

NONPARAMETRIC COVARIATE-ADJUSTED RESPONSE-ADAPTIVE DESIGN BASED ON A FUNCTIONAL URN MODEL

BY GIACOMO ALETTI^{*,†}, ANDREA GHIGLIETTI[‡] AND
WILLIAM F. ROSENBERGER[§]

ADAMSS Center^{} and Università degli Studi di Milano,[†]
Università Cattolica del Sacro Cuore[‡]
and George Mason University[§]*

In this paper, we propose a general class of covariate-adjusted response-adaptive (CARA) designs based on a new functional urn model. We prove strong consistency concerning the functional urn proportion and the proportion of subjects assigned to the treatment groups, in the whole study and for each covariate profile, allowing the distribution of the responses conditioned on covariates to be estimated nonparametrically. In addition, we establish joint central limit theorems for the above quantities and the sufficient statistics of features of interest, which allow to construct procedures to make inference on the conditional response distributions. These results are then applied to typical situations concerning Gaussian and binary responses.

1. Introduction.

1.1. *Adaptive designs in clinical trials.* The scientific validation of new treatments or therapies in medicine is typically the result of careful controlled randomized clinical trials. The standard methodology is to sequentially assign the patients to the treatment groups and collect the corresponding responses for statistical inference. The design of such experiments typically involves several aspects, such as ethical objectives, reduction of costs and inferential properties. Initially, randomized adaptive procedures were used to increase the balance among the treatment groups in order to achieve an unbiased comparison (e.g., see [32, 34]). However, balance does not ensure efficiency or good ethical properties, except in very particular circumstances (see [35]). Hence, new procedures, called response-adaptive, have been considered that use the accrued information on previous subjects' response to treatments to skew the probabilities of assignment away from 1/2 toward specific target values. An exhaustive review on response-adaptive procedures can be found for instance in [8, 23]. These designs are typically constructed to satisfy certain optimality criteria related to their performance, with respect to ethical

Received June 2017; revised November 2017.

MSC2010 subject classifications. Primary 62L20, 60F05, 62E20, 62L05; secondary 62F12, 62P10.

Key words and phrases. Clinical trials, covariate-adjusted analysis, inference, large sample theory, personalized medicine, randomization.

aspects, such as minimizing the expected number of failures, or statistical properties, such as maximizing the power of the test. Specifically, a standard methodology is to consider a desired asymptotic allocation proportion ρ depending on the response distributions that satisfies those optimality criteria. In a common situation, the response distributions depend on a vector of parameters $\theta \in \Theta$, and the target allocation is defined as a function $\rho : \Theta \rightarrow \mathcal{S}$, where \mathcal{S} is the simplex of dimension equal to the number of treatments. Thus, the response-adaptive design is constructed such that the proportion of subjects assigned to the treatments asymptotically target the desired target proportion, that is, $N_n/n \xrightarrow{\text{a.s.}} \rho(\theta)$.

When the information on significant covariates is available, adaptive designs based solely on the patient's response to treatments are inadequate to implement the randomized assignments. For instance, when a covariate has a strong influence on the response to a treatment, it may be inappropriate to use responses observed by subjects with a specific covariate profile to determine patients' allocation probability with a different profile. As described in [35], there are multiple ways for taking into account the effect of covariates in clinical trials and, in general, there is no agreement about how to implement such designs and what should be the main purpose of these procedures.

1.2. *Covariate-adjusted randomization designs.* A natural way to incorporate covariates in randomized procedures is to adopt stratification to force balance on certain important covariates. In fact, although randomization reduces the probability that the presence of the covariates are strongly different in the treatment groups (see [34]), perfect balance is reached only asymptotically and for small samples a significant imbalance can easily occur. Hence, a standard methodology is to stratify on few important known covariates, then to use restricted randomization within each stratum, and finally to let randomization handle the less influential and unknown covariates. Since this approach is possible only for a small number of covariates, several procedures have been proposed to individualize which covariates should be considered in the trial and how to implement the sequential assignments. An early work on these covariate-adjusted randomization designs is represented by the deterministic procedure proposed in [39] to minimize imbalances on strata and its extensions that include randomization: the biased coin design of [31] and the marginal urn design formulated by [42]. In this context, different types of imbalances have been controlled in the covariate-adaptive designs proposed in [25], by using the positive recurrence of the Markov process of the within-stratum imbalances. Moreover, a wide class of covariate-adaptive designs aimed at balancing the allocations has been recently presented in [13], in which the main asymptotic properties have been established. However, as shown in [35], balance does not guarantee either to have a design with good statistical properties or to assign more patients to the superior treatment. An alternative approach was proposed in [6] to incorporate treatment-by-covariate interactions and continuous covariates. Specifically, the subjects are assigned to treatment groups in

order to minimize the variance of the treatment responses conditionally on their covariate profiles. This procedure is based on the D_A -optimality criterion with linear models and, unlike the previous designs, it performs well also for correlated covariates. However, the minimum variance of the treatment responses can be obtained through balancing covariates in marginal strata only in linear models with homoscedastic errors, but it is not valid in nonlinear models (see, e.g., logistic regression [28]) or when errors present a general covariance structure. Hence, information on the covariates may not be sufficient to achieve efficiency or ethical goals, and using the responses to treatments in the allocation phase can be essential to this task. For this reason, adaptive procedures that use either the covariate profile of the subjects and the performances of the treatments have been considered in literature under the name covariate-adjusted response-adaptive (CARA) designs (see [36]).

1.3. *CARA designs.* In general, a CARA design is defined as a procedure that randomly assigns subjects to the treatment groups with a probability that depends on their own covariate profiles and on the previous patients' covariates, allocations and responses. The literature on this class of designs is not very long and early steps in this context can be found in [27] and [36].

Concerning binary responses and polytomous covariates, a relevant CARA design based on the randomized play-the-winner rule has been proposed in [14]. The case of binary responses and two competing treatments has been also considered in [36] for different types of covariates. The allocation rule proposed in [36] uses a suitable mapping based on a logistic regression model that implements the interaction among covariates and responses to treatments. Specifically, each subject is assigned according to the odds ratio comparing treatments in correspondence of his own covariate profile. The main properties of this adaptive allocation rule has been investigated through simulation, highlighting a significant reduction of the expected treatment failures. Nevertheless, theoretical results on the design performances have not been derived in [36]. A two-stage CARA design for binary responses based on logistic regression model was implemented in [16], in which the patients are assigned initially using a restricted procedure to compute the adaptive parameter estimators and then using a probability that depends on such estimators and the corresponding covariate profile.

The case of continuous responses and two treatments has been considered in [15], in which an adaptive design with limiting allocation proportion has been proposed by using a linear model to incorporate covariate information. However, the probability of assigning the next patient does not depend on its covariate profile, and hence the design proposed in [15] cannot be included in the classical CARA framework. The procedure in [15] has been improved in [7] by considering an adaptive biased-coin design for normal responses based on a generalized D_A -optimal criterion that takes into account both statistical and ethical purposes.

Although the performances of these allocation rules have been studied through simulation, their theoretical properties have not been proved.

Ideally, the analysis of the ethical and inferential properties of the experimental designs should be based on theoretical results concerning the asymptotic behavior of the allocation proportion and adaptive estimators, and none of the previous work on CARA designs is able to provide such results. In fact, since the allocation and the estimation process depend on both the responses and the covariates, CARA designs are very complex to be formulated in a rigorous mathematical setting. Two papers, in particular, formalize CARA in a rigorous mathematical framework. The first of these is the groundbreaking paper of [48], in which consistency and second-order asymptotic results concerning both adaptive estimators and allocation proportions have been proved for a very wide class of CARA designs. In the second [12], compound optimal design theory was used to find target allocations of interest, and these target allocations are attained using an accelerated biased coin design.

The procedures considered in [48] can be interpreted as generalized biased coin designs in which the probability of allocation is given by a known target function evaluated at the adaptive estimators of a finite number of parameters related to the responses means conditioned on the covariates. For this reason, this design has been applied to generalized linear models for different types of responses and covariates (discrete and continuous) and for more than two treatments. Some recent papers have extended the class of CARA designs presented in [48] in order to improve its inferential properties. For instance, the class of designs proposed in [50] allows inference also for common parameters in the response distributions. The unified family of designs presented in [24] takes into account both efficiency and medical ethics. The distribution of the parameters estimators in reduced generalized linear models established in [19] allows inference for separately testing the main effects, the covariate effects or their intersections. One of the aims of this paper is to construct a new framework for CARA designs in which the probability of allocation may depend by nonparametric or semiparametric estimates of the generic conditional response distribution. To this purpose, the proposed design is based on the other very popular class of randomized procedures: urn models.

1.4. Urn models. The history of urn models as probabilistic tools to describe random phenomena is very long and deep in several fields of scientific research (e.g., [17] in economics, [18] in genetics, [1, 2] in network analysis). Starting from the Pólya urn proposed in [21] to model contagious disease, many variants have been considered (see, e.g., [3, 5, 11, 22, 30, 47, 49]). In this paper, we focus on the broad class of urn schemes known as generalized Pólya urn (GPU); its asymptotic behavior has been the objects of several important works: starting from the asymptotic results proved in [5] by embedding the urn process in a continuous-time branching process, other significant theoretical results have been derived, for instance in [9–11, 26, 38, 47, 49]. Concerning applications, urn models have known

a great popularity as adaptive designs in clinical trials for several reasons: (i) the process of colors sampled from the urn represents a natural way to model the sequential randomized assignments of subjects to treatments groups, (ii) the quantity of balls replaced in the urn may depend on previous patients' information so that different types of adaptive procedures can be constructed, (iii) their updating rule composed by extractions and replacements is intuitive and easy to be implemented by clinicians. In clinical trials, urn models are mostly adopted as response-adaptive designs and covariates are not considered. For a review of urn models adopted as response-adaptive procedures, see [20, 23, 30].

