MIT Libraries | DSpace@MIT

MIT Open Access Articles

Covariate-Adaptive Optimization in Online Clinical Trials

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation: Bertsimas, Dimitris et al. "Covariate-Adaptive Optimization in Online Clinical Trials." Operations Research 67, 4 (May 2019): 905-1208 © 2019 The Author(s)

As Published: 10.1287/OPRE.2018.1818

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: https://hdl.handle.net/1721.1/129812

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike

Massachusetts Institute of Technology

Submitted to *Operations Research*
manuscript OPRE-2017-06-327.R2

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Covariate-Adaptive Optimization in Online Clinical Trials

Dimitris Bertsimas, Nikita Korolko, Alexander M. Weinstein
Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139,
bertsimas@mit.edu, nikita.korolko@alum.mit.edu, amw22@alum.mit.edu

The decision of how to allocate subjects to treatment groups is of great importance in experimental clinical trials for novel investigational drugs, a multi-billion dollar industry. Statistical power, the ability of an experiment to detect a positive treatment effect when one exists, depends in part on the similarity of the groups in terms of measurable covariates that affect the treatment response. We present a novel algorithm for online allocation that leverages robust mixed-integer optimization. In all tested scenarios, the proposed method yields statistical power at least as high as, and sometimes significantly higher than, state-of-the-art covariate-adaptive randomization approaches. We present a setting in which our algorithm achieves a desired level of power at a sample size 25-50% smaller than that required with randomization-based approaches. Correspondingly, we expect that our covariate-adaptive optimization approach could substantially reduce both the duration and operating costs of clinical trials in many commonly observed settings, while maintaining computational efficiency and protection against experimental bias.

Key words: Clinical trials; Covariate-adaptive randomization; Mixed-integer optimization; Online allocation; Robust optimization; Statistical power.

1. Introduction

A central task in formal experimentation for scientific inquiry is to identify unique groups of subjects which should be administered alternate treatment conditions. In particular, the groups should be as similar as possible as defined by measurable *covariates*, or attributes, which affect

an individual's response to treatment, so as distinguish the effect of treatment from confounding factors.

The decision of how to allocate subjects to treatment groups is especially impactful in the domain of experimental clinical trials for novel investigational drugs. Statistical power, the ability of an experiment to detect a positive treatment effect when one exists, depends on the size of the groups, the size of the treatment effect, and the chosen significance level, as well as, crucially, the similarity of the groups in terms of covariates that affect the treatment response. The benefits of achieving high power without increasing sample size are potentially massive. In 2010, global pharmaceutical companies spent $32.5 billion to conduct clinical trials, out of $46.4 total expenditure on research and development (Berndt and Cockburn 2013). Furthermore, drug approval requires multiple phases of clinical trials that typically take years to complete (U.S. Food and Drug Administration 2015). The financial and time pressures of this challenging process often require qualified individuals to be enrolled in a trial and quickly assigned and administered a treatment without waiting for information about the characteristics of future enrollees; the urgency is especially pressing in instances of potentially life-saving treatments.

By leveraging robust, mixed-integer optimization (MIO), we present a novel algorithm that can decrease the sample size needed to conduct a clinical trial by as much as 25-50% in certain settings, and thus substantially reduce both the cost of conducting clinical trials and the time it takes for novel effective therapies to reach patients. We extend the MIO approach of Bertsimas et al. (2015) to the online setting in which subjects arrive sequentially.

The study of treatment allocation for controlled experiments dates back to Fisher (1935). Randomization has been favored historically as a way to control for selection bias. However, randomization can yield another accidental bias identified by Efron (1971), in which there is an imbalance in the distributions of known or hidden covariates across randomly assigned treatment groups. There have been many attempts in the literature to address this accidental bias in both the offline and online allocation settings. For the offline problem, some prominent mechanisms are pairwise

matching (Rosenbaum and Rubin 1985, Greevy et al. 2004), rerandomization (Morgan and Rubin 2012), and the finite selection model (Morris 1979). Bertsimas et al. (2015) used an alternative offline optimization-based approach.

For the online sequential allocation problem, Rosenberger and Sverdlov (2008) provide an excellent review of the available heuristics for covariate-adaptive randomization, including prestratification and biased coin designs. Many of the existing heuristics stem from variations of the biased coin design first introduced by Efron (1971), including nonrandomized minimization (Taves 1974), randomized minimization (Pocock and Simon 1975), and designs that attempt to minimize the variance of the treatment effect (Atkinson 1982) or minimize loss of information (Antognini and Zagoraiou 2011). These biased coin designs outperform pure randomization and represent the current state of the art for online allocation. More recently, Kapelner and Krieger (2014) introduced a pooled sequential matching algorithm, which discards covariate data as soon as each subject is matched. Bhat, Farias, and Moallemi (2015) propose a dynamic programming algorithm for sequential allocation that comes with computational challenges for which they provide an approximation algorithm. We believe that simple randomization or biased coin designs are most commonly used in practice. While more advanced statistical methods exist, we demonstrate that robust, mixed-integer optimization has the potential to improve upon existing methods in many instances, by taking into account the uncertainty about future subjects.

In this paper, we develop a novel covariate-adaptive optimization mechanism for online allocation, which outperforms state-of-the-art covariate-adaptive randomization methods. We extend the offline mixed-integer optimization approach presented in Bertsimas et al. (2015) to the online setting in which patients arrive sequentially and each patient's covariate data cannot be observed until the time of her arrival. The new algorithm takes the form of a sequence of mixed-integer nonlinear optimization problems. We adopt a robust optimization technique with a quadratic uncertainty set (Ben-Tal, Nemirovski, and Roos 2002, Bertsimas, Brown, and Caramanis 2011), which serves as a natural approach to model the uncertainty about future subjects.

The new method, henceforth referred to as the covariate-adaptive robust optimization (CA-RO) algorithm, delivers the following benefits:

1. In all tested scenarios, the CA-RO method achieved statistical power at least as high as, and sometimes significantly higher than, covariate-adaptive randomization (CA-RAND) approaches. We present an example of a nonlinear covariate-response setting for which the CA-RO method achieved a desired level of statistical power at a sample size 25-50% smaller than that required with the best CA-RAND approach.

2. We present theoretical and empirical evidence that the optimization approach compares favorably with CA-RAND methods with respect to three advantages of complete randomization described by Efron (1971): freedom from selection bias; freedom from accidental bias with respect to observed and hidden covariates; and, a reasoned basis for inference.

3. The algorithm is sufficiently general to produce assignments among multiple groups $p = 1, \ldots, m$ with multiple observed covariates per subject. The CA-RO algorithm can also be extended to the setting where it is possible to aggregate subjects into small clusters of size r prior to making group assignments.

4. The CA-RO algorithm is computationally tractable for instances of practical size, despite taking the form of a nonlinear mixed-integer optimization problem that cannot be solved using off-the-shelf commercial solvers. Our choice of uncertainty set allows us to extract a closed-form solution for the robust constraints. Hence, we are able to solve the optimization by enumeration with complexity $\mathcal{O}(m^r)$, which does not depend on the sample size N. In all observed instances, CA-RO provides the decision-maker with a high-quality assignment recommendation instantaneously via enumeration.

The rest of the paper is organized as follows. In Section 2, we briefly revisit the optimization-based allocation algorithm for the offline setting from Bertsimas, Johnson, and Kallus (2015). This offline algorithm will form the basis for the online CA-RO approach we develop in Section 3. At the end of Section 3, we present computational results from experiments demonstrating the

effectiveness of CA-RO in reducing between-group covariate imbalance. In Section 4, we provide empirical evidence that the CA-RO algorithm achieves a high level of statistical power with much smaller sample size as compared to CA-RAND methods when the covariate-response relationship is nonlinear. In Section 5, we discuss the experimental bias and inference properties of the CA-RO approach and demonstrate that it compares favorably with CA-RAND methods. Section 6 contains concluding remarks.

2. Offline Optimization Approach

In this section, we describe a MIO approach to assign groups for the setting when pre-treatment covariate values of all subjects are known ahead of time (Bertsimas et al. 2015). The decision-maker knows *a priori* the total number of subjects N in the experiment and the respective covariates $\mathbf{w} = (w_1, \ldots, w_N)$ of all subjects. Thus, she can make treatment allocations using this full information. This may be the case, for example, in laboratory cancer drug testing on mice.

The decision-maker will assign $k := N/m$ subjects to each of $m \geq 2$ treatment groups. The objective of the assignment is to minimize the maximum discrepancy between any two groups in the weighted sum of the first and second moments of the covariates. Without loss of generality, we assume that the vector of covariates $\mathbf{w}$ is normalized and has zero sample mean and unit sample variance. The parameter ρ regulates the relative weight of the first and the second moments. The binary decision variables are $\mathbf{x} = \{x_{ip} \,|\, i = 1, \ldots, N, p = 1, \ldots, m\}$, where $x_{ip} = 1$ if subject i is assigned to group p, and $x_{ip} = 0$, otherwise. We can express the mean and second moment of each of the groups $p \in \{1, \ldots, m\}$ as follows:

$$\mu_p(\mathbf{x}) = \frac{1}{k} \sum_{i=1}^{N} w_i \, x_{ip} \quad \text{and} \quad \sigma_p^2(\mathbf{x}) = \frac{1}{k} \sum_{i=1}^{N} w_i^2 \, x_{ip}.$$

Hence, the optimal offline assignment can be found using the following MIO problem, which we henceforth refer to as the OPT algorithm:

$$\min_{\mathbf{x}} \max_{p<q} \quad |\mu_p(\mathbf{x}) - \mu_q(\mathbf{x})| + \rho |\sigma_p^2(\mathbf{x}) - \sigma_q^2(\mathbf{x})| \quad =$$

$$
\begin{aligned}
\min_{\mathbf{x},d} \quad & d \\
\text{s.t.} \quad & \forall p < q = 1, \ldots, m: \\
& \qquad d \geq \mu_p(\mathbf{x}) - \mu_q(\mathbf{x}) + \rho\sigma_p^2(\mathbf{x}) - \rho\sigma_q^2(\mathbf{x}) \\
& \qquad d \geq \mu_p(\mathbf{x}) - \mu_q(\mathbf{x}) + \rho\sigma_q^2(\mathbf{x}) - \rho\sigma_p^2(\mathbf{x}) \\
& \qquad d \geq \mu_q\mathbf{x}) - \mu_p(\mathbf{x}) + \rho\sigma_p^2(\mathbf{x}) - \rho\sigma_q^2(\mathbf{x}) \\
& \qquad d \geq \mu_q(\mathbf{x}) - \mu_p(\mathbf{x}) + \rho\sigma_q^2(\mathbf{x}) - \rho\sigma_p^2(\mathbf{x}) \\
& x_{ip} \in \{0,1\} \\
& \sum_{i=1}^{N} x_{ip} = k, \quad \forall p = 1, \ldots, m \\
& \sum_{p=1}^{m} x_{ip} = 1, \quad \forall i = 1, \ldots, N \\
& x_{ip} = 0 \quad \forall i < p.
\end{aligned}
$$

The final constraint reduces the redundancy due to permutation symmetry in group numbering.

In all tested scenarios from Bertsimas et al. (2015), the OPT method generates groups with covariate discrepancy that is exponentially lower in the group size k than those created by randomization. The expected average covariate discrepancy decreases from $O(k^{-1/2})$ for randomization to $O(2^{-k})$ for the OPT algorithm. Furthermore, the OPT algorithm demonstrates exceptional precision in estimating small treatment effects and superior statistical power given a fixed treatment effect.

For the remainder of this paper, the OPT algorithm will serve as a prescient benchmark for the performance of methods in the setting of sequential online allocation.

3. Covariate-Adaptive Optimization Algorithms

In this section, we introduce the CA-RO algorithm, develop an extension in which aggregation of decisions is allowed, and describe the results of empirical experiments comparing the covariate balance of CA-RO versus CA-RAND methods.

3.1. CA-RO Algorithm

To extend the model from Section 2 to the online setting, we consider the problem of N subjects arriving sequentially. The decision-maker knows *a priori* the number of subjects k that will be assigned to each of $m \geq 2$ treatment groups, such that $N = km$. In practical settings with sequential arrivals, it is often reasonable to assume that investigators will have at least an estimate of the number of subjects that will enroll based on the timeline, budget constraints, and prevalence of disease in the population.

