffouf D, Rish I, Aggarwal C (2020) Survey on applications of multi-armed and contextual bandits. *2020 IEEE Congress Evolutionary Comput.* (IEEE, Piscataway, NJ), 1–8.

Cakmak S, Zhou E, Gao S (2021) Contextual ranking and selection with Gaussian processes. Kim S, Feng B, Smith K, Masoud S, Zheng Z, eds. *Proc. 2021 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 1–12.

Chen CH, Lee LH (2011) *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*, vol. 1 (World Scientific, Hackensack, NJ).

Chen X, Wang Y, Zhou Y (2020) Dynamic assortment optimization with changing contextual information. *J. Machine Learn. Res.* 21(1):8918–8961.

Chen X, Owen Z, Pixton C, Simchi-Levi D (2022) A statistical learning approach to personalization in revenue management. *Management Sci.* 68(3):1923–1937.

Ding L, Hong LJ, Shen H, Zhang X (2022) Knowledge gradient for selection with covariates: Consistency and computation. *Naval Res. Logist.* 69(3):496–507.

Du J, Gao S, Chen CH (2022) A contextual ranking and selection method for personalized medicine. Preprint, submitted June 25, https://arxiv.org/abs/2206.12640.

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

Downloaded from informs.org by [98.253.32.243] on 04 February 2026, at 14:54 . For personal use only, all rights reserved.

Eckman DJ, Henderson SG (2018) Guarantees on the probability of good selection. Rabe M, Juan AA, Mustafee N, Skoogh A, Jian S, Johansson B, eds. *Proc. 2018 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 351–365.

Elmachtoub AN, Grigas P (2022) Smart "predict, then optimize". *Management Sci.* 68(1):9–26.

Gao S, Du J, Chen CH (2019) Selecting the optimal system design under covariates. Reveliotis S, Capalleri D, Dimarogonas D, Dotoli M, Fanti M, Lutz P, Seatzu C, Xie X, eds. *Proc. 15th Internat. Conf. Automation Sci. Engrg.* (IEEE, Piscataway, NJ), 547–552.

Hunter SR, Nelson BL (2017) Parallel ranking and selection. Tolk A, Fowler J, Shao G, Yücesan E, eds. *Adv. Modeling Simulation: Seminal Research 50 Years Winter Simulation Conf.* (Springer, New York), 249–275.

Kallus N, Udell M (2020) Dynamic assortment personalization in high dimensions. *Oper. Res.* 68(4):1020–1037.

Keslin G, Nelson BL, Plumlee M, Pagnoncelli BK, Rahimian H (2022) A classification method for ranking and selection with covariates. Feng B, Yijie P, Pedrielli G, Eunhye S, Shashaani S, Corlu CG, eds. *Proc. 2022 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 1–12.

Kim SH, Nelson BL (2001) A fully sequential procedure for indifference-zone selection in simulation. *ACM Trans. Model. Comput. Simulation* 11(3):251–273.

Li H, Lam H, Peng Y (2024) Efficient learning for clustering and optimizing context-dependent designs. *Oper. Res.* 72(2):617–638.

Li X, Zhang X, Zheng Z (2018) Data-driven ranking and selection: High-dimensional covariates and general dependence. Rabe M, Juan AA, Mustafee N, Skoogh A, Jian S, Johansson B, eds. *Proc. 2018 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 1933–1944.

Li H, Lam H, Liang Z, Peng Y (2020) Context-dependent ranking and selection under a Bayesian framework. Bae KHG, Feng B, Kim S, Lazarova-Molnar S, Zheng Z, Roeder T, Thiesing R, eds. *Proc. 2020 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 2060–2070.

Liu H, Xiao H, Li H, Lee LH, Chew EP (2022) Offline sequential learning via simulation. *IISE Trans.* 54(11):1019–1032.

Lobel I (2021) Revenue management and the rise of the algorithmic economy. *Management Sci.* 67(9):5389–5398.

Miao S, Chao X (2022) Online personalized assortment optimization with high-dimensional customer contextual data. *Manufacturing Service Oper. Management* 24(5):2741–2760.

Nelson BL, Matejcik FJ (1995) Using common random numbers for indifference-zone selection and multiple comparisons in simulation. *Management Sci.* 41(12):1935–1945.

Rhodes-Leader L, Nelson B, Onggo BS, Worthington D (2022) A multi-fidelity modelling approach for airline disruption management using simulation. *J. Oper. Res. Soc.* 73(10):2228–2241.

Rusmevichientong P, Shen ZJM, Shmoys DB (2010) Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Oper. Res.* 58(6):1666–1680.

Ryzhov IO (2016) On the convergence rates of expected improvement methods. *Oper. Res.* 64(6):1515–1528.

Sauré D, Zeevi A (2013) Optimal dynamic assortment planning with demand learning. *Manufacturing Service Oper. Management* 15(3):387–404.

Shen H, Hong LJ, Zhang X (2021) Ranking and selection with covariates for personalized decision making. *INFORMS J. Comput.* 33(4):1500–1519.

Vovk V, Gammerman A, Shafer G (2005) *Algorithmic Learning in a Random World*, vol. 29 (Springer, New York).

Wilder B, Ewing E, Dilkina B, Tambe M (2019) End to end learning and optimization on graphs. Wallach H, Larochelle H, Beygelzimer A, d'Alch'e-Buc F, Fox E, Garnett R, eds. *Adv. Neural Inform. Processing Systems*, vol. 32 (Curran Associates, Inc., Red Hook, NY), 4674–4685.

Zhang G, Chen S, Peng Y (2023) Efficient learning for selecting top-m context-dependent designs. Preprint, submitted May 6, https://arxiv.org/abs/2305.04086.

Zhong Y, Hong LJ (2018) Fully sequential ranking and selection procedures with PAC guarantee. Rabe M, Juan AA, Mustafee N, Skoogh A, Jian S, Johansson B, eds. *Proc. 2018 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 1898–1908.

Gregory Keslin is a PhD student in the Department of Industrial Engineering and Management Sciences at Northwestern University. His research interests are simulation optimization and statistics.

Barry L. Nelson is the Walter P. Murphy Professor Emeritus in the Department of Industrial Engineering and Management Sciences at Northwestern University. His research is in the design and analysis of stochastic simulation experiments with a particular focus on simulation optimization, and he is a Fellow of INFORMS and IISE.

Bernardo Pagnoncelli is a full professor at SKEMA Business School in Lille, France. His research focuses on theory and applications of stochastic optimization.

Matthew Plumlee is a principal applied scientist at Amazon with expertise in computer experiments, statistics, uncertainty quantification and machine learning. His contribution is unrelated to his role at Amazon.

Hamed Rahimian is an assistant professor in the Department of Industrial Engineering at Clemson University. His research focuses on stochastic and robust optimization.