The response-adaptive GPU design presented in [47] is the following: consider an urn containing balls of d colors, each one associated with a specific treatment; the composition of the urn at time $n \geq 0$ is represented by the vector $\mathbb{Y}_n = (\mathbb{Y}_n^1, \dots, \mathbb{Y}_n^d)^\top \in \mathbb{R}_+^d$. At any time $n \geq 1$, a subject enters the trial, a ball is sampled at random from the urn and its color is observed; formally, let $\mathbb{X}_n \in \{0, 1\}^d \cap \mathcal{S}$ represent the color sampled at time n : for each $j \in \{1, \dots, d\}$, $\mathbb{X}_n^j = 1$ indicates that the sampled ball is of color j , $\mathbb{X}_n^j = 0$ otherwise. Then, letting $k \in \{1, \dots, d\}$ be the sampled color, the subject is assigned to treatment k and a response ξ_n^k is observed. The potential responses $\{\xi_n^k; n \geq 1\}$, $k \in \{1, \dots, d\}$, are defined as d independent sequences of identically distributed (i.i.d.) random variables. Moreover, for any $i, j \in \{1, \dots, d\}$, the model uses a function $u^{ij} : S^j \rightarrow \mathbb{R}_+$ (S^j indicates the support of ξ_n^j) to convert the responses into the reinforcements; then the sampled ball of color k is returned to the urn together with $u^{ik}(\xi_n^k)$ new balls of color $i \in \{1, \dots, d\}$. Formally, denoting \mathbb{D}_n the $d \times d$ replacement matrix defined as $\mathbb{D}_n^{ij} := u^{ij}(\xi_n^j)$ for any $i, j \in \{1, \dots, d\}$, the urn is updated as follows:

$$(1.1) \quad \mathbb{Y}_n = \mathbb{Y}_{n-1} + \mathbb{D}_n \mathbb{X}_n.$$

Note that \mathbb{D}_n^{ik} can be computed only if the subject receives treatment k , because ξ_n^k is needed to obtain $\mathbb{D}_n^{ik} = u^{ik}(\xi_n^k)$. However, the updating rule expressed in (1.1) does not require the value of \mathbb{D}_n^{ik} if the subject receives a treatment different than k , since in that case $\mathbb{X}_n^k = 0$.

In the GPU model, it is well known that the asymptotic behavior of the urn proportion is related to the limit of the conditional expectation of the replacement matrices, that is, $\mathbb{H} := \lim_{n \rightarrow \infty} \mathbb{E}[\mathbb{D}_n | \mathcal{F}_{n-1}]$ a.s. (e.g., [5, 9–11, 47, 49]), where \mathcal{F}_{n-1} indicates the quantities observed up to time $(n - 1)$. Specifically, under suitable conditions, we have that $\mathbb{Y}_n/w(\mathbb{Y}_n) \xrightarrow{\text{a.s.}} \mathbb{V}$, where $w(\mathbb{Y}_n) = \sum_{j=1}^d \mathbb{Y}_n^j$ and $\mathbb{V} \in \mathcal{S}$ is the right eigenvector of \mathbb{H} associated with its maximum eigenvalue and such that $w(\mathbb{V}) = 1$. Hence, in order to target a specific proportion $\rho(\theta)$, one needs to define appropriately the replacement matrices \mathbb{D}_n such that \mathbb{H} guarantees $\rho(\theta) \equiv \mathbb{V}$. To this end, the functions u^{ij} should depend on the parameters θ , which are usually unknown in practice. Hence, in [47] the parameter θ is replaced in u^{ij} by the adaptive estimator $\hat{\theta}_{n-1}$ computed with the information available up to time $(n - 1)$ obtaining \hat{u}^{ij} , so that the replacement matrix is represented by $\mathbb{D}_n^{ij} = \hat{u}^{ij}(\xi_n^j)$ for

$i, j \in \{1, \dots, d\}$. This model proposed in [47] is called a *sequential estimation-adjusted urn* model (SEU) and it has been proved that the model targets any desired limiting proportion $\rho(\theta)$.

1.5. *Aim and organization of the paper.* Although the SEU model represents a very powerful urn design, as with most response-adaptive procedures, it does not incorporate the covariates in the randomization process. Hence, the probability of allocating a subject to a treatment group is independent of his or her covariate profile; moreover, all patients are asymptotically assigned with the same target probability $\rho(\theta)$, regardless of their covariate profiles. In this paper, we want to extend the SEU model by introducing information on the covariates in the urn scheme, so obtaining a CARA urn design. Analogously to the CARA coin designs described in [48], where the covariates are considered in the trial, each subject is assigned with a probability that depends on his or her own covariate profile. Formally, let τ be the covariate space, which could be finite, countable or continuous. Conditionally on the covariate profile $t \in \tau$, we consider different response distributions π_t^1, \dots, π_t^d and the corresponding parameters θ_t . Thus, in this framework the main goal of the design is to asymptotically assign all the subjects with covariate profile t with a desired probability $\rho(\theta_t)$, that is, the design targets the desired functional allocation $\rho(t)$. We will see that the model presented in this paper achieves this goal. It is worth noticing that in the design proposed in [48], θ_t characterizes the means of the responses conditionally on t and it is defined as $\theta_t = f(\beta, t)$, where f is a known function and β is a finite set of parameters. As a consequence, in [48] the target function $\rho(t)$ actually depends on a finite number of unknowns represented by β . The framework considered in this paper is different since θ_t may represent general features of the conditional response distributions and it can be estimated nonparametrically; the target function $\rho(t)$ depends on an infinite number of unknowns. Another difference is that in the existing literature the probability distribution of the covariates is the same for all the patients, while in this paper this distribution is allowed to be adaptively modified by the experimenter using the information collected during the trial.

The CARA design we propose consists of a functional urn model in which the urn composition is a d -dimensional multivariate function of the covariates. Each subject is assigned to the treatment group according to the color sampled from the urn identified by his or her covariate profile. After any allocation, the entire functional urn composition is updated, even if only the response associated with the patient's covariate profile has been observed. For this reason, a crucial point is the definition of the functional objects \mathbf{X}_n and D_n that extend the multivariate objects \mathbb{X}_n and \mathbb{D}_n in the updating rule (1.1).

In this paper, we establish first- and second-order asymptotic results concerning the following quantities:

- (i) the probability of allocation of the subjects for each covariate profile, $\mathbf{Z}_n(t)$;

(ii) the proportion of subjects assigned to the treatment groups for each covariate profile, $\mathbf{N}_{t,n}/w(\mathbf{N}_{t,n})$;

(iii) the proportion of subjects assigned to the treatment groups in the trial, \mathbf{N}_n/n ;

(iv) the adaptive estimators $\hat{\boldsymbol{\theta}}_{t,n}$ of features of interest $\boldsymbol{\theta}_t$ related with the distribution of the treatments responses conditionally on each covariate profiles (required for the inference based on covariate-stratification approach);

(v) the adaptive estimators $\hat{\boldsymbol{\beta}}_n$ of features of interest $\boldsymbol{\beta}$ related with the entire family of response distributions conditionally on the covariates (required for the inference based on covariate-adjusted approach).

In particular, we prove strong consistency of the above quantities, allowing the distribution of the responses conditioned on the covariates to be estimated non-parametrically. In addition, we establish joint central limit theorems (CLTs) which provide the essential probabilistic tools to construct inferential procedures on conditional response distributions. These results are then applied to typical situations concerning Gaussian and binary responses.

In Section 2, we present the CARA urn design based on the functional urn model. Section 3 is concerned with assumptions and main results. In Sections 4 and 5, the CARA urn design is applied to different practical scenarios and several response distributions are considered. Concluding remarks and future developments are discussed in Section 6. Technical details, including proofs of the theorems, are presented in Supplementary Material [4].

2. The functional urn model for CARA designs. In this section, we describe the design based on a functional urn model. We start by defining the quantities related to the process of subjects that sequentially enter the trial.

2.1. Notation. Consider a trial in which patients are sequentially and randomly assigned to $d \geq 2$ treatments. For any $n \geq 1$, let $\bar{\mathbf{X}}_n \in \{0, 1\}^d \cap \mathcal{S}$ represent the treatment assigned to subject n : for any $j \in \{1, \dots, d\}$, $\bar{X}_n^j = 1$ indicates that the assigned treatment is j , $\bar{X}_n^j = 0$ otherwise. Each subject n is identified by a vector $(T_n, \xi_n^1, \dots, \xi_n^d) \in (\tau \times S^1 \times \dots \times S^d)$, in which T_n indicates his or her covariate profile and, for any $j \in \{1, \dots, d\}$, ξ_n^j indicates the patient's potential response to treatment j . In general, the covariate profiles of the patients $\{T_n; n \geq 1\}$ are a sequence of independent but nonidentically distributed random variables, whose distributions can adaptively depend on the information collected during the trial: the covariate profiles, the allocations and the responses. This allows a generalization of the typical CARA design, by incorporating a probabilistic mechanism to select patients with a particular covariate profile if patient selection is also a random process. This includes as a special case the typical assumption of $\{T_n; n \geq 1\}$ as i.i.d. random variables, in which the information collected in the trial does not

affect the choice of the future covariate profiles. This represents the standard clinical trial where the clinician has no control over the patient recruitment process. We will denote by μ_{n-1} the probability distribution of T_n conditioned on \mathcal{F}_{n-1} , that is, the information collected up to time $(n - 1)$.