At each time-step $t = 1, \ldots, N$, where t indexes both the period and the subject, the decision-maker observes the covariate vector $\mathbf{w}_t \in \mathbb{R}^S$, where S is the number of covariates observed for each subject. We assume that this sequence of random covariate vectors is exchangeable, such that any ordering of the subjects' arrival is equally likely. The decision-maker then sets a decision $\{x_{tp}\}_{p=1}^m \in \{0,1\}^m$, where $x_{tp} = 1$, if the decision-maker assigns subject t to group $p \in \{1, \ldots, m\}$, and $x_{tp} = 0$, otherwise. In the CA-RO algorithm, the choice of $\{x_{tp}\}_{p=1}^m$ is made by solving one instance of robust MIO formulation (2) at each time-step. The data for the optimization at time-step t include the covariate observations $\{\mathbf{w}_i\}_{i=1}^t$ and assignments $\hat{\mathbf{x}} := \{\hat{x}_{ip} \,\big|\, i = 1, \ldots, t-1, p = 1, \ldots, m\} \in \{0,1\}^{(t-1)\times m}$ made at all previous time-steps. We define expressions for the sample mean $\bar{\mathbf{w}}_t$ and the empirical covariance matrix Σ_t at time-step t as follows:

$$\bar{\mathbf{w}}_t := \frac{1}{t}\sum_{i=1}^{t} \mathbf{w}_i \quad \text{and} \quad \Sigma_t := \frac{1}{t}\sum_{i=1}^{t} (\mathbf{w}_i - \bar{\mathbf{w}}_t)(\mathbf{w}_i - \bar{\mathbf{w}}_t)^\top.$$

We also define uncertain parameters $\tilde{\mathbf{w}} := \{\tilde{\mathbf{w}}_i \in \mathbb{R}^S\}_{i=t+1}^N$, which represent the unknown covariates for future subjects.

The objective of the CA-RO algorithm is to produce m groups whose covariate distributions are as similar as possible. We measure the proximity between two groups p and q in terms of the mean μ_p^s and *approximated* variance σ_p^s of group $p = 1, \ldots, m$ with respect to covariate $s = 1, \ldots, S$. At time-step $1 \leq t \leq N$, these sample statistics are defined as follows:

$$\mu_p^s := \frac{1}{k}\Big\{ \sum_{i=1}^{t-1} w_i^s \hat{x}_{ip} + w_t^s x_{tp} + \sum_{i=t+1}^{N} \tilde{w}_i^s x_{ip} \Big\},$$

$$\sigma_p^s := \frac{1}{k}\Big\{\sum_{i=1}^{t-1}(w_i^s-\bar{w}_t^s)^2\hat{x}_{ip}+(w_t^s-\bar{w}_t^s)^2x_{tp}+\sum_{i=t+1}^{N}(\tilde{w}_i^s-\bar{w}_t^s)^2x_{ip}\Big\},$$

where $\mathbf{x} := \{x_{ip}\in\{0,1\}\,|\,i=t,\dots,N,\,p=1,\dots,m\}$ are the binary assignment decision variables. We model the decision at each time-step $t=1,\dots,N$ by the following optimization problem:

$$\min_{\mathbf{x}}\max_{p<q}\sum_{s=1}^{S}|\mu_p^s-\mu_q^s|+\rho|\sigma_p^s-\sigma_q^s|. \tag{1}$$

Given that the values of future covariates $\tilde{\mathbf{w}}$ are unknown, we employ robust optimization (Ben-Tal et al. 2002) to model formulation (1) under uncertainty:

$$\begin{aligned}
\min_{\mathbf{x},\mathbf{M},\mathbf{V},z}\quad & z\\
\text{s.t.}\quad & z\geq\sum_{s=1}^{S}M_{pq}^s+\rho V_{pq}^s, \qquad \forall p<q\\
& M_{pq}^s\geq\mu_p^s-\mu_q^s, \qquad \forall p<q,\,s=1,\dots,S, \quad \forall\tilde{\mathbf{w}}\in U_w\\
& M_{pq}^s\geq\mu_q^s-\mu_p^s, \qquad \forall p<q,\,s=1,\dots,S, \quad \forall\tilde{\mathbf{w}}\in U_w\\
& V_{pq}^s\geq\sigma_p^s-\sigma_q^s, \qquad \forall p<q,\,s=1,\dots,S, \quad \forall\tilde{\mathbf{w}}\in U_w\\
& V_{pq}^s\geq\sigma_q^s-\sigma_p^s, \qquad \forall p<q,\,s=1,\dots,S, \quad \forall\tilde{\mathbf{w}}\in U_w\\
& \sum_{i=1}^{t-1}\hat{x}_{ip}+x_{tp}+\sum_{i=t+1}^{N}x_{ip}=k, \quad \forall p=1,\dots,m\\
& \sum_{p=1}^{m}x_{ip}=1, \quad \forall i=t,\dots,N
\end{aligned} \tag{2}$$

In this formulation, we use the uncertainty set U_w defined as follows:

$$U_w=\Big\{\tilde{\mathbf{w}}\in\mathbb{R}^{(N-t)\times S}\quad\Big|\quad\tilde{\mathbf{w}}_i=\bar{\mathbf{w}}_t+(\Sigma_t)^{\frac{1}{2}}\boldsymbol{\varepsilon}_i,\,i=t+1,\dots,N,\boldsymbol{\varepsilon}\in U_\varepsilon\Big\},$$

where perturbation vector $\boldsymbol{\varepsilon}=(\boldsymbol{\varepsilon}_{t+1},\dots,\boldsymbol{\varepsilon}_N)$ belongs to the ellipsoidal uncertainty set U_ε:

$$U_\varepsilon=\Big\{\boldsymbol{\varepsilon}\in\mathbb{R}^{(N-t)\times S}\quad\Big|\quad||\boldsymbol{\varepsilon}||_2=\sqrt{\sum_{i=t+1}^{N}\sum_{s=1}^{S}(\varepsilon_i^s)^2}\leq\Gamma\sqrt{(N-t)S}\Big\}. \tag{3}$$

The robustness parameter Γ controls the size of the ellipsoid and represents the level of conservatism of the uncertainty set. In order to protect against experimental biases, we suggest that Γ should be chosen independently at random for each time-step (see Section 5).

Formulation (2) takes the form of a mixed-binary quadratic robust optimization problem with conic uncertainty set, which cannot be solved using off-the-shelf commercial solvers. We overcome this computational challenge by finding an efficient way to solve the following auxiliary optimization problems with respect to the uncertain variables $\tilde{\mathbf{w}}$:

$$\max_{\tilde{\mathbf{w}} \in U_w} \left(\mu_p^s - \mu_q^s\right) \quad \text{and} \quad \max_{\tilde{\mathbf{w}} \in U_w} \left(\sigma_p^s - \sigma_q^s\right). \tag{4}$$

The objectives of optimization problems (4) are to maximize a linear or quadratic function, respectively, over an ellipsoid. The unique structure of these problems allows us to derive closed-form solutions by applying Karush-Kuhn-Tucker conditions along with eigenvalue optimization. Therefore, formulation (2) is equivalent to a mixed-binary optimization problem, described in (EC.7) and (EC.8) in section EC.2 of the e-companion, that can be solved via simple enumeration of at most m scenarios.

3.2. Aggregated CA-RO algorithm

The development of a partially online method is motivated by the opportunity presented when multiple subjects enroll in a clinical trial within a short period of time. Under these circumstances, the decision-maker may be able to make a joint decision regarding the simultaneous assignment of this sub-cohort of subjects to treatment groups.

For this analysis, we will distinguish the notion of time from the arrival of subjects. Time will be indexed by periods $t = 1, \ldots, T$. Subjects will be indexed separately by $i = 1, \ldots, N$ with covariate vectors $\mathbf{w}_i \in \mathbb{R}^S$, where $N \geq T$. Both the number of periods T and the number of subjects N are known *a priori*. We assume that at time t the decision-maker has observed the covariate values for $r_t \geq 1$ unassigned subjects who have arrived during time period t. Let us define $n_t := \sum_{j=1}^{t} r_j$ to represent the number of subjects who have arrived as of time t. We also introduce the vector $\mathbf{r}_t := \{r_j\}_{j=1}^{t}$. We can then define the following expressions to represent the sample mean and approximated variance of group p with respect to covariate s at time t:

$$\mu_p^s(\mathbf{r}_t) = \frac{1}{k}\Big\{ \sum_{i=1}^{n_{t-1}} w_i^s \hat{x}_{ip} + \sum_{i=n_{t-1}+1}^{n_t} w_i^s x_{ip} + \sum_{i=n_t+1}^{N} \tilde{w}_i^s x_{ip} \Big\}, \quad \text{and}$$

$$\sigma_p^s(\mathbf{r}_t) = \frac{1}{k}\Big\{ \sum_{i=1}^{n_{t-1}} (w_i^s - \bar{w}_t^s)^2 \hat{x}_{ip} + \sum_{i=n_{t-1}+1}^{n_t} (w_i^s - \bar{w}_t^s)^2 x_{ip} + \sum_{i=n_t+1}^{N} (\tilde{w}_i^s - \bar{w}_t^s)^2 x_{ip} \Big\}.$$

In the aggregated CA-RO algorithm, we solve formulation (2) at each time-step t, but we replace the expressions μ_p^s and σ_p^s with their generalized counterparts $\mu_p^s(\mathbf{r}_t)$ and $\sigma_p^s(\mathbf{r}_t)$, respectively. The optimal solutions $\{x_{ip}^* \in \{0,1\} \,|\, i = n_t - r_t + 1, \ldots, n_t,\, p = 1, \ldots, m\}$ to the corresponding MIO problem are used to make the assignments at period t for r_t subjects. The problem can be solved at time t by enumeration with complexity $\mathcal{O}(m^{r_t})$, and is therefore computationally tractable for instances of practical size.

If the aggregation level is uniform across time such that $r_t = r$ for all $t = 1, \ldots, T-1$ and $r_T = N - (T-1)r$, we define the CA-RO(r) algorithm with aggregation level r. We observe that the CA-RO(1) algorithm is equivalent to the fully online CA-RO algorithm and the CA-RO(N) algorithm is equivalent to the OPT algorithm from Section 2.

It is reasonable to assume that larger values of the aggregation parameter r lead to better performance of the partially online algorithm in terms of both covariate balance and statistical power. With a higher level of aggregation, the decision-maker has more information at the time of each decision. In Section 3.4, we provide empirical evidence for this relationship.

3.3. Practical Considerations

When using the CA-RO algorithm in practice, we suggest a few modifications and parameter selection guidelines.

1. At the beginning of the assignment process, group indices $p = 1, \ldots, m$ can be randomly assigned to each of the treatment conditions. In this way, the CA-RO algorithm is used to identify groups that are well-balanced with respect to observed covariates, but plays no role in determining which group should receive which treatment.

2. In the objective of formulation (2), the parameter ρ controls the tradeoff between imbalance in the sample mean and the approximated variance. In practice, to facilitate an intuitive choice of ρ, it is convenient to substitute the objective $\max_{p<q} \sum_{s=1}^{S} \left[M_{pq}^s + \rho \sqrt{V_{pq}^s} \right]$, which puts the expressions for

first and second moments on the same scale. This substitution by a nonlinear objective is tractable because we are able to solve the optimization efficiently by enumeration. In the experiments that follow, we use this nonlinear objective with $\rho = 6$, which we found to yield strong results that were robust to perturbations of ρ across many instances.

3. At the beginning of the time horizon, we ensure that all groups have been assigned at least one subject at random before we apply the optimization in formulation (2).

4. Toward the end of the time horizon, we set the robustness parameter $\Gamma = 0$ so as to make our algorithm more greedy and avoid overly conservative assignment decisions.

5. In all tested experiments, the CA-RO(r) algorithms for $r \in \{1, 3, 5\}$ produced assignment recommendations instantaneously (i.e. under 1 second), which suggests that the method can be used not just for clinical trials, but also for settings requiring real-time decisions, such as experimentation for website design or marketing.

6. One can imagine a situation in which it is beneficial to exclude a given patient from the trial on the basis of extreme covariate values rather than make an assignment. While we do not consider this scenario directly in this paper, this question could be addressed through a pre-filter outlier detection process. There are modifications that could be made to the MIO formulation to allow for selective exclusion, but we leave this exploration to future work.

3.4. Empirical Performance

In this subsection, we evaluate the empirical performance of the CA-RO algorithm. First, we review four state-of-the-art CA-RAND methods, which serve as benchmarks for the CA-RO algorithm. Second, we compare the performance of the CA-RO algorithm at various aggregation levels with pure randomization and these CA-RAND methods.