Since any subject receives one treatment, for any $n \geq 1$ only one value among $\{\xi_n^1, \dots, \xi_n^d\}$ can be observed during the trial and we will denote it $\bar{\xi}_n$. For each $j \in \{1, \dots, d\}$, we assume that $\{\xi_n^j; n \geq 1\}$ is a sequence of independent random variables whose distribution depends on the sequence of covariate profiles $\{T_n; n \geq 1\}$. Specifically, we define a family of probability distributions $\{\pi_t^j; t \in \tau\}$, where each one represents the distribution of ξ_n^j conditioned on the event $\{T_n = t\}$: π_t^j indicates the probability law of the response to treatment j observed from a subject whose covariate profile is equal to t . Note that $\{T_n = t\}$ could have measure zero. In the paper, we will also use the corresponding families of cumulative distribution functions (CDFs) $\{F_t^j; t \in \tau\}$ and the families of quantile functions (QFs) $\{Q_t^j; t \in \tau\}$.

2.2. *The model.* For any $n \geq 0$, let $\mathbf{Y}_n = (Y_n^1, \dots, Y_n^d)^\top$ be a d -dimensional vector of nonnegative bounded functions and let $\mathbf{Z}_n = \mathbf{Y}_n/w(\mathbf{Y}_n)$, where $w(\mathbf{Y}_n) := \sum_{j=1}^d Y_n^j$. For any $t \in \tau$, $\mathbf{Y}_n(t) \in (0, 1)^d$ represents an urn containing $Y_n^j(t)$ balls of color $j \in \{1, \dots, d\}$ and $\mathbf{Z}_n(t) \in \mathcal{S}$ indicates the proportions of the colors in the urn at time n . To avoid inessential complications, we consider a uniform initial composition $\mathbf{Y}_0 = \mathbf{1}$. For any $n \geq 1$, let \mathcal{F}_{n-1} be the σ -algebra composed by the information related with the first $(n - 1)$ patients, that is, their covariate profiles, allocations and responses:

$$(2.1) \quad \mathcal{F}_{n-1} := \sigma(T_1, \bar{\mathbf{X}}_1, \bar{\xi}_1, \dots, T_{n-1}, \bar{\mathbf{X}}_{n-1}, \bar{\xi}_{n-1}).$$

When subject n enters the trial, his or her covariate profile T_n is observed, and we now operate on the conditioning set $\{\mathcal{F}_{n-1}, T_n\}$, consistent with the definition of CARA designs found in [23]. Then a ball is sampled at random from the urn identified by T_n ; that is, with proportions

$$\mathbf{Z}_{n-1}(T_n) = (Z_{n-1}^1(T_n), \dots, Z_{n-1}^d(T_n))^\top$$

and its color is observed; thus, the subject n receives the treatment associated with the sampled color and a response $\bar{\xi}_n$ is collected. In order to update the functional urn, we construct a weighting function \mathbf{X}_n and a functional replacement matrix D_n that extend \mathbb{X}_n and \mathbb{D}_n in the classical updating rule; see (1.1).

First, we define the weighting function \mathbf{X}_n . Let U_n be a uniform $(0, 1)$ random variable independent of \mathcal{F}_{n-1} and T_n , and define, for any $t \in \tau$ and for any $j \in \{1, \dots, d\}$

$$(2.2) \quad \check{X}_n^j(t) := \mathbb{1}_{\{\sum_{i=1}^{j-1} Z_{n-1}^i(t) < U_n \leq \sum_{i=1}^j Z_{n-1}^i(t)\}},$$

where we use the convention $\sum_{i=1}^0(\cdot) = 0$. Notice that $\check{\mathbf{X}}_n(t)$ represents the color of the ball that would be sampled if we used the urn identified by t ; that is, if the covariate profile T_n were equal to t . Thus, $\check{\mathbf{X}}_n := \check{\mathbf{X}}_n(T_n)$ represents the color actually sampled from the functional urn at time n : for each $j \in \{1, \dots, d\}$, $\check{X}_n^j = 1$ indicates that the sampled ball is of color j , $\check{X}_n^j = 0$ otherwise. Hence, conditionally on \mathcal{F}_{n-1} and T_n , \check{X}_n^j is Bernoulli distributed with parameter $Z_{n-1}^j(T_n)$. Since $\check{\mathbf{X}}_n(t)$ models the color hypothetically sampled from the urn identified by t [i.e., $\mathbf{Y}_{n-1}(t)$], we define the weighting function \mathbf{X}_n as the expected value of $\check{\mathbf{X}}_n$ conditioned on the information of the color sampled from the urn identified by T_n , that is,

$$(2.3) \quad \mathbf{X}_n := \mathbb{E}[\check{\mathbf{X}}_n | \mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n].$$

An analytic expression of \mathbf{X}_n derived from (2.3) is provided in (A.1) in the Supplementary Material [4]. Note that by (2.2) and (2.3) we have $\mathbf{X}_n(t) \in \mathcal{S}$ for any $t \in \tau$, since $\mathbf{X}_n(t) \in [0, 1]^d$ and

$$w(\mathbf{X}_n) = w(\mathbb{E}[\check{\mathbf{X}}_n | \mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n]) = \mathbb{E}[w(\check{\mathbf{X}}_n) | \mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n] = 1.$$

Moreover, we also have that

$$(2.4) \quad \mathbb{E}[\mathbf{X}_n | \mathcal{F}_{n-1}] = \mathbf{Z}_n,$$

since by the law of total expectation $\mathbb{E}[\mathbf{X}_n | \mathcal{F}_{n-1}] = \mathbb{E}[\check{\mathbf{X}}_n | \mathcal{F}_{n-1}] = \mathbf{Z}_{n-1}$.

We now define the functional replacement matrix D_n . First, for any $i, j \in \{1, \dots, d\}$ and $t \in \tau$, let $u_t^{ij} : \mathcal{S}^j \rightarrow \mathbb{R}_+$ be a function that converts the responses into the reinforcements. Analogously to the classical updating rule (1.1), for any $t \in \tau$, the urn identified by t should be ideally updated by $u_t^{ij}(\xi_n^j)$ balls of color $i \in \{1, \dots, d\}$ when treatment j is assigned, where ξ_n^j represents the response observed from a subject with covariate profile t . Hence, the urn identified by t should be updated by $u_t^{ij}(W_t^j)$ balls of color $i \in \{1, \dots, d\}$, where W_t^j represents a random variable with probability distribution π_t^j and that, conditionally on \mathcal{F}_{n-1} and T_n , is independent of $\bar{\mathbf{X}}_n$. This can be equivalently formalized by introducing a uniform $(0, 1)$ random variable V_n independent of U_n, T_n and \mathcal{F}_{n-1} , and defining $W_t^j = Q_t^j(V_n)$ for any $j \in \{1, \dots, d\}$, where we recall that Q_t^j is the QF associated with the probability distribution π_t^j . In fact, by definition, we have that $Q_t^j(V_n) \sim \pi_t^j$ when $V_n \sim U(0, 1)$. Thus, the replacements in the urn identified by $t \in \tau$ should be defined by the following random matrix:

$$(2.5) \quad \check{D}_n^{ij}(t) := u_t^{ij}(Q_t^j(V_n)) \quad \forall i, j \in \{1, \dots, d\}.$$

However, when the subject n with covariate profile T_n is assigned to a treatment j and the response ξ_n^j is observed, we can only compute $\bar{D}_n^{ij} := u_{T_n}^{ij}(\xi_n^j)$ that corresponds to $\check{D}_n^{ij}(T_n)$. Nevertheless, the response ξ_n^j , associated with the covariate

profile T_n , contains the information on the quantile V_n that can be taken into account to update all the urns $t \in \tau$ using (2.5). Specifically, the $d \times d$ replacement matrix of bounded functions should be defined as the expected value of the potential replacement matrix \check{D}_n for the urn t , conditionally on the information of the response observed to treatment $k \in \{1, \dots, d\}$, that we call $\bar{\xi}_n = \xi_n^k$, from a subject with covariate profile T_n , that is,

$$(2.6) \quad D_n^* := \mathbb{E}[\check{D}_n | T_n, \bar{\mathbf{X}}_n, \bar{\xi}_n].$$

An explicit expression of (2.6) can be derived as follows: for any $s \in \tau$, $k \in \{1, \dots, d\}$ and $y \in S^k$,

$$(2.7) \quad \begin{aligned} D_n^{*ij} &= \mathbb{E}[\check{D}_n^{ij} | \{T_n = s\}, \{\bar{X}_n^k = 1\}, \{\bar{\xi}_n = y\}] \\ &= \mathbb{E}[u_t^{ij}(Q_t^j(V_n)) | \{V_n \in (Q_s^k)^{-1}(y)\}], \end{aligned}$$

where

$$(Q_s^k)^{-1}(y) := \{v \in (0, 1) : Q_s^k(v) = y\},$$

and we recall that Q_s^k is the QF associated with the probability distribution π_s^k . Note from (2.6) that D_n^* depends on quantities that are unknown at time n . Specifically, the expression in (2.6) contains the conditional QFs Q_t^j and Q_s^k ; moreover, as mentioned in the Introduction, u_t^{ij} typically depends on the response distribution π_t^j in order to obtain some desired asymptotic properties from the design. Hence, since the conditional distributions π_t^j are typically unknown, we compute the corresponding functional estimators \hat{u}_t^{ij} , \hat{Q}_t^j and \hat{Q}_s^k by using the information related with the first $(n - 1)$ subjects. Thus, the $d \times d$ replacement matrix D_n of bounded functions is defined, on the sets $\{T_n = s\}, \{\bar{X}_n^k = 1\}, \{\xi_n^k = y\}$, as follows:

$$(2.8) \quad D_n^{ij} := \mathbb{E}[\hat{u}_t^{ij}(\hat{Q}_t^j(V_n)) | \{V_n \in (\hat{Q}_s^k)^{-1}(y)\}].$$

Note that $D_n(T_n) = \bar{D}_n$. The analytic expression of D_n depends on the specific family of probability distribution $\{\pi_t^j; t \in \tau\}$, $j \in \{1, \dots, d\}$, that models the relation among the response ξ_n^j and the covariate profile T_n . See Section A in the Supplementary Material [4].