When evaluating the performance of CA-RO, we consider pure randomization (RAND) along with the matching on-the-fly algorithm of Kapelner and Krieger (2014) (KK), and three biased coin designs: the minimization method of Pocock and Simon (1975) (PS), the D_A-optimal design of

Atkinson (1982) (DA), and the covariate-adaptive biased coin design of Antognini and Zagoraiou (2011) (AZ). The biased coin design methodology with $m=2$, generically defined as

$$\phi_t = \Pr\Big(x_{t1}=1 \Big| \hat{\mathbf{x}}_1,\dots,\hat{\mathbf{x}}_{t-1};\, \mathbf{w}_1,\dots,\mathbf{w}_t\Big) = F\Big(\hat{\mathbf{x}}_1,\dots,\hat{\mathbf{x}}_{t-1};\, \mathbf{w}_1,\dots,\mathbf{w}_t\Big),$$

forms the basis of the PS, DA and AZ methods, with function $F(\cdot)$ defined separately for each method. For the PS method, the bias ϕ_t of the coin used in randomized allocation is defined as

$$\phi_t = \begin{cases} \frac{1}{2}, & \text{if } D(t)=0, \\ p, & \text{if } D(t)<0, \\ 1-p, & \text{if } D(t)>0, \end{cases}$$

where $p \geq 0.5$ is the bias parameter and $D(t) = 2\sum_{i=1}^{t-1} \hat{x}_i - (t-1)$ represents the covariate imbalance between the two groups after $t-1$ subjects have been assigned (with $D(1)=0$). For the DA method,

$$\phi_t = \frac{(1-\zeta)^2}{(1-\zeta)^2+(1+\zeta)^2}, \quad \text{where} \quad \zeta = \frac{w_t \sum_{i=1}^{t-1} w_i(x_{i1}-x_{i2})}{\sum_{i=1}^{t-1} w_i^2}$$

for the case of one-dimensional covariates. In this definition, the value ζ represents the normalized mismatch between the two groups after $t-1$ assignments. For the AZ method, we first define $J+1$ discrete levels ω_j of the covariate space. If subject t arrives from level ω_j, then the probability of assignment of this subject into group 1 is $\phi_t = G_j(D_t(\omega_j))$, where non-increasing and symmetric functions G_j are defined as follows:

$$G_j(\zeta) = \begin{cases} \frac{1}{2}, & 0 \leq \zeta \leq 1, \\ (\zeta^J+1)^{-1}, & \zeta > 1, \end{cases} \qquad \text{and} \quad G_j(-\zeta)+G_j(\zeta)=1,\ \forall \zeta \in, \mathbb{Z},\ j=0,\dots J.$$

In this description, $D_t(\omega_j)$ denotes the imbalance between the two groups within the level ω_j. In the KK method, subjects are either randomized to treatment groups or paired via a matching criterion based on the pairwise Mahalanobis distance. In the latter case, the new paired subject is assigned to the treatment opposite its pair in order to balance the groups.

We now discuss the empirical performance of the various algorithms with respect to covariate balance. For this analysis we took data from a Mayo Clinic trial of treatment for primary biliary cirrhosis of the liver conducted between 1974 and 1984 (Therneau and Grambsch 2000) with a total of 312 patients described by 20 covariates. After running a simple logistic regression model predicting the outcome of the treatment, we identified three statistically significant covariates with the largest effect: age (C1), alkaline phosphatase in U/liter (C2), and prothrombin time in seconds (C3). We evaluated nine algorithms - RAND, PS, DA, AZ, KK and CA-RO(r) (with four different values of r) - to measure the average worst pairwise difference in generalized moments across groups with respect to three covariates described above (Table 1). For this and all subsequent experiments when evaluating the CA-RO algorithm at any level of aggregation, we chose the robustness parameter Γ in uncertainty set (3) independently and uniformly at random from the interval [0.5, 4] at each time-step. In terms of the discrepancy in the first moment, CA-RO outperformed other CA-RAND algorithms for all standardized covariates. The discrepancy in the second moment, which closely approximates the discrepancy in the variance in this setting, was also always lower for CA-RO than for the best CA-RAND method. The discrepancy in higher moments, as well as generalized moments of $\log(|w|)$ and $1/w$, for CA-RO methods was always comparable with other CA-RAND algorithms. As we expected, the advantage of optimization increases with r, and the offline OPT algorithm has significantly better performance than other approaches.

We found similar (and often stronger) results from additional experiments with synthetic data in which the covariates were generated from different types of distributions, including normal, uniform, and long-tailed Cauchy distributions. One of these experiments is discussed in Section EC.1 of the e-companion.

4. Statistical Power of CA-RO Algorithm

A common pre-condition for the approval of any clinical trial is to demonstrate that the trial will have a sample size sufficient to make sound statistical inferences with high probability. These inferences include both statistical power, the ability to detect a positive treatment effect when

one exists, and a low type I error rate, the ability to correctly identify an ineffective treatment. In classical statistical models, the power of a randomized controlled trial can be derived from the sample size and significance level, given an estimated treatment effect. Randomized allocation can yield an accidental imbalance in covariates between treatment groups that can impact the ability to make experimental inferences. Traditionally, when estimating treatment effects, practitioners have been satisfied to control for this covariate imbalance *a posteriori* via regression methods (Lin 2013).

We provide strong empirical evidence that such post hoc adjustments may produce suboptimal effect estimation, particularly when the relationship between covariates and response is nonlinear. By testing a variety of covariate-response models, we show that, at any fixed sample size, the statistical power of a clinical trial is at least as high when covariate-adaptive optimization is used rather than covariate-adaptive randomization. In settings where the covariate-response relationship is nonlinear, we observe that the power under the CA-RO algorithm is significantly higher than for state-of-the-art CA-RAND methods. Therefore, in certain settings, the use of covariate-adaptive optimization could allow decision-makers to achieve desired levels of statistical power with significantly smaller sample size as compared with CA-RAND mechanisms. Given the high cost of enrolling human subjects in clinical trials, the ability to achieve needed statistical power with much smaller sample size can result in significant cost savings for the healthcare industry and society at large.

4.1. Test for Statistical Power

In order to compare statistical power under CA-RAND and CA-RO online allocation procedures, we apply a hypothesis testing framework based on simulation (Bertsimas et al. 2015).

Let us assume there are $m = 2$ groups: a treatment group, which will be administered a given therapy, and a control group, which will be administered a placebo. There are N subjects in the trial, such that $k = N/2$ subjects will be assigned to each of the groups. At each time-step $t = 1, \ldots, N$, the decision-maker observes the values of a covariate vector $\mathbf{w}_t$ and makes a binary

assignment x_t, where $x_t = 1$ indicates the treatment group (1) and $x_t = 0$ indicates the control group (0). Let v_t be the response measured after the assigned treatment was administered for subject t. We adopt the potential outcomes framework of Rosenbaum and Rubin (1983), such that each subject has a pair of potential outcomes $(v_t^{(1)}, v_t^{(0)})$, where the superscript indicates treatment or control and only one of these two outcomes can be observed. Under this framework, we have the following relationship between the observed response and the potential outcomes: $v_t = v_t^{(1)} x_t + v_t^{(0)}(1 - x_t)$.

Given v_t for each subject $t = 1, \ldots, N$, we can estimate the average treatment effect $\hat{\delta}$. We adopt two estimators for $\hat{\delta}$, unadjusted and regression-adjusted, from Lin (2013):

1. $\hat{\delta}_{\text{unadj}} := \frac{1}{k}\Big[\sum_{t=1}^{N} v_t x_t - \sum_{t=1}^{N} v_t(1 - x_t)\Big]$
2. $\hat{\delta}_{\text{adj}} := \beta_x$, where β_x is the estimated coefficient on x_t in the ordinary least squares regression $v_t = \beta_0 + \beta_x x_t + \boldsymbol{\beta}_w^\top \mathbf{w}_t$.

To test the significance of this observed effect $\hat{\delta}$, we adopt Fisher's sharp null hypothesis (Fisher 1935), which states that every subject $t = 1, \ldots, N$ would have had the same response to treatment regardless of which treatment was assigned. In other words, under the sharp null hypothesis, we have $v_t = v_t^{(1)} = v_t^{(0)}$. Equipped with a complete set of potential outcomes for each subject, we can estimate the average treatment effect under alternative random allocations of subjects $1, \ldots, N$. If we compute the average treatment effect δ_b as our test statistic for each alternative allocation $b = 1, \ldots, B$, we can then estimate the p-value for our observed $\hat{\delta}$, using a two-sided test, as

$$p = \frac{1}{1 + B}\Big(1 + \sum_{b=1}^{B} \mathbb{I}\Big[|\delta_b| \geq |\hat{\delta}|\Big]\Big).$$

We reject the null hypothesis if $p \leq \alpha$ for some pre-specified significance level α; otherwise, we accept the null hypothesis.

In order to estimate the statistical power under a given algorithm $\mathbb{A}$, we generate Q random samples of N subjects with covariates drawn *i.i.d.* from a fixed distribution. We apply the hypothesis test described above for all random samples, and measure the number of samples Q_{reject} for which the null hypothesis is rejected. We evaluate the probability that the null hypothesis will be

rejected by computing the ratio $\lambda := Q_{\text{reject}}/Q$. If the true treatment effect δ_0 is nonzero, then λ estimates the power of the experiment; otherwise, λ estimates the type I error rate.

The alternative allocations for the hypothesis test can be generated randomly using Monte Carlo simulation to approximate the distribution of possible allocations under random assignment mechanism $\mathbb{A}$. If Monte Carlo simulation does not yield a sufficiently diverse set of allocations within computational limits, one can generate bootstrapped resamples of covariate vectors $\mathbf{w}_t^b$, $t = 1, \ldots, N$ drawn uniformly at random from the set $\mathcal{W} = \{\mathbf{w}_1, \ldots, \mathbf{w}_N\}$ (Efron and Tibshirani 1994). Based on the observations from the original experiment and under the null hypothesis, we have complete mappings $v^{(1)}(\cdot) : \mathcal{W} \to \mathbb{R}$ and $v^{(0)}(\cdot) : \mathcal{W} \to \mathbb{R}$, which represent the potential outcomes under treatment and control, respectively, for individuals with covariates in $\mathcal{W}$. Therefore, for each subject in a given bootstrapped sample, we observe the response under her random allocation x_t^b as $v_t^b = v^{(1)}(\mathbf{w}_t^b) \cdot x_t^b + v^{(0)}(\mathbf{w}_t^b) \cdot (1 - x_t^b)$.

4.2. Computational Results

To evaluate the statistical power of the CA-RO algorithm relative to CA-RAND methods, we simulated clinical trials under three different hidden realities, each characterized by a unique model relating treatment response to subject covariates. We assumed each subject $t = 1, \ldots, N$ had a covariate vector $\mathbf{w}_t = (w_t^1, w_t^2)$ of dimension $S = 2$, whose components were drawn *i.i.d.* from a standard normal distribution. The covariate-response models were as follows:

- Nonlinear (NL): $v_t = \delta_0 x_t + (w_t^1)^2 - (w_t^2)^2 + \epsilon_t$,
- Linear (LIN): $v_t = \delta_0 x_t + 2(w_t^1) + 2(w_t^2) + \epsilon_t$,
- No relationship (NR): $v_t = \delta_0 x_t + \epsilon_t$,

where δ_0 is the ground-truth additive treatment effect and ϵ_t is a Gaussian noise term with mean 0 and standard deviation 0.75.

For each covariate-response model, we evaluated statistical power λ under the CA-RAND and CA-RO algorithms by applying the hypothesis test described in Section 4.1 with both estimators $\hat{\delta}_{\text{unadj}}$ and $\hat{\delta}_{\text{adj}}$ (Figure 1). We considered $N \in \{40, 80, 120\}$ with $\delta_0 = 0.5$, $Q = 800$, $B = 500$, and significance level $\alpha = 0.05$. For all scenarios, the power of the experiment increases with N.

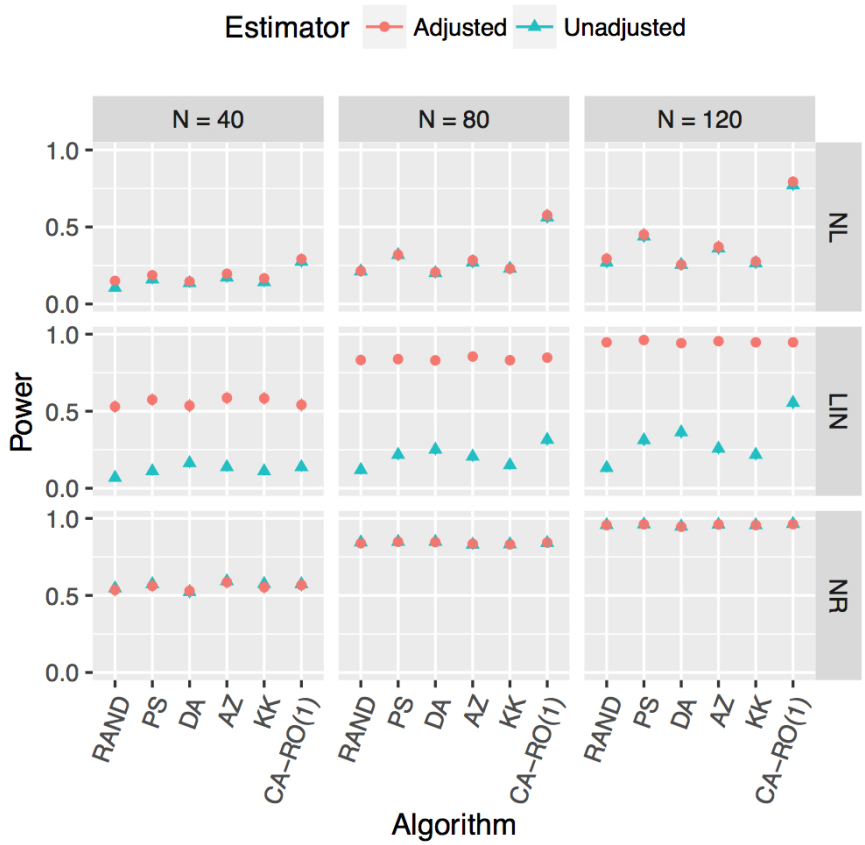

Figure 1 Statistical power (with 95% confidence intervals) under CA-RO(1) vs. CA-RAND methods for $N \in \{40, 80, 120\}$ under various response models (NL, LIN, NR), using both adjusted and unadjusted treatment effect estimators.