Summarizing, for any $t \in \tau$ and $i \in \{1, \dots, d\}$, we replace in the urn identified by t a number of balls of color i equal to $\sum_{j=1}^d D_n^{ij}(t) X_n^j(t)$. Hence, for any $n \geq 1$ the functional urn is updated as follows:

$$(2.9) \quad \mathbf{Y}_n = \mathbf{Y}_{n-1} + D_n \mathbf{X}_n,$$

and we set $\mathbf{Z}_n = \mathbf{Y}_n/w(\mathbf{Y}_n)$. Finally, we define the σ -algebra \mathcal{F}_n generated by the quantities related with the first n subjects:

$$\mathcal{F}_n := \sigma(\mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n, \bar{\xi}_n),$$

and we compute the probability distribution of the covariate profile of the next patient $\mu_n(dt) := \mathbb{P}(T_{n+1} \in dt | \mathcal{F}_n)$ with the information in \mathcal{F}_n .

The key feature of the design is that quantile functions are used to update *all* urns, not just the urn for which $T_n = t$. In theory, there could be an uncountably infinite number of urns, with only a finite subset of them used for patient allocation. However, in clinical practice, mathematically “continuous” covariates are really not continuous [34]; for instance, cholesterol is represented by integer values, likely in some range that would for all intents and purposes, make it a finite discrete covariate. However, the procedure is well defined for uncountably infinite urns, and first-order asymptotic properties can be obtained, although some of the covariate-specific metrics do not make sense in that context. When we move to second-order asymptotics, we partition τ into K strata, which could be intervals of a continuous set.

REMARK 2.1. Suppose π_t^1, \dots, π_t^d are known, then D_n could be replaced by D_n^* , which does not depend on \mathcal{F}_{n-1} . In that case, the distribution of \mathbf{Y}_n , conditionally on \mathcal{F}_{n-1} , depends only on \mathbf{Y}_{n-1} : the functional urn composition $\{\mathbf{Y}_n; n \geq 1\}$ is a Markov process. However, for any $t_0 \in \tau$, the distribution of the real random variable $\mathbf{Y}_n(t_0)$, conditionally on \mathcal{F}_{n-1} , depends on all the quantities contained in \mathcal{F}_{n-1} given in (2.1) and not only on $\mathbf{Y}_{n-1}(t_0)$. Hence, for any $t_0 \in \tau$, the real-valued sequence of the urn composition $\{\mathbf{Y}_n(t_0); n \geq 1\}$ is not a Markov process. This emphasizes our choice of a functional urn model.

3. Assumptions and main results. This section is concerned with the assumptions and the main results of the design described in Section 2. Specifically, we are interested in the asymptotic behavior of the following processes:

(i) the probability of allocation of the subjects for each covariate profile: $\{\mathbf{Z}_n(t); t \in \tau\}$;

(ii) the proportion of subjects associated with each covariate profile assigned to the treatments: $\{\mathbf{N}_{t,n}/w(\mathbf{N}_{t,n}); t \in \tau\}$, where $\mathbf{N}_{t,n} := \sum_{i=1}^n \bar{\mathbf{X}}_i \mathbb{1}_{\{T_i=t\}}$;

(iii) the proportion of subjects assigned to the treatments: \mathbf{N}_n/n , where $\mathbf{N}_n := \sum_{i=1}^n \bar{\mathbf{X}}_i$;

(iv) the adaptive estimators of features of interest related with the distribution of the treatments responses conditionally on each covariate profile:

$$\{\hat{\theta}_{t,n}; t \in \tau\} := \{(\hat{\theta}_{t,n}^j, j \in \{1, \dots, d\})^\top; t \in \tau\},$$

where each estimator $\hat{\theta}_{t,n}^j$ is computed with the responses of the first $N_{t,n}^j$ subjects assigned to treatment j with covariate profile t , that is, $\{1 \leq i \leq n : \bar{X}_i^j \mathbb{1}_{\{T_i=t\}} = 1\}$;

(v) the adaptive estimators of features of interest related with the entire family of distribution of the treatments responses conditionally on the covariates:

$$\hat{\beta}_n := (\hat{\beta}_n^j, j \in \{1, \dots, d\})^\top;$$

each estimator $\hat{\beta}_n^j$ is now computed with the responses of the first N_n^j subjects assigned to treatment j , that is, $\{1 \leq i \leq n : \bar{X}_i^j = 1\}$.

REMARK 3.1. In the case where τ is continuous, metrics in (ii) and (iv) have no meaning.

3.1. *First-order asymptotic properties.*

3.1.1. *Assumptions.* We start by providing the main assumptions that are required for establishing the first-order asymptotic properties.

(A1) *Constant balance and positiveness of replacement matrices.* Let $D_n^{ij}(t) > 0$ for any $i, j \in \{1, \dots, d\}$ and $t \in \tau$, which is equivalent to require that $u_t^{ij}(y) > 0$ for any $y \in S^j$. Moreover, denoting $\check{D}^{\cdot j}(t)$ the j th column of $\check{D}(t)$, we require that there exists a function $c(t)$ such that $\inf_{t \in \tau} c(t) > 0$ and for any $t \in \tau$

$$(3.1) \quad \mathbb{P}(w(\check{D}^{\cdot 1}(t)) = w(\check{D}^{\cdot 2}(t)) = \dots = w(\check{D}^{\cdot d}(t)) = c(t)) = 1.$$

Since by (2.5) $\check{D}_n^{ij}(t) = u_t^{ij}(W_t^j)$ with $W_t^j \sim \pi_t^j$, (3.1) holds when the function u_t^{ij} is chosen such that $\sum_{i=1}^d u_t^{ij}(W_t^j)$ is equal to $c(t)$ with probability one for all $j \in \{1, \dots, d\}$ and $t \in \tau$. To avoid unessential complications, without loss of generality we assume throughout all the paper that $c(t) = 1$.

(A2) *Limiting generating matrix.* For any $t \in \tau$, let $H(t) := \mathbb{E}[\check{D}_1(t)]$ and $H_n(t) := \mathbb{E}[D_n(t)|\mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n]$; then we assume that $H(t)$ is irreducible, diagonalizable and there exists $\alpha > 0$ independent of $t \in \tau$ such that

$$(3.2) \quad \mathbb{E}[|H_n(t) - H(t)||\mathcal{F}_{n-1}] = O(n^{-\alpha}) \quad \text{a.s.}$$

We will refer to H_n as *generating matrix* and to H as *limiting generating matrix*.

A simple interpretation of (3.2) can be obtained by noticing that H can be expressed as $\mathbb{E}[D_n^*|\mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n]$ for all $n \geq 1$, from which it follows that the assumption in (3.2) is related to the properties of consistency of the adaptive estimators. To see that $H \equiv \mathbb{E}[D_n^*|\mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n]$, observe that $\mathbb{E}[D_n^*|\mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n] = \mathbb{E}[\check{D}_n|\mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n]$ and by (2.5) $\check{D}_n^{ij}(t) = u_t^{ij}(Q_t^j(V_n))$, with $V_n \sim U(0, 1)$ independent of \mathcal{F}_{n-1}, T_n and $\bar{\mathbf{X}}_n$, which implies

$$\mathbb{E}[\check{D}_n^{ij}(t)|\mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n] = \mathbb{E}[u_t^{ij}(Q_t^j(V_n))] = H^{ij}(t).$$

REMARK 3.2. Assumption (A2) ensures that the conditional increments of the urn composition $\mathbb{E}[\mathbf{Y}_n - \mathbf{Y}_{n-1}|\mathcal{F}_{n-1}]$ are asymptotically equal to $H\mathbf{Z}_{n-1}$. Indeed, combining (2.4) and (2.9), we have

$$\mathbb{E}[\mathbf{Y}_n - \mathbf{Y}_{n-1}|\mathcal{F}_{n-1}] = \mathbb{E}[H_n\mathbf{X}_n|\mathcal{F}_{n-1}] = O(n^{-\alpha}) + H\mathbf{Z}_{n-1}.$$

From a probabilistic point of view, this is a key point to develop a functional urn asymptotic theory.

REMARK 3.3. Our functional urn model defines H_n as $\mathbb{E}[D_n|\mathcal{F}_{n-1}, T_n, \bar{\mathbf{X}}_n]$ instead of $\mathbb{E}[D_n|\mathcal{F}_{n-1}]$, typically given in the literature (e.g., see [9, 11, 29]). The two definitions coincide only when no functional dependence occurs. However, Assumption (A2) ensures that the limiting generating matrix is still the same:

$$\mathbb{E}[D_n|\mathcal{F}_{n-1}] = \mathbb{E}[H_n|\mathcal{F}_{n-1}] = H + \mathbb{E}[H_n - H|\mathcal{F}_{n-1}] \xrightarrow{\text{a.s.}} H.$$

REMARK 3.4. Assumption (A2) is verified under mild conditions on the consistency of the adaptive estimators in \hat{u}_t^{ij} and $\hat{\pi}_t^j$.

3.1.2. *First-order asymptotic results.* The main consistency results concerning the design are collected in the following theorem.

THEOREM 3.1. *Let $\mathbf{v}(t) \in S$ be the right eigenvector of $H(t)$ associated with $\lambda = 1$ and assume (A1), (A2). Then:*

(a) *for any probability measure ν on τ , we have*

$$\int_{\tau} \|\mathbf{Z}_n(t) - \mathbf{v}(t)\| \nu(dt) \xrightarrow{\text{a.s.}} 0,$$

and hence $\mathbf{Z}_n(t) \xrightarrow{\text{a.s.}} \mathbf{v}(t)$ for any $t \in \tau$;

(b) *for any $t \in \tau$ such that $\sum_{i=1}^n \mu_{i-1}(\{t\}) \xrightarrow{\text{a.s.}} \infty$, we have*

$$\|\mathbf{N}_{t,n}/w(\mathbf{N}_{t,n}) - \mathbf{v}(t)\| \xrightarrow{\text{a.s.}} 0;$$

(c) *if there exists a probability measure μ on τ such that $\int_{\tau} |\mu_n(dt) - \mu(dt)| \xrightarrow{\text{a.s.}} 0$, then we have that*

$$\left\| \mathbf{N}_n/n - \int_{\tau} \mathbf{v}(t)\mu(dt) \right\| \xrightarrow{\text{a.s.}} 0.$$

REMARK 3.5. In the special case that the covariate profiles of the subjects $\{T_n; n \geq 1\}$ is a sequence of i.i.d. random variables, we have $\mu_i = \mu$ for any $i \geq 0$, and hence the assumptions of (b) and (c) in Theorem 3.1 are immediately satisfied.