- In the NR scenario, post hoc regression adjustment does not improve power for any of the methods. All methods yield similar power since there is no benefit from covariate balance.
- Conversely, in the linear response setting (LIN), regression adjustment increases statistical power substantially for all methods. When using the $\hat{\delta}_{\text{unadj}}$ estimator, CA-RO(1) yields higher power relative to randomization and CA-RAND methods. However, post hoc regression adjustment using ordinary least squares, which exactly replicates the covariate-response model with additive treatment effect, reduces the need for the *a priori* covariate balance provided by CA-RO. Power evaluated using $\hat{\delta}_{\text{adj}}$ is equally high across all methods.
- Finally, in the nonlinear response scenario described above (NL), there is virtually no benefit to using regression adjustment. In this setting, CA-RO(1) yields much higher statistical power than

pure randomization and CA-RAND methods. The advantage of CA-RO grows with the sample size N.

We conducted additional experiments under a variety of nonlinear models and found, in all tested scenarios, that CA-RO had power at least as high as (and often higher than) randomization-based methods. The NL scenario is an example in which the benefit of CA-RO was particularly dramatic.

We also ran simulations in which the p-values were estimated using a one-sided test rather than a two-sided test. As one might expect, for fixed δ_0 and distribution of noise ϵ, power was higher for all methods under the one-sided test. However, in all tested scenarios, the CA-RO algorithm maintained a similar advantage relative to other methods.

In Table 2, we show the results for the NL setting under the CA-RO(r) assignment mechanism for $N = 40$ with aggregation levels $r \in \{1, 3, 5\}$ along with OPT, which is equivalent to CA-RO(N). As we expect, the power increases with the aggregation level r.

In Figure 1, we show that, under some covariate-response models, CA-RO yields higher power at fixed sample sizes than other methods. This motivates a complementary question: What is the sample size required to achieve a desired level of statistical power? We considered the NL scenario and tested values of δ_0 from 0.75 to 1.75 to estimate $N^*_{\mathbb{A}}(\delta_0)$, the minimum number of subjects per group needed to achieve power of at least 80% when assignment mechanism $\mathbb{A}$ is employed (Figure 2a). With a large effect size of $\delta_0 = 1.75$, statistical power of 80% was achieved with a sample size of 22 using the CA-RO(1) algorithm compared with a sample size of 30 using the best CA-RAND method (in this case, PS). With a small effect size of $\delta_0 = 0.75$, the advantage of optimization was even bigger; a sample size of 58 was sufficient to achieve 80% power, compared with a sample size of 122 using the best CA-RAND method (again, PS). For a given treatment effect δ_0, the threshold sample size needed to achieve 80% power under the CA-RO algorithm was reduced by at least 25% relative to the best CA-RAND method (Figure 2b). If we consider the NL setting with $\delta_0 = 0.75$ as an example, the CA-RO method may enable the execution and analysis of some clinical trials that would otherwise be infeasible given the prohibitively large sample size required to achieve a sufficient level of statistical power when CA-RAND methods are employed.

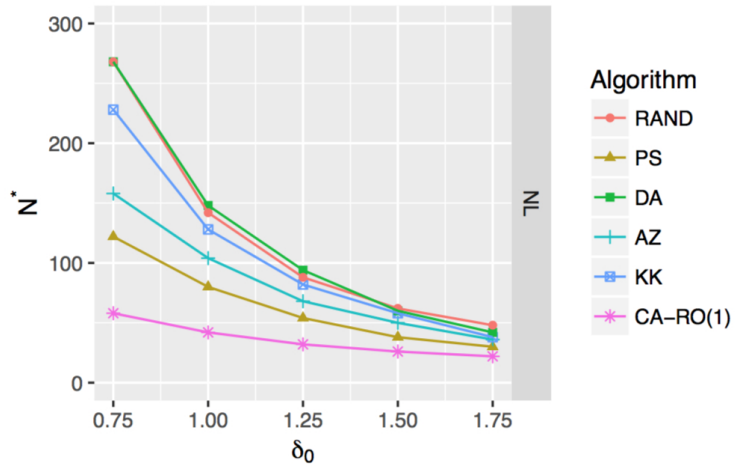

(a) Sample size needed to achieve statistical power of at least 80%.

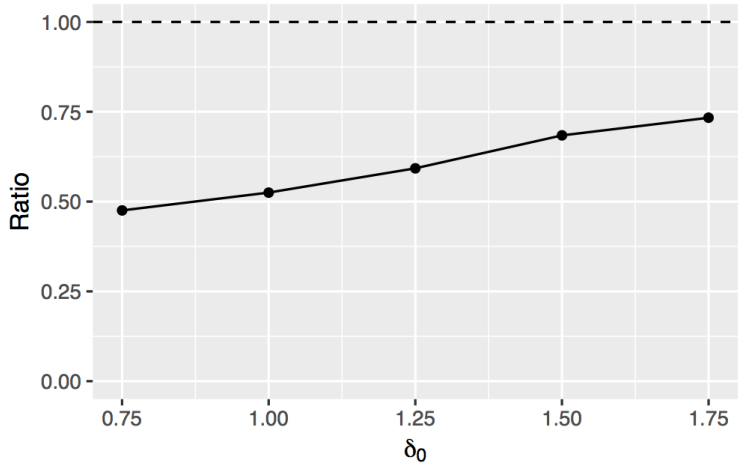

(b) Ratio of $N^*(\delta_0)$ for CA-RO(1) vs. best CA-RAND algorithm (PS).

Figure 2 Statistical power under CA-RO vs. CA-RAND methods.

Table 3 demonstrates that the minimum sample size $N^*_{\mathbb{A}}(\delta_0)$ decreases further as the aggregation level r of CA-RO(r) algorithm grows. Relative to state-of-the-art CA-RAND methods, the CA-RO approach can dramatically reduce the number of subjects enrolled in a trial without sacrificing statistical power.

We also evaluate the rate of type I errors for CA-RO(1) with $\delta_0 = 0$ with $Q = 800$ and $B = 500$ for $N \in \{40, 80, 120\}$ and for each of the three covariate-response scenarios described above, using the regression-adjusted treatment effect estimator (Table 4). We observe that, for each setting, the type I error is a decreasing function of N. Type I error rates for all algorithms tested are shown in Figure EC.1 in the e-companion. PS, the CA-RAND method with the best statistical power in this experiment, had a mean type I error rate that was uniformly higher than that produced by CA-RO(1).

5. Unbiasedness of CA-RO Approach

In this section, we provide empirical and theoretical evidence that the CA-RO algorithm introduced in Section 3 exhibits the same statistical advantages ascribed to complete randomization by Efron (1971): freedom from selection bias, freedom from accidental bias with respect to observed and hidden covariates, and a reasoned basis for inference.

5.1. Freedom from selection bias

The CA-RO algorithm protects against selection bias, the possibility that an investigator could consciously or unconsciously influence the order of subject enrollment based on deterministic knowledge of the next treatment assignment. Through computer simulation, we demonstrate that, by selecting the robustness parameter Γ independently and uniformly at random with support [0.5, 4] at each time-step, the CA-RO method yields sufficiently random treatment assignments as to protect against this type of selection bias. For N from 30 to 100, we randomly generated 30 unique sequences of covariates $\mathbf{w} \in \mathbb{R}^N$ drawn independently from $\mathcal{N}(0, 1)$. We used the CA-RO(1) algorithm to generate 3,000 random assignments of the N subjects to two groups. We observe that one cannot determine the sequence of future assignments based on knowledge of the algorithm because, on average, the total number of possible allocations is large (Figure 3a) and no individual assignment sequence has a likelihood higher than 6% (Figure 3b).

A concern that directly competes with selection bias is certifiability. When used with a fixed and predefined sequence of robustness parameters Γ_t, $t = 1, \ldots, N$, the CA-RO algorithm is a sequence

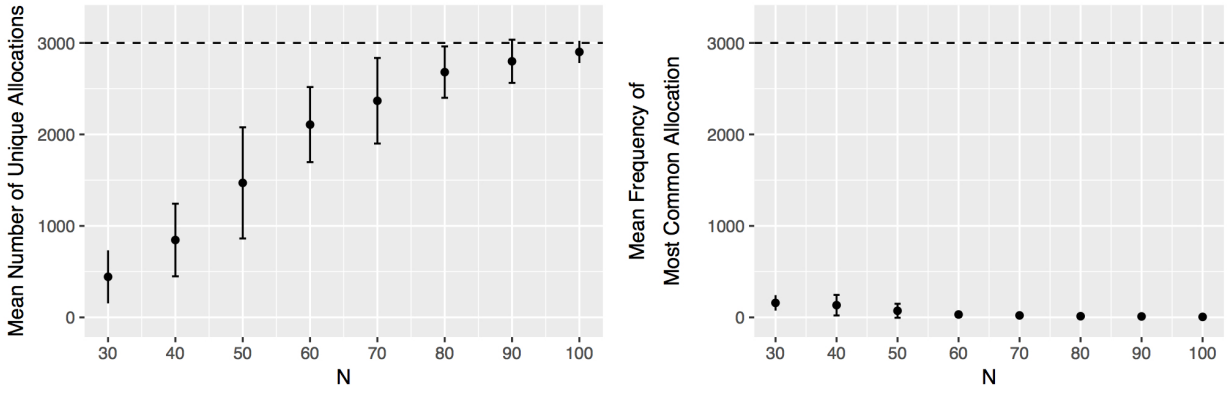

(a) Mean (with standard deviation) number of unique allocations among 3,000 simulations.

(b) Mean (with standard deviation) frequency of most common allocation among 3,000 simulations.

Figure 3 Analysis of distribution of unique allocations under CA-RO(1).

of deterministic optimization problems, each of which can be reproduced. This reproducibility provides a natural method for certifying *a posteriori* that the algorithm's recommendation was followed, given knowledge of the subjects' covariates and arrival order. If certifiability is deemed to be of greater concern than selection bias in the context of a particular experimental setting, one can apply the CA-RO algorithm with fixed robustness parameters in order to achieve full certifiability of assignments. Conversely, certifiability is not achievable using randomized methods unless the random seed used to initialize the algorithm is provided.

5.2. Freedom from accidental covariate imbalance

We have shown in the empirical results from Section 3.4 that the CA-RO method produces consistently better balance in the first two moments across groups compared with simple randomization and other existing CA-RAND approaches. In this subsection, we show that, despite only considering the observed covariates $\mathbf{w} \in \mathbb{R}^N$ when making assignment decisions, the CA-RO algorithm provides the same level of protection as CA-RAND methods against irregular allocation with respect to other, potentially unseen factors.

We consider two natural cases for the dependence of hidden factors on the observed covariates $\mathbf{w}$: no correlation and continuous dependence.

1. If there is a hidden factor that is uncorrelated with observed covariates $\mathbf{w}$, the CA-RO algorithm generates an allocation which is as random with respect to the hidden covariates as that produced by randomized methods.

2. The second case, when the unseen covariate is a continuous function of the observed covariate, warrants further discussion. We see empirically that, when unseen factor f has a polynomial or logarithmic conditional expectation in scalar random variable w, the discrepancy in higher moments and generalized moments $f = \log(|w|)$ and $f = 1/w$ for CA-RO methods is always comparable with (and often lower than) the mismatch produced by CA-RAND algorithms (Table 1). In the remainder of this subsection, we present formal theoretical evidence that this empirical relationship extends to the general case of continuous dependence.

To examine this general case, we assume that there are two different assignment algorithms $\mathbb{A}$ and $\mathbb{B}$ (e.g. CA-RO(1) and PS), and an unseen factor f that can be modeled in the form

$$f = g(w) + \epsilon,$$

where $g(\cdot)$ is a Lipschitz function with constant L and ϵ is some noise function. When generating groups of size k by algorithm $\mathbb{A}$, let us denote the maximum discrepancy in means with respect to unseen covariate f by:

$$z^f_{\mathbb{A}} := \max_{p<q} \frac{1}{k} \Big| \sum_{i \in \mathrm{I}_p(\mathbb{A})} g(w_i) - \sum_{i \in \mathrm{I}_q(\mathbb{A})} g(w_i) \Big|,$$

where $I_p(\mathbb{A})$, $I_q(\mathbb{A}) \subset \{1, \ldots, N\}$ are disjoint index sets respectively describing groups p and q produced by algorithm $\mathbb{A}$. The maximum discrepancy $z^f_{\mathbb{B}}$ between groups generated by algorithm $\mathbb{B}$ is defined analogously.

In Proposition 1, we derive theoretical upper bounds on the the values of $z^f_{\mathbb{A}}$ and $|z^f_{\mathbb{A}} - z^f_{\mathbb{B}}|$, where $\mathbb{A}$ is the CA-RO(1) approach and $\mathbb{B}$ is any CA-RAND method. The first upper bound on $z^f_{\mathbb{A}}$, given by (7a), demonstrates that the maximum discrepancy in means with respect to unseen covariate f is controlled by the corresponding discrepancy with respect to the observed covariate w. The second upper bound on $|z^f_{\mathbb{A}} - z^f_{\mathbb{B}}|$, given by (7b), indicates that the maximum discrepancy in means with

respect to unseen covariate f is as well-controlled under CA-RO(1) as under any other CA-RAND algorithm.

Before we present Proposition 1, let us define some auxiliary notations. Sets of indices S_α^β, for $\alpha, \beta = 1, 2$, describing the differences between the groups produced by algorithms $\mathbb{A}$ and $\mathbb{B}$ have the form

$$S_1^1 = \mathrm{I}_1(\mathbb{A}) \cap \mathrm{I}_2(\mathbb{B}) \quad \text{and} \quad S_2^1 = \mathrm{I}_2(\mathbb{A}) \cap \mathrm{I}_1(\mathbb{B});$$

$$S_1^2 = \mathrm{I}_1(\mathbb{A}) \cap \mathrm{I}_1(\mathbb{B}) \quad \text{and} \quad S_2^2 = \mathrm{I}_2(\mathbb{A}) \cap \mathrm{I}_2(\mathbb{B}).$$

We also introduce the optimal objective value $\theta^*(\mathbb{A})$ of the auxiliary pairwise matching problem:

$$\begin{aligned} \theta^*(\mathbb{A}) := \min_{\mathbf{y}} \ & \tfrac{1}{k} \sum_{i \in \mathrm{I}_1(\mathbb{A})} \sum_{j \in \mathrm{I}_2(\mathbb{A})} |w_i - w_j| \, y_{ij} \\ \text{s.t.} \quad & \sum_{i \in \mathrm{I}_1(\mathbb{A})} y_{ij} = 1, \quad \forall j \in \mathrm{I}_2(\mathbb{A}) \\ & \sum_{j \in \mathrm{I}_2(\mathbb{A})} y_{ij} = 1, \quad \forall i \in \mathrm{I}_1(\mathbb{A}) \\ & y_{ij} \in \{0, 1\}. \end{aligned} \tag{5}$$

Finally, we define the value $\xi^*(\mathbb{A}, \mathbb{B}) := \min_{c=1,2} \xi_c(\mathbb{A}, \mathbb{B})$, where $\xi_c(\mathbb{A}, \mathbb{B})$ is the optimal value of the problem

$$\begin{aligned} \xi_c(\mathbb{A}, \mathbb{B}) := \min_{\mathbf{y}} \ & \sum_{i \in S_1^c} \sum_{j \in S_2^c} |w_i - w_j| \, y_{ij} \\ \text{s.t.} \quad & \sum_{i \in S_1^c} y_{ij} = 1, \quad \forall j \in S_2^c \\ & \sum_{j \in S_2^c} y_{ij} = 1, \quad \forall i \in S_1^c \\ & y_{ij} \in \{0, 1\}. \end{aligned} \tag{6}$$

Proposition 1. Let us consider the simplest case where subjects with scalar covariates $w_i, i = 1, \ldots, 2k$ are assigned to $m = 2$ groups. For any assignment algorithms $\mathbb{A}$ and $\mathbb{B}$ that produce groups of equal size k, and for any Lipschitz function $g(\cdot) \in \mathrm{Lip}(L)$, the following inequalities hold:

$$z_{\mathbb{A}}^f \le L \cdot \theta^*(\mathbb{A}), \tag{7a}$$

$$|z_{\mathbb{A}}^f - z_{\mathbb{B}}^f| \le 2L \cdot \xi^*(\mathbb{A}, \mathbb{B}). \tag{7b}$$

Proof. The proof of Proposition 1 is presented in the e-companion, section EC.4. □

Having obtained theoretical upper bounds (7), we conducted numerical experiments to measure the values of the average pairwise distances $\theta^*(\mathbb{A})$ and $\xi^*(\mathbb{A},\mathbb{B})$ defined in (5) and (6), respectively. We fixed $\mathbb{A} = \text{CA-RO(1)}$ and chose $\mathbb{B}$ from among RAND, PS, DA, AZ and KK, where the randomization methods were modified to ensure they would produce equal-sized groups at the end of the horizon. We randomly generated populations of size N between 60 and 100, where each subject had a scalar standard normal covariate w_i. After executing both chosen algorithms $\mathbb{A}$ and $\mathbb{B}$, we identified the index sets S_α^β, for $\alpha,\beta = 1,2$ and solved the auxiliary optimization problems (5) and (6). After 3,000 simulations, we observed that, in more than 99% of instances, the average discrepancy $\theta^*\big(\text{CA-RO(1)}\big) \le 0.35$; the corresponding upper bounds for $\xi^*\big(\text{CA-RO(1)},\mathbb{B}\big)$ for $\mathbb{B} =$ RAND, PS, DA, AZ, KK are equal to 0.255, 0.187, 0.2, 0.185, 0.215, respectively.

Given that, by definition, the distances θ^* and ξ^* scale linearly with respect to the covariates w_i, $i = 1,\dots,N$, one may derive the empirical counterparts of upper bounds (7), which hold with high probability:

$$z^f_{\text{CA-RO(1)}} \le 0.35\, L\cdot\sigma$$

$$\max_{\mathbb{B}\in\{\text{RAND, PS, DA, AZ, KK}\}} |z^f_{\text{CA-RO(1)}} - z^f_{\mathbb{B}}| \le 0.51\, L\cdot\sigma$$

where σ is the standard deviation of attributes w. The constant 0.51 in the right-hand side of the second bound is derived from the maximum discrepancy $\xi^*\big(\text{CA-RO(1)},\mathbb{B}\big)$ among CA-RAND methods reported above.

The result of Proposition 1 can be extended to the case of general continuous functions $g(\cdot)$ under the assumption that the support $K \subset \mathbb{R}^S$ of covariates $\mathbf{w}$ is a compact set. Indeed, any continuously differentiable function $g(\cdot)$ defined on a compact set K (including any polynomial function) is in a Lipschitz class with $L = \max\limits_{x\in K} |g'(x)|$. Since any continuous function on K can be approximated with arbitrary precision by some polynomial according to the Weierstrass theorem, the upper bounds (7) hold for any continuous function $g(\cdot)$ on the set K.

5.3. Reasoned basis for inference

The results from Figures 3a and 3b, which demonstrate the variety of unique allocations that can result under the CA-RO approach, indicate that CA-RO provides a sufficient degree of randomization to be used as a reasoned basis for inference. While the probability distribution of these allocations does not appear to be uniform (see Figure 3b), the fact that diverse allocations arise motivates us to conduct randomization-inspired tests for statistical significance such that the power of the CA-RO method can be estimated under various scenarios in Section 4.

6. Conclusions

In this paper, we introduced a covariate-adaptive optimization algorithm for the problem of online allocation of subjects in randomized controlled trials. Our method leverages robust mixed-integer quadratic optimization to improve upon state-of-the-art covariate-adaptive randomization methods. We demonstrated many desirable properties of the new CA-RO approach, including computational tractability, smaller between-group covariate imbalance as compared with randomization-based methods, and a low potential for common experimental biases. In all tested scenarios, the CA-RO method performed competitively with CA-RAND approaches, and sometimes significantly outperformed these methods, as measured by statistical power. We presented a setting with a non-linear covariate-response model for which the CA-RO method achieved a desired level of statistical power at a sample size 25-50% smaller than the best CA-RAND method. Thus, the proposed CA-RO algorithm has significant potential to reduce both the cost and duration of clinical trials. The CA-RO algorithm can be used to make assignments to any arbitrary number of treatment groups and for any number of observed covariates. Finally, we constructed an extension of the CA-RO method for the setting in which it is possible to aggregate decision-making. We believe that the proposed CA-RO algorithm is an efficient alternative to covariate-adaptive randomization that can significantly strengthen experimental power in clinical trials and many other disciplines exploiting controlled experiments.

Acknowledgments

We would like to thank Professor Colin Fogarty of MIT Sloan School of Management for sharing his technical expertise. This research was partially supported by the Office of Naval Research, Grant 021152-00001.

Biographies

Dimitris Bertsimas is the Boeing Professor of Operations Research, Codirector of the Operations Research Center at Massachusetts Institute of Technology, and a member of the National Academy of Engineering. He works on optimization, statistics, stochastics, and their applications. The present paper is part of the authors research in the interface of optimization and statistics.

Nikita Korolko received his PhD degree from the Operations Research Center at Massachusetts Institute of Technology. His current research interests include robust optimization and its applications to online problems. He has held research and data science positions at Mitsubishi Electric Research Laboratories, Tesla Motors and Uber Technologies.

Alexander M. Weinstein received his PhD degree from the Operations Research Center at Massachusetts Institute of Technology. His research interests include applications of optimization, statistics, and machine learning in domains, such as healthcare and transportation. He has worked in applied settings as a research scientist at Amazon, Zillow Group, and, most recently, Lyft.

References

Antognini AB, Zagoraiou M (2011) The covariate-adaptive biased coin design for balancing clinical trials in the presence of prognostic factors. *Biometrika* 98(3):519–535.

Atkinson AC (1982) Optimum biased coin designs for sequential clinical trials with prognostic factors. *Biometrika* 69(1):61–67.

Ben-Tal A, Nemirovski A, Roos C (2002) Robust solutions of uncertain quadratic and conic-quadratic problems. *SIAM Journal on Optimization* 13(2):535–560.

Berndt ER, Cockburn IM (2013) Price indexes for clinical trial research: a feasibility study. Technical report, National Bureau of Economic Research.

Bertsimas D, Brown DB, Caramanis C (2011) Theory and applications of robust optimization. *SIAM review* 53(3):464–501.

Bertsimas D, Johnson M, Kallus N (2015) The power of optimization over randomization in designing experiments involving small samples. *Operations Research* 63(4):868–876.

Bhat N, Farias VF, Moallemi CC (2015) Optimal A-B testing. *Working paper* .

Efron B (1971) Forcing a sequential experiment to be balanced. *Biometrika* 58(3):403–417.

Efron B, Tibshirani RJ (1994) *An Introduction to the Bootstrap* (CRC Press).

Fisher RA (1935) *The Design of Experiments* (Oliver and Boyd, Edinburgh).

Greevy R, Lu B, Silber JH, Rosenbaum P (2004) Optimal multivariate matching before randomization. *Biostatistics* 5(2):263–275.

Kapelner A, Krieger A (2014) Matching on-the-fly: Sequential allocation with higher power and efficiency. *Biometrics* 70(2):378–388.

Lin W (2013) Agnostic notes on regression adjustments to experimental data: Reexamining Freedman's critique. *The Annals of Applied Statistics* 7(1):295–318.

Morgan KL, Rubin DB (2012) Rerandomization to improve covariate balance in experiments. *The Annals of Statistics* 40(2):1263–1282.

Morris C (1979) A finite selection model for experimental design of the health insurance study. *Journal of Econometrics* 11(1):43–61.

Pocock SJ, Simon R (1975) Sequential treatment assignment with balancing for prognostic factors in the controlled clinical trial. *Biometrics* 31(1):103–115.

Rosenbaum PR, Rubin DB (1983) The central role of the propensity score in observational studies for causal effects. *Biometrika* 41–55.

Rosenbaum PR, Rubin DB (1985) Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician* 39(1):33–38.

Rosenberger WF, Sverdlov O (2008) Handling covariates in the design of clinical trials. *Statistical Science* 23(3):404–419.

Taves DR (1974) Minimization: a new method of assigning patients to treatment and control groups. *Clinical pharmacology and therapeutics* 15(5):443.

Therneau T, Grambsch P (2000) *Modeling survival data: extending the Cox model* (Springer-Verlag, New York).

US Food and Drug Administration (2015) *The Drug Development Process: Clinical Research*. Http://www.fda.gov/ForPatients/Approvals/Drugs/ucm405622.htm.

Table 1: Average absolute between-group discrepancy in moments under allocation algorithms for dataset with 3 covariates ($m = 2, S = 3$).