REMARK 3.6. Conditioning on any specific covariate profile, these results are consistent with well-known asymptotic results found in the literature on urn models (e.g., see [5, 9–11, 46, 47, 49]), in which the urn proportion and the allocation proportion converges a.s. to the normalized right eigenvector of the limiting irreducible mean replacement matrix associated with $\lambda = 1$.

3.2. *Second-order asymptotic properties.* The convergence results proved in Section 3.1 consider a general covariate space τ . In order to show second-order properties, we now partition τ into K finite elements, which could, for instance, be K intervals of a continuous covariate space. This partitioning induces K urns used to allocate subjects with covariate profiles in the set $\{1, \dots, K\}$. In clinical trials practice, K must be considerably smaller than the total sample size.

3.2.1. *Assumptions.* We now present further assumptions that are required for establishing the second-order asymptotic properties.

(A3) *Finite partition of the covariate space.* We assume that the covariate space τ is composed by a finite number $K \in \mathbb{N}$ of distinct elements. When τ contains infinite elements, we can take a partition of τ , that is, $\{\tau_1, \dots, \tau_K\}$ such that $\bigcup_k \tau_k = \tau$ and $\tau_{k_1} \cap \tau_{k_2} = \emptyset$ for $k_1 \neq k_2$, and consider these sets to be the elements of τ , that is, $\tau := \{\tau_1, \dots, \tau_K\}$. To facilitate the notation, without loss of generality in the sequel we redefine $\tau = \{1, \dots, K\}$ and $\mu_{n-1}(t) = \mu_{n-1}(\{t\}) = \mathbb{P}(T_n = t | \mathcal{F}_{n-1})$ for any $t \in \tau$.

(A4) *Conditional response distributions.* The analog of the null hypothesis in classical inferential statistics is given here by assuming that the conditional response distributions π_t^1, \dots, π_t^d are known for any $t \in \tau$. As a direct consequence, we have that $D_n = D_n^*$ and $H_n = H$ with probability one for any $n \geq 1$.

(A5) *Eigenvalues of the limiting generating matrix.* Denoting $\lambda_H^*(t)$ the eigenvalue of $\text{Sp}(H(t)) \setminus \{1\}$ with largest real part, assume that $\max_{t \in \tau} \text{Re}(\lambda_H^*(t)) < 1/2$.

(A6) *Dynamics of adaptive estimators.*

(A6a) (*Covariate-stratification approach*). For some $t \in \tau$ and $j \in \{1, \dots, d\}$, consider that there are features of interest θ_t^j related with the distribution π_t^j of the responses to treatment j conditionally on the covariate profile t . Then we assume that the corresponding adaptive estimator $\hat{\theta}_{t,n}^j$ is strongly consistent and its dynamics can be expressed as follows: there exists $n_0 \geq 1$ such that for any $n \geq n_0$,

$$(3.3) \quad \hat{\theta}_{t,n}^j - \hat{\theta}_{t,n-1}^j = -\frac{\bar{X}_n^j \mathbb{1}_{\{T_n=t\}}}{N_{t,n}^j} (f_{t,j}(\hat{\theta}_{t,n-1}^j) - \Delta \mathbf{M}_{t,j,n} - \mathbf{R}_{t,j,n}),$$

where:

- (i) $f_{t,j}$ is a Lipschitz continuous function such that $f_{t,j}(\theta_t^j) = 0$;
- (ii) $\Delta \mathbf{M}_{t,j,n} \in \mathcal{F}_n$ is a martingale increment such that $\mathbb{E}[\Delta \mathbf{M}_{t,j,n} | \mathcal{F}_{n-1}, T_n, \bar{X}_n^j] = 0$, and it converges stably to $\Delta \mathbf{M}_{t,j}$ with kernel K independent of \mathcal{F}_{n-1} : $\mathcal{L}(\Delta \mathbf{M}_{t,j,n} | \mathcal{F}_{n-1}, T_n = t, \bar{X}_n^j = 1) \xrightarrow{\text{a.s.}} K(t, j)$;
- (iii) $\mathbf{R}_{t,j,n} \in \mathcal{F}_n$ is such that $n \mathbb{E}[\|\mathbf{R}_{t,j,n}\|^2] \rightarrow 0$.

Moreover, let $f_{t,j}$ be differentiable at θ_t^j , denote by $\lambda_{\theta_t^j}^*$ the eigenvalue of $\text{Sp}(\mathcal{D}f_{t,j}(\theta_t^j))$ with largest real part and assume that $\min_{t \in \tau} \text{Re}(\lambda_{\theta_t^j}^*) > 1/2$. We also assume that for some $\delta > 0$,

$$(3.4) \quad \sup_{n \geq 1} \mathbb{E}[\|\Delta y \mathbf{M}_{t,j,n}\|^{2+\delta} | \mathcal{F}_{n-1}] < +\infty \quad \text{a.s.}$$

and

$$(3.5) \quad \mathbb{E}[\Delta \mathbf{M}_{t,j,n} (\Delta \mathbf{M}_{t,j,n})^\top | \mathcal{F}_{n-1}] \xrightarrow[n \rightarrow +\infty]{\text{a.s.}} \Gamma_{t,j},$$

where $\Gamma_{t,j}$ is a symmetric positive matrix.

(A6b) (*Covariate-adjusted approach*). For some $j \in \{1, \dots, d\}$, consider that there are features of interest β^j related with the entire family of distributions $\{\pi_t^j; t \in \tau\}$ of the responses to treatment j conditionally on the covariates. Then we assume that the corresponding adaptive estimator $\hat{\beta}_n^j$ is strongly consistent and its dynamics can be expressed as follows:

$$(3.6) \quad \hat{\beta}_n^j - \hat{\beta}_{n-1}^j = -\frac{\bar{X}_n^j}{N_n^j} (f_j(\hat{\beta}_{n-1}^j) - \Delta \mathbf{M}_{j,n} - \mathbf{R}_{j,n}),$$

where the quantities in (3.6) fulfill the same conditions presented above for the dynamics (3.3).

REMARK 3.7. Assumption (A6) is usually satisfied in most relevant cases (see [40] for the generalized urn model). For instance, a sufficient condition for the consistency of $\hat{\theta}_{t,n}^j$, which is generally true for most practical situations, is that the symmetric part of $\mathcal{D}f_{t,j}$ is positive definite in the entire parameter space. This can be proved using analogous arguments to those used in the proof of part (a) of Theorem 3.1. As an example, whenever $\theta_t^j = \mathbb{E}[g(\xi_n^j) | T_n = t]$ for some function g , the sample mean estimator

$$\hat{\theta}_{t,n}^j = (N_{t,n}^j)^{-1} \sum_{i=1}^n \mathbb{1}_{\{T_i=t\}} \bar{X}_i^j g(\bar{\xi}_i),$$

satisfies (A6a) with $\Delta M_{t,j,n} = (g(\bar{\xi}_n) - \theta_t^j)$ and $f_{t,j}(\hat{\theta}_{t,n-1}^j) = (\hat{\theta}_{t,n-1}^j - \theta_t^j)$, which implies $\mathcal{D}f_{t,j} = I$.

REMARK 3.8. When some parameters are related to the response distribution of more than one treatment, the corresponding estimators satisfy analogous conditions to those in (A6). Specifically, if one parameter does not depend on the treatment, instead of (3.3) we can consider

$$(3.7) \quad \hat{\theta}_{t,n} - \hat{\theta}_{t,n-1} = -\frac{\mathbb{1}_{\{T_n=t\}}}{w(\mathbf{N}_{t,n})} (f_t(\hat{\theta}_{t,n-1}) - \Delta \mathbf{M}_{t,n} - \mathbf{R}_{t,n}),$$

while instead of (3.6) we can consider

$$(3.8) \quad \hat{\beta}_n - \hat{\beta}_{n-1} = -\frac{1}{n} (f(\hat{\beta}_{n-1}) - \Delta \mathbf{M}_n - \mathbf{R}_n).$$

Further dynamics may be considered when some parameters depend, for instance, on a proper subset of the possible treatments.

(A7) *Conditional distribution of the covariates.*

(A7a) (*Covariate-stratification approach*) Let $\{\mu_n; n \geq 0\}$ be the sequence of probability measures on τ such that $\mu_n(t) = \mathbb{P}(T_{n+1} = t | \mathcal{F}_n)$ for any $t \in \tau$. Assume that there exists $n_0 \geq 1$ such that for any $n \geq n_0$,

$$(3.9) \quad \mu_n(t) = f_{\mu,t}(\hat{\theta}_{t,n}^j, \mathbf{N}_{t,n}/w(\mathbf{N}_{t,n})),$$

where $\{f_{\mu,t}; t \in \tau\}$ are differentiable functions, $f_{\mu,t}(\cdot) \geq \varepsilon$ for some $\varepsilon > 0$ and $\sum_{t=1}^K f_{\mu,t}(\cdot) = 1$. Denoting λ_μ^* the eigenvalue of $\text{Sp}(\sum_{s=1}^K \mathbf{v}(s) \mathcal{D}_N f_{\mu,s}(\mathbf{x}_0, \boldsymbol{\beta})^\top)$ with the largest real part, we assume that $\max_{t \in \tau} \text{Re}(\lambda_\mu^*(t)) < 1/2$.

(A7b) (*Covariate-adjusted approach*). Let $\{\mu_n; n \geq 0\}$ be the sequence of probability measures on τ such that $\mu_n(t) = \mathbb{P}(T_{n+1} = t | \mathcal{F}_n)$. Assume that there exists $n_0 \geq 1$ such that for any $n \geq n_0$,

$$(3.10) \quad \mu_n(t) = f_{\mu,t}(\hat{\boldsymbol{\beta}}_n^j, \mathbf{N}_n/n),$$

where $\{f_{\mu,t}; t \in \tau\}$ are differentiable functions, $f_{\mu,t}(\cdot) \geq 0$ and $\sum_{t=1}^K f_{\mu,t}(\cdot) = 1$. Moreover, we assume there exists an internal point $\mathbf{x}_0 \in \mathcal{S}$ that verifies $\mathbf{x}_0 = \sum_{s=1}^K f_{\mu,s}(\mathbf{x}_0, \boldsymbol{\beta}) \mathbf{v}(s)$. Denoting λ_μ^* the eigenvalue of $\text{Sp}(\sum_{s=1}^K \mathbf{v}(s) \mathcal{D}_N f_{\mu,s}(\mathbf{x}_0, \boldsymbol{\beta})^\top)$ with the largest real part, we assume that $\text{Re}(\lambda_\mu^*) < 1/2$.

REMARK 3.9. The condition $f_{\mu,t}(\cdot) \geq \varepsilon > 0$ ensures that any covariate profile $t \in \tau$ is asymptotically observable with positive probability. This assumption is essential to study the behavior of $\hat{\theta}_{t,n}^j$ and $\mathbf{N}_{t,n}$, while it is not necessary for $\hat{\boldsymbol{\beta}}_n^j$ and \mathbf{N}_n . For this reason, it will be required in Theorem 3.2 but not in Theorem 3.3.

REMARK 3.10. In the special case that the covariate profiles of the subjects $\{T_n; n \geq 1\}$ is a sequence of i.i.d. random variables, we have $\mu_i = \mu$ for any $i \geq 0$, and hence the above assumptions are satisfied in a straightforward manner with $f_{\mu,t}(\cdot) = \mu(t)$ and $\mathcal{D}_N f_{\mu,t} = \mathcal{D}_N f_{\boldsymbol{\beta},t} = 0$ for any $t \in \tau$.

3.2.2. *Second-order asymptotic results.* We first provide the convergence rate and the joint asymptotic distribution concerning the quantities of interest in the design in the framework of covariate-stratification response-adaptive designs. This result is established in the following central limit theorem. We introduce the variables independent of $\sigma(\mathcal{F}_n; n \geq 1)$: $T \in \tau$ with distribution $\mu(t)$, $\bar{\mathbf{X}} \in \{0, 1\}^d \in \mathcal{S}$ such that $\mathbb{P}(\bar{X}^j = 1 | T) = v^j(T)$, $D := \mathbb{E}[\bar{D} | T, \bar{\mathbf{X}}, \bar{\boldsymbol{\xi}}]$, where the distribution of $\bar{\boldsymbol{\xi}}$ conditioned on $\{T = t\}$ and $\{\bar{X}^j = 1\}$ is π_t^j .

THEOREM 3.2. Define $\mathbf{W}_n := (\mathbf{Z}_n(t), \mathbf{N}_{t,n}/w(\mathbf{N}_{t,n}), \hat{\theta}_{t,n}^j, t \in \tau)^\top$, $\mathbf{W} := (\mathbf{v}(t), \mathbf{v}(t), \boldsymbol{\theta}_t, t \in \tau)^\top$ and assume (A1)–(A5), (A6a), (A7a). Then

$$(3.11) \quad \mu_n(t) \xrightarrow{a.s.} \mu(t) = f_{\mu,t}(\mathbf{v}(t), \boldsymbol{\theta}_t), \quad \mathbf{W}_n \xrightarrow{a.s.} \mathbf{W},$$

$$(3.12) \quad \sqrt{n}(\mathbf{W}_n - \mathbf{W}) \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \Sigma), \quad \Sigma := \int_0^\infty e^{u(\frac{1}{2}-A)} \Gamma e^{u(\frac{1}{2}-A^\top)} du,$$

where

$$A := \begin{pmatrix} A_{ZZ} & 0 & 0 \\ -I & I & 0 \\ 0 & 0 & A_{\theta\theta} \end{pmatrix}, \quad \Gamma := \begin{pmatrix} \Gamma_{ZZ} & \Gamma_{ZN} & \Gamma_{Z\theta} \\ \Gamma_{ZN}^\top & \Gamma_{NN} & 0 \\ \Gamma_{Z\theta}^\top & 0 & \Gamma_{\theta\theta} \end{pmatrix},$$

and $A_{ZZ}, A_{\theta\theta}, \Gamma_{NN}, \Gamma_{\theta\theta}$ are block-diagonal matrices whose t th block is:

- (i) $A_{ZZ}^t = (I - H(t) + \mathbf{v}(t)\mathbf{1}^\top)$;
- (ii) $A_{\theta\theta}^t$ is a block-diagonal matrices whose j th block is $[A_{\theta\theta}^t]^{jj} := \mathcal{D}f_{t,j}(\boldsymbol{\theta}_t^j)$;
- (iii) $\Gamma_{NN}^t := \mu^{-1}(t)(\text{diag}(\mathbf{v}(t)) - \mathbf{v}(t)\mathbf{v}^\top(t))$;
- (iv) $\Gamma_{\theta\theta}^t$ is a block-diagonal matrices whose j th block is $[\Gamma_{\theta\theta}^t]^{jj} := (v^j(t)\mu(t))^{-1}\mathbb{E}[\Delta\mathbf{M}_{t,j}(\Delta\mathbf{M}_{t,j})^\top | T = t, \bar{X}^j = 1]$;

and $\Gamma_{ZZ}, \Gamma_{ZN}, \Gamma_{Z\theta}$ are matrices defined as follows: for any $t_1, t_2 \in \tau$,

- (v) $\Gamma_{ZZ}^{t_1 t_2} := \mathbb{E}[D(t_1)\mathbf{g}(t_1, T, \bar{\mathbf{X}})\mathbf{g}^\top(t_2, T, \bar{\mathbf{X}})D^\top(t_2)] - \mathbf{v}(t_1)\mathbf{v}^\top(t_2)$;
- (vi) $\Gamma_{ZN}^{t_1 t_2} := H(t_1)G(t_1, t_2)\text{diag}(\mathbf{v}(t_2)) - \mathbf{v}(t_1)\mathbf{v}^\top(t_2)$;
- (vii) $[\Gamma_{Z\theta}^{t_1 t_2}]^j := \mathbb{E}[D(t_1)\mathbf{g}(t_1, t_2, \mathbf{e}_j)\Delta\mathbf{M}_{t_2,j}^\top | T = t_2, \bar{X}^j = 1]$;

where \mathbf{g} is a d -multivariate function with values in \mathcal{S} defined in (B.16) in the Supplementary Material [4], and $G(t_1, t_2)$ is a matrix with columns $\{\mathbf{g}(t_1, t_2, \mathbf{e}_j); j \in \{1, \dots, d\}\}$.

We now provide the convergence rate and the joint asymptotic distribution of the quantities interest in the design in the framework of covariate-adjusted response-adaptive designs. This result is established in the following central limit theorem.

THEOREM 3.3. Define $\mathbf{W}_n := (\mathbf{Z}_n(t), t \in \tau, \mathbf{N}_n/n, \hat{\boldsymbol{\beta}}_n)^\top$, $\mathbf{W} := (\mathbf{v}(t), t \in \tau, \mathbf{x}_0, \boldsymbol{\beta})^\top$ and assume (A1)–(A5), (A6b), (A7b). Then

$$(3.13) \quad \mu_n(t) \xrightarrow{a.s.} \mu(t) = f_{\mu,t}(\mathbf{x}_0, \boldsymbol{\beta}), \quad \mathbf{W}_n \xrightarrow{a.s.} \mathbf{W},$$

$$(3.14) \quad \sqrt{n}(\mathbf{W}_n - \mathbf{W}) \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \Sigma), \quad \Sigma := \int_0^\infty e^{u(\frac{1}{2}-A)}\Gamma e^{u(\frac{1}{2}-A^\top)} du$$

and

$$A := \begin{pmatrix} A_{ZZ} & 0 & 0 \\ A_{NZ} & A_{NN} & A_{N\beta} \\ 0 & 0 & A_{\beta\beta} \end{pmatrix}, \quad \Gamma := \begin{pmatrix} \Gamma_{ZZ} & \Gamma_{ZN} & \Gamma_{Z\beta} \\ \Gamma_{ZN}^\top & \Gamma_{NN} & 0 \\ \Gamma_{Z\beta}^\top & 0 & \Gamma_{\beta\beta} \end{pmatrix},$$

where again and $A_{ZZ}, A_{\beta\beta}, \Gamma_{\beta\beta}$ are block-diagonal matrices whose t th or j th block is:

- (i) $A_{ZZ}^t = (I - H(t) + \mathbf{v}(t)\mathbf{1}^\top)$;
- (ii) $A_{\beta\beta}^{jj} = \mathcal{D}f_j(\boldsymbol{\beta}^j)$;
- (iii) $\Gamma_{\beta\beta}^{jj} := (\mathbb{E}[v^j(T)])^{-1}\mathbb{E}[\Delta\mathbf{M}_j(\Delta\mathbf{M}_j)^\top | \bar{X}^j = 1]$;

and

- (iv) $A_{NN} := I - \sum_{s=1}^K \mathbf{v}(s) \mathcal{D}_N f_{\mu,s}(\mathbf{x}_0, \boldsymbol{\beta})^\top$;
- (v) $A_{N\beta} := - \sum_{s=1}^K \mathbf{v}(s) \mathcal{D}_\beta f_{\mu,s}(\mathbf{x}_0, \boldsymbol{\beta})^\top$;
- (vi) $\Gamma_{NN} := \text{diag}(\mathbb{E}[\mathbf{v}(T)]) - \mathbb{E}[\mathbf{v}(T)]\mathbb{E}[\mathbf{v}^\top(T)]$;

and $A_{NZ}, \Gamma_{ZZ}, \Gamma_{ZN}, \Gamma_{Z\beta}$ are matrices defined as follows: for any $t_1, t_2 \in \tau$,