		Moment						
Covariate	Algorithm	1	2	3	4	5	$\log(\lvert w\rvert)$	$1/w$
(C1)	RAND	0.093	0.105	0.265	0.503	1.263	0.089	1.061
	PS	0.039	0.092	0.214	0.484	1.188	0.066	0.920
	DA	0.039	0.138	0.178	0.546	1.074	0.182	1.050
	AZ	0.057	0.107	0.226	0.503	1.179	0.101	1.004
	KK	0.070	0.103	0.227	0.473	1.151	0.092	1.034
	CA-RO(1)	0.024	0.070	0.158	0.331	0.969	0.083	1.091
	CA-RO(3)	0.024	0.059	0.162	0.329	0.997	0.086	1.051
	CA-RO(5)	0.026	0.048	0.160	0.311	0.974	0.079	1.053
	OPT	0.005	7.52×10^{-5}	0.150	0.537	1.064	0.260	1.112
(C2)	RAND	0.089	0.306	1.414	6.914	35.419	0.098	1.947
	PS	0.051	0.250	1.238	6.319	33.313	0.068	1.888
	DA	0.040	0.209	1.014	5.401	29.513	0.149	1.973
	AZ	0.073	0.299	1.408	6.994	36.190	0.105	1.879
	KK	0.048	0.171	0.897	4.934	27.866	0.082	2.058
	CA-RO(1)	0.028	0.093	0.660	4.206	25.515	0.097	1.918
	CA-RO(3)	0.028	0.083	0.620	4.066	25.036	0.092	1.940
	CA-RO(5)	0.028	0.073	0.561	3.790	23.788	0.095	1.999
	OPT	0.004	7.91×10^{-5}	0.567	3.779	23.593	0.302	1.885
(C3)	RAND	0.092	0.275	1.656	10.488	66.429	0.088	0.684
	PS	0.047	0.257	1.600	10.303	65.643	0.058	0.509
	DA	0.041	0.250	1.526	10.102	64.966	0.166	0.653

AZ	0.060	0.257	1.621	10.384	66.111	0.107	0.599
KK	0.068	0.208	1.376	9.262	60.995	0.086	0.703
CA-RO(1)	0.025	0.101	1.247	9.009	61.335	0.114	0.703
CA-RO(3)	0.026	0.089	1.214	8.886	60.911	0.118	0.696
CA-RO(5)	0.028	0.085	1.192	8.784	60.324	0.118	0.702
OPT	0.005	7.92×10^{-5}	1.345	9.717	66.206	0.180	0.728

Table 2 Statistical power under CA-RO(r) for NL scenario with $N = 40$.

Aggregation level, r	1	3	5	N
Power, λ	29.1%	29.8%	31.9%	36.4%

Table 3 Minimum number of subjects per group $N^*_{\mathbb{A}}(\delta_0)$ needed for power over 80%.

	Treatment effect, δ_0				
Algorithm, $\mathbb{A}$	0.75	1	1.25	1.5	1.75
RAND	268	142	88	62	48
PS	122	80	54	38	30
DA	268	148	94	60	42
AZ	158	104	68	50	36
KK	228	128	82	58	38
CA-RO(1)	58	42	32	26	22
CA-RO(3)	52	36	28	22	18
CA-RO(5)	48	32	26	22	18
OPT	26	22	18	16	14

Table 4 Type I error under CA-RO(1).

	Sample size, N		
Scenario	40	80	120
Nonlinear (NL)	7.1%	4.6%	4.5%
Linear (LIN)	7.0%	6.0%	4.5%
No relationship (NR)	6.5%	5.3%	5.6%

Electronic Companion to Covariate-Adaptive Optimization in Online Clinical Trials

EC.1. Empirical performance: Synthetic data

In this section, we discuss the empirical performance of the various algorithms with respect to covariate balance using synthetic data. For $N \in \{20, 60, 100\}$ and $m = 2$, we simulated 3,000 unique sets of covariate values drawn *i.i.d.* from a standard normal distribution. We evaluated nine algorithms - RAND, PS, DA, AZ, KK and CA-RO(r) (with four different values of r) - to measure the average worst pairwise difference in generalized moments across groups (Table EC.1). Similarly as before, when evaluating the CA-RO algorithm at any level of aggregation, we chose the robustness parameter Γ in uncertainty set (3) independently and uniformly at random from the interval [0.5, 4] at each time-step. In terms of the discrepancy in the first moment, CA-RO was always among the best methods. The discrepancy in the second moment, which closely approximates the discrepancy in the variance in this setting, was always lower for CA-RO than for the best CA-RAND method. The discrepancy in higher moments, as well as generalized moments of $\log(|w|)$ and $1/w$, for CA-RO methods was always comparable with other CA-RAND algorithms. As we expected, the advantage of optimization increases with r, and the offline OPT algorithm has starkly better performance than other approaches.

EC.2. Tractability of the CA-RO algorithm

Lemma 1. Robust optimization problem (2) with ellipsoidal uncertainty set U_ε as defined in (3) is equivalent to mixed-binary closed-form optimization problem (EC.7) in case $S \geq 2$ and (EC.8) in case $S = 1$.

Proof. In order to model the constraints for each p, q, s from optimization problem (2) that should hold for all possible realizations of uncertain vector $\tilde{\mathbf{w}} \in U_w$, we will find a closed-form solution to the following auxiliary optimization problems, repeated from (4):

$$\max_{\tilde{\mathbf{w}} \in U_w} (\mu_p^s - \mu_q^s) \quad \text{and} \quad \max_{\tilde{\mathbf{w}} \in U_w} (\sigma_p^s - \sigma_q^s).$$

Step 1. Optimization of the linear term.

Let us define a parameter $\tilde{\Gamma} := \Gamma^2(N-t)S$, where Γ is the robustness parameter from (3). Then, for any fixed values of p, q, s, and $\tilde{\Gamma}$, we consider the optimization problem

$$\max_{\tilde{\mathbf{w}} \in U_w} (\mu_p^s - \mu_q^s). \tag{EC.1}$$

We have

$$k(\mu_p^s - \mu_q^s) = \sum_{i=1}^{t-1} w_i^s (\hat{x}_{ip} - \hat{x}_{iq}) + w_t^s (x_{tp} - x_{tq}) + \sum_{i=t+1}^{N} \tilde{w}_i^s (x_{ip} - x_{iq}),$$

where only the last term of the right-hand side depends on uncertain $\tilde{\mathbf{w}}$. Therefore, we need to solve the following optimization problem for fixed values of components of $\mathbf{x}$:

$$\begin{aligned}\max_{\tilde{\mathbf{w}} \in U_w} \sum_{i=t+1}^{N} \tilde{w}_i^s (x_{ip} - x_{iq}) &= \max_{\boldsymbol{\varepsilon} \in U_\varepsilon} \sum_{i=t+1}^{N} (\bar{w}_t^s + \mathbf{v}_{(s)}^\top \boldsymbol{\varepsilon}_i)(x_{ip} - x_{iq}) \\ &= \bar{w}_t^s \sum_{i=t+1}^{N} (x_{ip} - x_{iq}) + \max_{\boldsymbol{\varepsilon} \in U_\varepsilon} \sum_{i=t+1}^{N} (\mathbf{v}_{(s)}^\top \boldsymbol{\varepsilon}_i)(x_{ip} - x_{iq}).\end{aligned}$$

where $\mathbf{v}_{(s)}$ denotes the s-th row of the matrix $(\Sigma_t)^{\frac{1}{2}}$. The optimization problem

$$\max_{\boldsymbol{\varepsilon} \in U_\varepsilon} \sum_{i=t+1}^{N} (\mathbf{v}_{(s)}^\top \boldsymbol{\varepsilon}_i)(x_{ip} - x_{iq})$$

can be rewritten in the following form:

$$\begin{aligned} \max_{\boldsymbol{\varepsilon}} \quad & (\mathbf{a}^{pqs})^\top \boldsymbol{\varepsilon} \\ \text{s.t.} \quad & \boldsymbol{\varepsilon}^\top \boldsymbol{\varepsilon} \le \tilde{\Gamma}, \end{aligned} \tag{EC.2}$$

where vector $\mathbf{a}^{pqs}$ of dimension $(N-t) \times S$ is defined by $(\mathbf{a}^{pqs})_{is'} = (x_{ip} - x_{iq})(\Sigma_t^{\frac{1}{2}})_{ss'}$ for $i = t+1, \ldots, N$, $s' = 1, \ldots, S$.

Application of the Karush-Kuhn-Tucker conditions yields that the optimal value of the optimization problem (EC.2) is equal to

$$\sqrt{\tilde{\Gamma}} \cdot \|\mathbf{a}^{pqs}\|_2 = \sqrt{\tilde{\Gamma}} \sqrt{\sum_{s'=1}^{S} \sum_{i=t+1}^{N} (x_{ip} - x_{iq})^2 ((\Sigma_t^{\frac{1}{2}})_{ss'})^2} = \sqrt{\tilde{\Gamma}} \|\mathbf{v}_{(s)}\|_2 \sqrt{\sum_{i=t+1}^{N} (x_{ip} - x_{iq})^2}.$$

The last factor can be simplified and expressed in terms of the current time-step decision variables as follows:

$$\sum_{i=t+1}^{N}(x_{ip}-x_{iq})^2=\sum_{i=t+1}^{N}(x_{ip}^2-2x_{ip}x_{iq}+x_{iq}^2)=\sum_{i=t+1}^{N}(x_{ip}+x_{iq})=2k-\sum_{i=1}^{t-1}(\hat{x}_{ip}+\hat{x}_{iq})-(x_{tp}+x_{tq}),$$

where the second equality is due to the fact that x_{ip} and x_{iq} are binary variables with $x_{ip}x_{iq}=0$. Thus, the analysis of optimization problem (EC.1) allows us to write a closed-form counterpart of the linear terms in (2) that depends only on the current time-step decision variables x_{tp} for $p=1,\dots,m$, such that:

$$\begin{aligned}
&M^s_{pq}\geq \mu^s_p-\mu^s_q,\quad \forall\tilde{\mathbf{w}}\in U_w \quad\Longleftrightarrow\\
&k\,M^s_{pq}\geq\sum_{i=1}^{t-1}(w^s_i-\bar{w}^s_t)(\hat{x}_{ip}-\hat{x}_{iq})+(w^s_t-\bar{w}^s_t)(x_{tp}-x_{tq})+\\
&\qquad\qquad+\sqrt{\tilde{\Gamma}}\|\mathbf{v}_{(s)}\|_2\sqrt{2k-\sum_{i=1}^{t-1}(\hat{x}_{ip}+\hat{x}_{iq})-(x_{tp}+x_{tq})}.
\end{aligned}$$

Step 2. Optimization of the variance term.

Similarly to Step 1, we fix values of p,q,s and $\tilde{\Gamma}$ and consider the optimization problem

$$\max_{\tilde{\mathbf{w}}\in U_w}(\sigma^s_p-\sigma^s_q). \tag{EC.3}$$

As before, only the term representing the future time periods depends on the uncertain parameters $\boldsymbol{\varepsilon}$. Therefore, the primary goal of this step is to find a closed-form solution to the auxiliary optimization problem

$$\max_{\boldsymbol{\varepsilon}\in U_\varepsilon}\sum_{i=t+1}^{N}(\mathbf{v}^\top_{(s)}\boldsymbol{\varepsilon}_i)^2(x_{ip}-x_{iq})=\max_{\|\boldsymbol{\varepsilon}\|_2^2\leq\tilde{\Gamma}}\boldsymbol{\varepsilon}^\top A\boldsymbol{\varepsilon}=\tilde{\Gamma}\cdot\lambda_{\max}(A). \tag{EC.4}$$

In (EC.4), $\lambda_{\max}(A)$ denotes the maximum eigenvalue of the square block matrix A:

$$A=\begin{bmatrix}
(x_{t+1,p}-x_{t+1,q})B & 0 & 0\dots & 0\\
0 & (x_{t+2,p}-x_{t+2,q})B & 0\dots & 0\\
\dots & \dots & \dots & \dots\\
0 & 0 & 0\dots & (x_{Np}-x_{Nq})B
\end{bmatrix},$$

where matrix $B = \mathbf{v}_{(s)}(\mathbf{v}_{(s)})^\top$.

The maximum eigenvalue $\lambda_{\max}(A)$ depends not only on the values of $\mathbf{x}$, but also on the dimension S of the covariate space.

- **Case 1.** $S \geq 2$. In this case, the eigenvalues of matrix B are 0 and $\|\mathbf{v}_{(s)}\|_2^2$, and the maximum eigenvalue of matrix A can be determined as a function of $\mathbf{x}$ as follows:

$$\lambda_{\max}(A) = \begin{cases} 0, & \text{if } x_{ip} - x_{iq} \leq 0 \text{ for all } i = t+1, \dots, N, \\ \|\mathbf{v}_{(s)}\|_2^2, & \text{if } x_{ip} - x_{iq} = 1 \text{ for at least one } i = t+1, \dots, N. \end{cases}$$

By construction, the condition $x_{ip} - x_{iq} = 1$ for at least one $i = t+1, \dots, N$ holds if and only if group p is not full after the current time-step assignment, i.e.,

$$k - \sum_{i=1}^{t-1} \hat{x}_{ip} - x_{tp} \geq 1.$$

Thus, optimization problem (EC.4) has the following closed-form solution that depends only on the current time-step decision variables:

$$\max_{\boldsymbol{\varepsilon} \in U_\varepsilon} \sum_{i=t+1}^{N} (\mathbf{v}_{(s)}^\top \boldsymbol{\varepsilon}_i)^2 (x_{ip} - x_{iq}) = \tilde{\Gamma} \cdot \|\mathbf{v}_{(s)}\|_2^2 \cdot \Psi_p(\hat{\mathbf{x}}, \mathbf{x}),$$

where

$$\Psi_p(\hat{\mathbf{x}}, \mathbf{x}) = \begin{cases} 1, & \text{if } k - \sum\limits_{i=1}^{t-1} \hat{x}_{ip} - x_{tp} \geq 1, \\ 0, & \text{otherwise.} \end{cases} \tag{EC.5}$$

Now we can exploit the closed-form solution for optimization problem (EC.3) within (2), as follows:

$$\begin{aligned} V_{pq}^s \geq \sigma_p^s - \sigma_q^s, \quad \forall \tilde{\mathbf{w}} \in U_w \quad &\iff \\ k\, V_{pq}^s \geq \sum_{i=1}^{t-1} (w_i^s - \bar{w}_t^s)^2 (\hat{x}_{ip} - \hat{x}_{iq}) + (w_t^s - \bar{w}_t^s)^2 (x_{tp} - x_{tq}) + & \\ + \tilde{\Gamma} \cdot \|\mathbf{v}_{(s)}\|_2^2 \cdot \Psi_p(\hat{\mathbf{x}}, \mathbf{x}). & \end{aligned}$$

- **Case 2.** $S=1$. In this case, matrix B is one-dimensional and its only eigenvalue is $\|\mathbf{v}_{(s)}\|_2^2$.