- (vii) $A_{NZ}^{t_2} := -\mu(t_2)I$;
- (viii) $\Gamma_{ZZ}^{t_1 t_2} := \mathbb{E}[D(t_1)\mathbf{g}(t_1, T, \bar{\mathbf{X}})\mathbf{g}^\top(t_2, T, \bar{\mathbf{X}})D^\top(t_2)] - \mathbf{v}(t_1)\mathbf{v}^\top(t_2)$;
- (ix) $\Gamma_{ZN}^{t_1} := H(t_1)\mathbb{E}[G(t_1, T)\text{diag}(\mathbf{v}(T))] - \mathbf{v}(t_1)\mathbb{E}[\mathbf{v}^\top(T)]$;
- (x) $\Gamma_{Z\beta}^{t_1 j} := \mathbb{E}[D(t)\mathbf{g}(t_1, T, j)\Delta\mathbf{M}_j^\top | \bar{X}^j = 1]$,

where we recall that \mathbf{g} is a d -multivariate function with values in \mathcal{S} defined in (B.16) in the Supplementary Material [4], and $G(t_1, t_2)$ is a matrix with columns $\{\mathbf{g}(t_1, t_2, \mathbf{e}_j); j \in \{1, \dots, d\}\}$.

REMARK 3.11. We recall that Theorem 3.2 allows inferential procedures based on stratified estimators, while Theorem 3.3 allows inference on covariate-adjusted regression parameters representing the covariate-adjusted treatment effect.

REMARK 3.12. In the hypothesis of Remark 3.8 when $\hat{\boldsymbol{\theta}}_t$ is defined as in (3.7), Theorem 3.2 holds with (iv) and (vii) in Γ replaced by:

- (iv) $\Gamma_{\theta\theta}^{tt} := (\mu(t))^{-1}\mathbb{E}[\Delta\mathbf{M}_t(\Delta\mathbf{M}_t)^\top | T = t]$;
- (vii) $\Gamma_{Z\theta}^{t_1 t_2} := \mathbb{E}[D(t_1)\mathbf{g}(t_1, t_2, \bar{\mathbf{X}})\Delta\mathbf{M}_{t_2}^\top | T = t_2]$.

Analogously, when $\hat{\boldsymbol{\beta}}$ is defined as in (3.7), Theorem 3.3 holds with (iii) and (x) in Γ replaced by:

- (iii) $\Gamma_{\beta\beta} := \mathbb{E}[\Delta\mathbf{M}(\Delta\mathbf{M})^\top]$;
- (x) $\Gamma_{Z\beta}^t := \mathbb{E}[D(t)\mathbf{g}(t, T, \bar{\mathbf{X}})\Delta\mathbf{M}^\top]$.

REMARK 3.13. If there are no covariates (i.e., τ is a singleton), these results reduce to the model investigated in [46]. In this case, $(A - I/2)$ corresponds to Q^\top in [46], and indeed

$$A_{ZZ} = (I - H + \mathbf{v}\mathbf{1}^\top) = I/2 - Q_{ZZ}^\top,$$

$$A_{NN} = I = I/2 - Q_{NN}^\top \quad \text{and} \quad A_{ZN} = -I = -Q_{ZN}^\top.$$

We compute

$$\Gamma_{NN} = \text{diag}(\mathbf{v}) - \mathbf{v}\mathbf{v}^\top = \Sigma_1,$$

$$\Gamma_{ZN} = H \text{diag}(\mathbf{v}) - \mathbf{v}\mathbf{v}^\top = H\Sigma_1,$$

$$\Gamma_{ZZ} = \mathbb{E}[D\bar{\mathbf{X}}\bar{\mathbf{X}}D^\top] - \mathbf{v}\mathbf{v}^\top = H\Sigma_1H^\top + \Sigma_2.$$

The last equation follows, since using $H\mathbf{v} = \mathbf{v}$ and denoting V_j the covariance matrix of $\mathbf{D}_{\cdot j}$, that is, the j th column of D , we have

$$\mathbb{E}[D\bar{\mathbf{X}}\bar{\mathbf{X}}D^\top] = \sum_{j=1}^d v^j \mathbb{E}[\mathbf{D}_{\cdot j}\mathbf{D}_{\cdot j}^\top] = \sum_{j=1}^d v^j (V_j + \mathbf{H}_{\cdot j}\mathbf{H}_{\cdot j}^\top) = \Sigma_2 + H \operatorname{diag}(\mathbf{v})H^\top$$

and

$$\mathbf{v}\mathbf{v}^\top = H \operatorname{diag}(\mathbf{v})H^\top - H\Sigma_1H^\top.$$

EXAMPLE 3.1. Consider the inferential problem of testing the equivalence of the effects of $d = 2$ treatments in presence $K = 2$ covariate profiles, under the following CDFs of the responses: $F_1^j(y) = y^{1/\alpha}$ when $T = 1$ and $F_2^j(y) = y^{1/\beta}$ when $T = 2$, where α and β are positive parameters. Since $\log(1/\bar{\xi}_i)$ is an exponential random variable with mean α when $T = 1$, we consider the following adaptive estimator:

$$\hat{\alpha}_n = \frac{\sum_{i=1}^n \mathbb{1}_{\{T_i=1\}} \log(\bar{\xi}_i^{-1})}{\sum_{i=1}^n \mathbb{1}_{\{T_i=1\}}},$$

which satisfies (3.7) in (A6) with $f_1(\hat{\alpha}_{n-1}) = (\hat{\alpha}_{n-1} - \alpha)$ and $\Delta M_{1,n} = (\log(\bar{\xi}_n^{-1}) - \alpha)$. Analogous arguments hold to construct the estimator $\hat{\beta}_n$ of β . In this case, $\mathcal{D}f_1 = \mathcal{D}f_2 = 1$, we have $A_{\theta\theta} = I$ and we need to compute $H(t)$ to find A_{ZZ} . To this end note that, since the QFs are $Q_1^j(v) = v^\alpha$ and $Q_2^j(v) = v^\beta$, we have

$$D_n(1) = \begin{pmatrix} V_n^\alpha & 1 - V_n^\alpha \\ 1 - V_n^\alpha & V_n^\alpha \end{pmatrix}, \quad D_n(2) = \begin{pmatrix} V_n^\beta & 1 - V_n^\beta \\ 1 - V_n^\beta & V_n^\beta \end{pmatrix},$$

where we recall that $V_n \sim U(0, 1)$, which implies

$$H(1) = (1 + \alpha)^{-1} \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix}, \quad H(2) = (1 + \beta)^{-1} \begin{pmatrix} 1 & \beta \\ \beta & 1 \end{pmatrix}.$$

From the structure of $H(t)$ above, we obtain $\mathbf{v}(t) = (1/2, 1/2)^\top$ for any $t \in \{1, 2\}$ and the condition $\{\max_t \lambda_{H^*}^*(t) < 1/2\}$ in (A5) is verified for $\alpha, \beta > 1/3$. Hence, we can compute

$$A_{ZZ}^{11} = \begin{pmatrix} 1/2 + \alpha & 1/2 - \alpha \\ 1/2 - \alpha & 1/2 + \alpha \end{pmatrix}, \quad A_{ZZ}^{22} = \begin{pmatrix} 1/2 + \beta & 1/2 - \beta \\ 1/2 - \beta & 1/2 + \beta \end{pmatrix}.$$

Since $\mathbf{v}(t) = (1/2, 1/2)^\top$ implies $\mathbf{g}(t_1, t_2, \mathbf{e}_j) = \mathbf{e}_j$ for any t_1, t_2 , we obtain

$$\Gamma_{ZZ}^{11} = [(2\alpha + 1)(\alpha + 1)]^{-1} \begin{pmatrix} 2\alpha^2 + \alpha + 1 & 2\alpha \\ 2\alpha & 2\alpha^2 + \alpha + 1 \end{pmatrix},$$

$$\Gamma_{ZZ}^{12} = [(\alpha + \beta + 1)(\alpha + 1)(\beta + 1)]^{-1}$$

$$\times \begin{pmatrix} \alpha^2 + \alpha + \beta^2 + \beta \\ 2(\alpha + 1)(\beta + 1) + (\alpha + \beta + 1)(\alpha\beta + 1) \\ 2(\alpha + 1)(\beta + 1) + (\alpha + \beta + 1)(\alpha\beta + 1) \\ \alpha^2 + \alpha + \beta^2 + \beta \end{pmatrix},$$

while Γ_{ZZ}^{22} is the same as Γ_{ZZ}^{11} with α replaced by β . Then, defining the 2×2 -matrix $J := (2I - \mathbf{1}\mathbf{1}^\top)$, we have $\Gamma_{NN}^{11} = (4\mu(1))^{-1}J$ and $\Gamma_{NN}^{22} = (4\mu(2))^{-1}J$, while for any $t = 1, 2$ we have $\Gamma_{ZN}^{1t} = (\alpha - 1)/(2\alpha + 2)J$ and $\Gamma_{ZN}^{2t} = (\beta - 1)/(2\beta + 2)J$. Moreover, $\Gamma_{\alpha\alpha} = (\mu(1))^{-1}\alpha^2$, $\Gamma_{\beta\beta} = (\mu(2))^{-1}\beta^2$. Finally, $\Gamma_{Z\alpha}^t = \Gamma_{Z\beta}^t = \mathbf{0}$ for any $t \in \tau$.