Hence,

$$\lambda_{\max}(A)=\begin{cases}\|\mathbf{v}_{(s)}\|_2^2, & \text{if } x_{ip}=1 \text{ for at least one } i=t+1,\dots,N,\\ -\|\mathbf{v}_{(s)}\|_2^2, & \text{if } x_{ip}=0 \text{ and } x_{iq}=1 \text{ for all } i=t+1,\dots,N,\\ 0, & \text{if } x_{ip}=0 \text{ for all } i=t+1,\dots,N \text{ and}\\ & \quad x_{iq}=0 \text{ for at least one } i=t+1,\dots,N.\end{cases}$$

This is equivalent to the formulation: $\lambda_{\max}(A)=\|\mathbf{v}_{(s)}\|_2^2\cdot\Theta_{pq}(\hat{\mathbf{x}},\mathbf{x})$, where

$$\Theta_{pq}(\hat{\mathbf{x}},\mathbf{x})=\begin{cases}1, & \text{if } k-\sum_{i=1}^{t-1}\hat{x}_{ip}-x_{tp}\geq 1,\\ -1, & \text{if } k-\sum_{i=1}^{t-1}\hat{x}_{ip}-x_{tp}=0 \text{ and } \sum_{i=1}^{t-1}\hat{x}_{iq}+x_{tq}+(N-t)=k,\\ 0, & \text{if } k-\sum_{i=1}^{t-1}\hat{x}_{ip}-x_{tp}=0 \text{ and } \sum_{i=1}^{t-1}\hat{x}_{iq}+x_{tq}+(N-t)>k.\end{cases} \tag{EC.6}$$

Thus, optimization problem (2) modeling the CA-RO algorithm with ellipsoidal uncertainty set has the following closed form for $S\geq 2$:

$$\begin{aligned}\min_{\mathbf{x},\mathbf{M},\mathbf{V},z}\quad & z\\ \text{s.t.}\quad & z\geq\sum_{s=1}^{S}M_{pq}^s+\rho V_{pq}^s, \qquad \forall p<q\\ & \forall p<q,\quad s=1,\dots,S:\\ & k\,M_{pq}^s\geq\sum_{i=1}^{t-1}(w_i^s-\bar{w}_t^s)(\hat{x}_{ip}-\hat{x}_{iq})+(w_t^s-\bar{w}_t^s)(x_{tp}-x_{tq})+\\ &\qquad\qquad +\sqrt{\tilde{\Gamma}}\|\mathbf{v}_{(s)}\|_2\sqrt{2k-\sum_{i=1}^{t-1}(\hat{x}_{ip}+\hat{x}_{iq})-(x_{tp}+x_{tq})}\\ & k\,M_{pq}^s\geq\sum_{i=1}^{t-1}(w_i^s-\bar{w}_t^s)(\hat{x}_{iq}-\hat{x}_{ip})+(w_t^s-\bar{w}_t^s)(x_{tq}-x_{tp})+\\ &\qquad\qquad +\sqrt{\tilde{\Gamma}}\|\mathbf{v}_{(s)}\|_2\sqrt{2k-\sum_{i=1}^{t-1}(\hat{x}_{ip}+\hat{x}_{iq})-(x_{tp}+x_{tq})}\\ & k\,V_{pq}^s\geq\sum_{i=1}^{t-1}(w_i^s-\bar{w}_t^s)^2(\hat{x}_{ip}-\hat{x}_{iq})+(w_t^s-\bar{w}_t^s)^2(x_{tp}-x_{tq})+\end{aligned} \tag{EC.7}$$

$$+\tilde{\Gamma}\cdot\|\mathbf{v}_{(s)}\|_2^2\cdot\mathbb{I}\Big\{k-\sum_{i=1}^{t-1}\hat{x}_{ip}-x_{tp}\geq 1\Big\}$$

$$k\,V^s_{pq}\geq\sum_{i=1}^{t-1}(w^s_i-\bar{w}^s_t)^2(\hat{x}_{iq}-\hat{x}_{ip})+(w^s_t-\bar{w}^s_t)^2(x_{tq}-x_{tp})+$$

$$+\tilde{\Gamma}\cdot\|\mathbf{v}_{(s)}\|_2^2\cdot\Psi_q(\hat{\mathbf{x}},\mathbf{x})$$

$$\sum_{i=1}^{t-1}\hat{x}_{ip}+x_{tp}\leq k,\quad\forall p=1,\dots,m$$

$$\sum_{p=1}^{m}x_{tp}=1$$

$$x_{ip}\in\{0,1\},\quad\forall i=t,\dots,N,\,p=1,\dots,m.$$

The second-to-last constraint guarantees that no group will be assigned more than k subjects and is therefore a sufficient replacement for the second-to-last constraint of formulation (2).

A similar formulation for the case $S=1$ is given by

$$\begin{aligned}
\min_{\mathbf{x},\mathbf{M},\mathbf{V},z}\quad & z\\
\text{s.t.}\quad & z\geq\sum_{s=1}^{S}M^s_{pq}+\rho V^s_{pq},\qquad\forall p<q\\
& \forall p<q,\quad s=1,\dots,S:\\
& k\,M^s_{pq}\geq\sum_{i=1}^{t-1}(w^s_i-\bar{w}^s_t)(\hat{x}_{ip}-\hat{x}_{iq})+(w^s_t-\bar{w}^s_t)(x_{tp}-x_{tq})+\\
& \qquad+\sqrt{\tilde{\Gamma}}\|\mathbf{v}_{(s)}\|_2\sqrt{2k-\sum_{i=1}^{t-1}(\hat{x}_{ip}+\hat{x}_{iq})-(x_{tp}+x_{tq})}\\
& k\,M^s_{pq}\geq\sum_{i=1}^{t-1}(w^s_i-\bar{w}^s_t)(\hat{x}_{iq}-\hat{x}_{ip})+(w^s_i-\bar{w}^s_t)(x_{iq}-x_{ip})+\\
& \qquad+\sqrt{\tilde{\Gamma}}\|\mathbf{v}_{(s)}\|_2\sqrt{2k-\sum_{i=1}^{t-1}(\hat{x}_{ip}+\hat{x}_{iq})-(x_{tp}+x_{tq})}\\
& k\,V^s_{pq}\geq\sum_{i=1}^{t-1}(w^s_i-\bar{w}^s_t)^2(\hat{x}_{ip}-\hat{x}_{iq})+(w^s_t-\bar{w}^s_t)^2(x_{tp}-x_{tq})+\\
& \qquad+\tilde{\Gamma}\cdot\|\mathbf{v}_{(s)}\|_2^2\cdot\Theta_{pq}(\hat{\mathbf{x}},\mathbf{x})\\
& k\,V^s_{pq}\geq\sum_{i=1}^{t-1}(w^s_i-\bar{w}^s_t)^2(\hat{x}_{iq}-\hat{x}_{ip})+(w^s_t-\bar{w}^s_t)^2(x_{tq}-x_{tp})+
\end{aligned}\tag{EC.8}$$

$$+\tilde{\Gamma}\cdot\|\mathbf{v}_{(s)}\|_2^2\cdot\Theta_{qp}(\hat{\mathbf{x}},\mathbf{x})$$

$$\sum_{i=1}^{t-1}\hat{x}_{ip}+x_{tp}\le k,\quad \forall p=1,\dots,m$$

$$\sum_{p=1}^{m}x_{tp}=1$$

$$x_{ip}\in\{0,1\},\quad \forall i=t,\dots,N,\ p=1,\dots,m,$$

where $\Psi_p(\hat{\mathbf{x}},\mathbf{x})$ and $\Psi_q(\hat{\mathbf{x}},\mathbf{x})$ are as defined in (EC.5); $\Theta_{pq}(\hat{\mathbf{x}},\mathbf{x})$ and $\Theta_{qp}(\hat{\mathbf{x}},\mathbf{x})$ are as defined in (EC.6). □

Corollary 1. To find the optimal objective value of discrete optimization problem (EC.7) and (EC.8) and the optimal current assignment at time t, it is sufficient to inspect the following easily specified set $\mathcal{X}$ consisting of not more than m points:

$$\mathcal{X}:=\bigcup_{p=1}^{m}\Big\{\, x_{tp}=1;$$

$$x_{tq}=0,\quad \forall q=1,\dots,m,\ q\neq p;$$

$$x_{iu}=0,\quad \forall i=t+1,\dots,N,\ u=1,\dots,m;$$

$$\sum_{i=1}^{t-1}\hat{x}_{ip}+x_{tp}\le k\,\Big\}.$$

Proof. Formulations (EC.7) and (EC.8) depend only on current time-step decisions x_{tp}, for $p=1,\dots,m$. Given that these variables are binary and the subject with index t must be assigned to exactly one group, it is sufficient to inspect the set $\mathcal{X}$, with cardinality at most m, to solve (2) for the optimal current assignment. □

EC.3. CARO algorithm with patient exclusion

In this section, we consider an extended clinical trial setting that allows a decision-maker to exclude some patients from a trial due to their extreme covariate values that may aggravate final between-group balance.

Following the notation from Section 3, we assume that in the new setting the total number of subjects in a trial is $N+N_0$, where N_0 is the maximum number of subjects to be excluded, as

specified by the decision-maker. The definitions of the mean μ_p^s and approximated variance σ_p^s of group $p=1,\dots,m$ with respect to covariate $s=1,\dots,S$ at time-step $1\le t\le N+N_0$ are almost unchanged:

$$\mu_p^s := \frac{1}{k}\Big\{\sum_{i=1}^{t-1} w_i^s \hat{x}_{ip} + w_t^s x_{tp} + \sum_{i=t+1}^{N+N_0} \tilde{w}_i^s x_{ip}\Big\},$$
$$\sigma_p^s := \frac{1}{k}\Big\{\sum_{i=1}^{t-1} (w_i^s - \bar{w}_t^s)^2 \hat{x}_{ip} + (w_t^s - \bar{w}_t^s)^2 x_{tp} + \sum_{i=t+1}^{N+N_0} (\tilde{w}_i^s - \bar{w}_t^s)^2 x_{ip}\Big\},$$

where $\mathbf{x} := \{x_{ip}\in\{0,1\} \,\big|\, i=t,\dots,N+N_0,\, p=1,\dots,m\}$ are the binary assignment decision variables. The only new decision variables that we introduce are binary indicators

$$y_i = \begin{cases} 1, & \text{if we exclude patient } i \text{ from the trial,} \\ 0, & \text{otherwise.} \end{cases}$$

for $i=t,\dots,N+N_0$. In this case, at each time-step $t=1,\dots,N+N_0$ we model the assignment or exclusion decision by the following robust optimization problem (that has the same structure as problem (2)):

$$\begin{aligned}
\min_{\mathbf{x},\mathbf{y},\mathbf{M},\mathbf{V},z} \quad & z \\
\text{s.t.} \quad & z \ge \sum_{s=1}^{S} M_{pq}^s + \rho V_{pq}^s, \qquad \forall p<q \\
& M_{pq}^s \ge \mu_p^s - \mu_q^s, \qquad \forall p<q,\, s=1,\dots,S, \quad \forall \tilde{\mathbf{w}}\in U_w \\
& M_{pq}^s \ge \mu_q^s - \mu_p^s, \qquad \forall p<q,\, s=1,\dots,S, \quad \forall \tilde{\mathbf{w}}\in U_w \\
& V_{pq}^s \ge \sigma_p^s - \sigma_q^s, \qquad \forall p<q,\, s=1,\dots,S, \quad \forall \tilde{\mathbf{w}}\in U_w \\
& V_{pq}^s \ge \sigma_q^s - \sigma_p^s, \qquad \forall p<q,\, s=1,\dots,S, \quad \forall \tilde{\mathbf{w}}\in U_w \\
& \sum_{i=1}^{t-1} \hat{x}_{ip} + x_{tp} + \sum_{i=t+1}^{N+N_0} x_{ip} = k, \quad \forall p=1,\dots,m \\
& y_i + \sum_{p=1}^{m} x_{ip} = 1, \quad \forall i=t,\dots,N+N_0 \\
& \sum_{i=1}^{t-1} \hat{y}_i + \sum_{i=t}^{N+N_0} y_i \le N_0.
\end{aligned} \tag{EC.9}$$

In this formulation, the last two constraints state that each subject should be either excluded from the trial or assigned to one of the m groups; and the total number of excluded subjects cannot exceed the upper bound N_0. The definition of uncertainty set U_w takes into account an increased number of time-steps:

$$U_w = \Big\{ \tilde{\mathbf{w}} \in \mathbb{R}^{(N+N_0-t)\times S} \quad \Big| \quad \tilde{\mathbf{w}}_i = \bar{\mathbf{w}}_t + (\Sigma_t)^{\frac{1}{2}}\, \boldsymbol{\varepsilon}_i,\, i = t+1, \ldots, N+N_0, \boldsymbol{\varepsilon} \in U_\varepsilon \Big\},$$

where perturbation vector $\boldsymbol{\varepsilon} = (\boldsymbol{\varepsilon}_{t+1}, \ldots, \boldsymbol{\varepsilon}_{N+N_0})$ belongs to the ellipsoidal uncertainty set U_ε:

$$U_\varepsilon = \Big\{ \boldsymbol{\varepsilon} \in \mathbb{R}^{(N+N_0-t)\times S} \quad \Big| \quad ||\boldsymbol{\varepsilon}||_2 = \sqrt{\sum_{i=t+1}^{N+N_0} \sum_{s=1}^{S} (\varepsilon_i^s)^2} \le \Gamma \sqrt{(N+N_0-t)S} \Big\}.$$

Optimization formulation (EC.9) can be solved using the same enumeration method that we designed for the initial formulation (2).

EC.4. Proof of Proposition 1

Proof. In order to verify inequality (7a) for $m = 2$, we note that optimization problem (5) uniquely determines a pairwise matching of sets $\mathrm{I}_1(\mathbb{A})$ and $\mathrm{I}_2(\mathbb{A})$ with minimum average distance between pairs. We denote the resulting pairs as $\{(i_l, j_l) : l = 1, \ldots, k\}$, where k is the number of indices in each set. By the definition of $z^f_{\mathbb{A}}$, we derive

$$z^f_{\mathbb{A}} = \frac{1}{k} \Big| \sum_{i \in \mathrm{I}_1(\mathbb{A})} g(w_i) - \sum_{i \in \mathrm{I}_2(\mathbb{A})} g(w_i) \Big| \le \frac{1}{k} \sum_{l=1}^{k} |g(w_{i_l}) - g(w_{j_l})| \le \frac{L}{k} \sum_{l=1}^{k} |w_{i_l} - w_{j_l}| = L \cdot \theta^*(\mathbb{A}). \quad \text{(EC.10)}$$

Similar reasoning is applicable for the second inequality (7b). First, it is easy to see that the cardinality of both sets S_1^1 and S_2^1 is the same:

$$\gamma := |S_1^1| = |S_2^1|.$$

By symmetry, the cardinalities of the complementary sets are also identical:

$$|S_1^2| = |S_2^2| = k - \gamma.$$

The next step is to express the between-group discrepancies in the means generated by algorithms $\mathbb{A}$ and $\mathbb{B}$ as follows:

$$z^f_{\mathbb{A}} = \frac{1}{k}\Big|\sum_{i\in S_1^2} g(w_i) + \sum_{i\in S_1^1} g(w_i) - \sum_{i\in S_2^1} g(w_i) - \sum_{i\in S_2^2} g(w_i)\Big| = |a+b|.$$

$$z^f_{\mathbb{B}} = \frac{1}{k}\Big|\sum_{i\in S_1^2} g(w_i) - \sum_{i\in S_1^1} g(w_i) + \sum_{i\in S_2^1} g(w_i) - \sum_{i\in S_2^2} g(w_i)\Big| = |a-b|,$$

where

$$a := \frac{1}{k}\Big(\sum_{i\in S_1^2} g(w_i) - \sum_{i\in S_2^2} g(w_i)\Big) \quad \text{and} \quad b := \frac{1}{k}\Big(\sum_{i\in S_1^1} g(w_i) - \sum_{i\in S_2^1} g(w_i)\Big).$$

Hence, analogously to argument (EC.10), one may obtain upper bounds:

$$|a| \le L\cdot\xi_2(\mathbb{A},\mathbb{B}) \quad \text{and} \quad |b| \le L\cdot\xi_1(\mathbb{A},\mathbb{B}). \tag{EC.11}$$

A simple corollary from the triangle inequality is that, for any a and b,

$$\big||a+b| - |a-b|\big| \le 2\min\big(|a|,|b|\big).$$

This corollary, taken together with (EC.11), implies that

$$|z^f_{\mathbb{A}} - z^f_{\mathbb{B}}| \le 2L\cdot\min\big\{\xi_1(\mathbb{A},\mathbb{B}),\, \xi_2(\mathbb{A},\mathbb{B})\big\} \le 2L\cdot\xi^*(\mathbb{A},\mathbb{B}).$$

This proposition has a straighforward extension to the cases of $m>2$ groups and multidimensional covariates. The proofs have a similar structure to the case considered here, and thus are omitted.

□

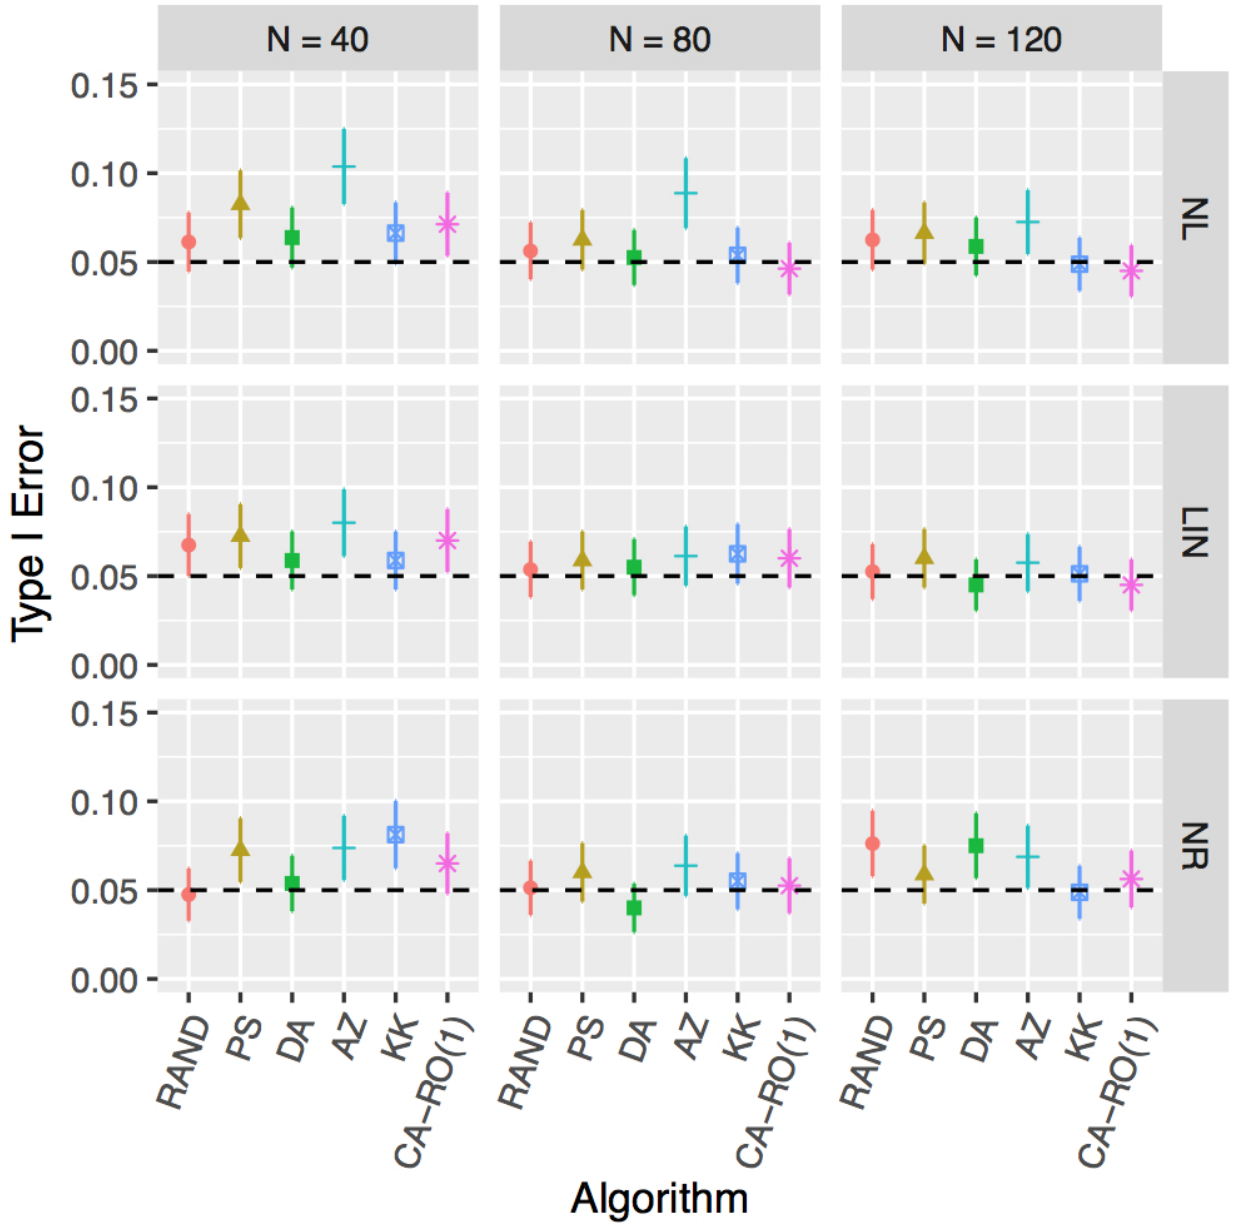

Figure EC.1 Type I error (with 95% confidence intervals) with CA-RO(1) vs. CA-RAND methods for $N \in \{40, 80, 120\}$ under various response models (NL, LIN, NR), using the adjusted treatment effect estimator. Dashed line indicates 0.05 significance level.

Table EC.1: Average absolute between-group discrepancy in moments under allocation algorithms ($m = 2, S = 1$).

		Moment						
N	Algorithm	1	2	3	4	5	$\log(\lvert w\rvert)$	$1/w$
20	RAND	0.358	0.498	1.321	2.955	8.108	0.387	13.521
	PS	0.260	0.509	1.120	2.872	7.389	0.524	13.150
	DA	0.167	0.616	0.931	3.172	6.953	0.719	13.492
	AZ	0.286	0.553	1.228	3.070	7.881	0.558	13.143
	KK	0.221	0.416	1.046	2.706	7.186	0.305	13.307
	CA-RO(1)	0.250	0.269	1.179	2.196	7.759	0.335	13.439
	CA-RO(3)	0.251	0.226	1.228	2.033	8.022	0.340	13.639
	CA-RO(5)	0.254	0.186	1.224	1.954	7.969	0.343	13.621
	OPT	0.024	0.010	0.960	1.517	7.348	0.354	13.702
60	RAND	0.205	0.292	0.793	1.869	5.396	0.224	10.020
	PS	0.125	0.250	0.625	1.640	4.746	0.257	9.850
	DA	0.092	0.350	0.560	1.935	4.700	0.408	9.990
	AZ	0.125	0.278	0.663	1.788	4.981	0.257	9.763
	KK	0.172	0.274	0.708	1.788	5.144	0.214	9.853
	CA-RO(1)	0.099	0.139	0.629	1.378	4.992	0.214	9.979
	CA-RO(3)	0.095	0.090	0.645	1.176	5.044	0.210	9.999
	CA-RO(5)	0.096	0.067	0.653	1.128	5.112	0.206	10.028
	OPT	0.001	3.33×10^{-4}	0.531	1.046	4.668	0.272	10.255
100	RAND	0.161	0.225	0.604	1.470	4.257	0.177	10.507
	PS	0.083	0.178	0.438	1.242	3.619	0.182	10.251
	DA	0.072	0.274	0.436	1.538	3.816	0.324	10.080

AZ	0.088	0.190	0.482	1.320	3.839	0.181	9.933
KK	0.133	0.218	0.544	1.449	4.107	0.171	10.195
CA-RO(1)	0.066	0.116	0.479	1.132	4.034	0.168	10.223
CA-RO(3)	0.063	0.073	0.488	0.984	4.091	0.165	10.368
CA-RO(5)	0.064	0.051	0.485	0.901	4.041	0.163	10.551
OPT	0.001	1.14×10^{-4}	0.402	0.806	3.692	0.218	10.240