4. Application to responses with Gaussian conditional distribution. In this section, we analyze the functional urn model in the case that the distribution of the responses to treatments, conditionally on the covariates, are Gaussian. In particular, consider the following model between the covariates and the responses to treatment j , $j \in \{1, \dots, d\}$:

$$(4.1) \quad \xi_n^j = g^j(T_n) + \varepsilon_n^j \quad \forall n \geq 1,$$

where $g^j \in L^2(\tau)$ and $\{\varepsilon_n^j; n \geq 1\} \sim \text{i.i.d. } \mathcal{N}(0, \sigma_j^2)$. We consider g^j and σ_j^2 unknown and we denote by \hat{g}^j and $\hat{\sigma}_j^2$ the corresponding consistent estimators. For instance, in a parametric setting we may assume $g^j(t) = \sum_{i=1}^M \beta_i^j \phi_i(t)$ for some $M \in \mathbb{N}$, $\beta_i^j \in \mathbb{R}$ and $\phi_i \in L^2(\tau)$. Then, letting $\sigma_j = \sigma$ for all $j \in \{1, \dots, d\}$, the model (4.1) represents the classical regression analysis with independent and homoscedastic errors. In this case, $\hat{g}^j(t) = \sum_{i=1}^M \hat{\beta}_i^j \phi_i(t)$, where $\hat{\beta}_i^j$ are the least square estimators.

From (4.1), we have that, conditionally on the set $\{T_n = t\}$, the response ξ_n^j is normally distributed with mean $g^j(t)$ and variance σ_j^2 . Hence, the family of probability distribution $\{\pi_t^j, t \in \tau\}$ is represented by

$$(4.2) \quad \pi_t^j = \mathcal{N}(g^j(t), \sigma_j^2).$$

Analogously, we can define the CDF for any $y \in \mathbb{R}$ and the QF for any $v \in (0, 1)$ as follows:

$$(4.3) \quad F_t^j(y) = \Phi\left(\frac{y - g^j(t)}{\sigma_j}\right), \quad Q_t^j(v) = g^j(t) + \sigma_j z_v,$$

where Φ and z_v are, respectively, the CDF and the QF of a standard normal variable.

4.1. *Convergence to target functions.* We now show how in the model (4.1) the probability of assigning a patient with covariate profile $t \in \tau$, that is, $\mathbf{Z}_n(t)$, can converge to any desired target $\boldsymbol{\rho}(t) = \boldsymbol{\rho}(\boldsymbol{\theta}_t)$, where $\boldsymbol{\theta}_t$ are parameters of the conditional distributions π_t^j , $j = \{1, \dots, d\}$. Since π_t^j admits a density function in \mathbb{R} , we can express the functional replacement matrix D_n as defined in (A.2) in the Supplementary Material [4]: conditionally on $\{T_n = s\}$, $\{\bar{X}_n^k = 1\}$ and $\{\xi_n^k = y\}$, we have that for any $t \in \tau$

$$D_n^{ij}(t) = \hat{u}^{ij}(\hat{g}^j(t) + \hat{\sigma}_j z_{\hat{f}_s^k(y)}) = \hat{u}^{ij} \left(\hat{g}^j(t) + \frac{\hat{\sigma}_j}{\hat{\sigma}_k} (y - \hat{g}^k(s)) \right).$$

The consistency of the estimators ensures that Assumption (A2) is satisfied and the asymptotic behavior of the urn process is determined by the limiting generating matrix H defined as

$$H^{ij}(t) = \mathbb{E}[\check{D}_n^{ij}(t)] = \mathbb{E}[u^{ij}(Q_t^j(V_n))] = \mathbb{E}[u^{ij}(g^j(t) + \sigma_j z_{V_n})],$$

where $z_{V_n} \sim \mathcal{N}(0, 1)$ since $V_n \sim U(0, 1)$. Thus, from Theorem 3.1 we have that $\mathbf{Z}_n(t) \xrightarrow{\text{a.s.}} \mathbf{v}(t)$, where \mathbf{v} is such that $w(\mathbf{v}) = 1$ and $H\mathbf{v} = \mathbf{v}$. Hence, the functions u^{ij} , $i, j \in \{1, \dots, d\}$, can be chosen such that $\mathbf{v}(t)$ coincides with the desired target function $(\rho^1(t), \dots, \rho^d(t))^\top$.

In the case when $d = 2$ treatments, we now consider the target proportion allocation proposed in [45] for responses distributed as $\mathcal{N}(m_1, \sigma_1)$ for treatment 1 and $\mathcal{N}(m_2, \sigma_2)$ for treatment 2:

$$(4.4) \quad \tilde{\rho}(m_1, m_2, \sigma_1, \sigma_2) = \frac{\sigma_1 \sqrt{m_2}}{\sigma_1 \sqrt{m_2} + \sigma_2 \sqrt{m_1}}, \quad m_1, m_2 > 0.$$

As described in [45], the allocation proportion (4.4) minimizes the total expected responses from all the subjects $[n(m_1 \cdot \tilde{\rho} + m_2 \cdot (1 - \tilde{\rho}))]$ with a fixed variance, $[n^{-1}(\sigma_1^2/\tilde{\rho} + \sigma_2^2/(1 - \tilde{\rho}))]$. In our framework, the target proportion is the function $(\rho(t), 1 - \rho(t))^\top$, where $\rho(t) = \tilde{\rho}(g^1(t), g^2(t), \sigma_1, \sigma_2)$. To achieve this limiting proportion, we need to define the functions $u_t^{11}, u_t^{12}, u_t^{21}$ and u_t^{22} such that the normalized right eigenvector of $H(t)$ is $\mathbf{v}(t) = (\rho(t), 1 - \rho(t))^\top$. For instance, a possible choice is the following: for any $y \in \mathbb{R}$,

$$\hat{u}_t^{11}(y) = \tilde{\rho}(\hat{g}^1(t), \hat{g}^2(t), \hat{\sigma}_1, \hat{\sigma}_2),$$

$\hat{u}_t^{12} = \hat{u}_t^{11}$ and $\hat{u}_t^{21} = \hat{u}_t^{22} = 1 - \hat{u}_t^{11}$. Hence, we have that $H^{11}(t) = \lim_{n \rightarrow \infty} \mathbb{E}[\hat{u}_t^{11}(\xi_n^1) | T_n = t] = \rho(t)$ a.s. and, analogously, $H^{12}(t) = \rho(t)$ and $H^{21}(t) = H^{22}(t) = 1 - \rho(t)$, which implies $\mathbf{v} = (\rho(t), 1 - \rho(t))^\top$.

4.2. *Inference on conditional response distribution.* We now analyze how to do inference in the model (4.1). Specifically, we consider the problem of testing the equivalence of the response means conditionally on the covariates, that is, $H_0 : g^j = g$ for any $j \in \{1, \dots, d\}$, with g given function in $L^2(\tau)$. Take $d = 2$

treatments, consider $K \geq 1$ possible covariate profiles and assume $\sigma_1 = \sigma_2 = \sigma \in (0, \infty)$ to be known. We set

$$u_t^{11}(y) = u_t^{22}(y) = \phi\left(\frac{y - g(t)}{\sigma}\right),$$

$u_t^{21} = 1 - u_t^{11}$ and $u_t^{12} = 1 - u_t^{22}$. From this choice, under H_0 we have that: $D_n^{jj}(t) = \phi((\xi_n^j - g(t))/\sigma) \sim U(0, 1)$ and for $i \neq j$ $D_n^{ij} = 1 - D_n^{jj} \sim U(0, 1)$, which implies $H^{ij}(t) = 1/2$ for any $i, j \in \{1, 2\}$, and hence $\mathbf{v}(t) = (1/2, 1/2)^\top$ and $\lambda_H^*(t) = 0$ for any $t \in \tau$. Then we can apply the CLT established in Theorem 3.3 to construct inferential procedures to test the null hypothesis. It is worth seeing how the dynamics of the functional urn model changes when H_0 does not hold. In particular, under $H_1 : \{g^1 = g + \Delta\}$, for some $\Delta \in L^2(\tau)$, we have

$$D_n^{11}(t) = \phi\left(\frac{\xi_n^1 - g^1(t)}{\sigma} + \frac{\Delta(t)}{\sigma}\right) \sim \phi\left(z_{V_n} + \frac{\Delta(t)}{\sigma}\right),$$

where $z_{V_n} \sim \mathcal{N}(0, 1)$.

5. Application to responses with Bernoulli conditional distribution. In this section, we analyze the functional urn model when the responses to treatments, conditionally on the covariates, are Bernoulli distributed. In particular, consider the following model between the covariates and the responses to treatment $j, j \in \{1, \dots, d\}$:

$$(5.1) \quad \xi_n^j = \mathbb{1}_{\{U_n^j \leq p^j(T_n)\}} \quad \forall n \geq 1,$$

where $U_n^j \sim$ i.i.d. $U(0, 1)$ and $0 < p^j(t) < 1$ for any $t \in \tau$. We consider p^j unknown and we denote by \hat{p}^j its consistent estimator. If we assume there exist $M \in \mathbb{M}, \beta_i^j \in \mathbb{R}$ and $\phi_i \in L^2(\tau)$ such that

$$(5.2) \quad \log\left(\frac{p^j(t)}{1 - p^j(t)}\right) = \sum_{i=1}^M \beta_i^j \phi_i(t),$$

the model (5.1) represents the classical logistic regression for binary responses. In this case, $\hat{p}^j(t) = \sum_{i=1}^M \hat{\beta}_i^j \phi_i(t)$, where $\hat{\beta}_i^j$ are the maximum likelihood estimators (MLEs) of β_i^j .

From (5.1), we have that, conditionally on the set $\{T_n = t\}, \{\xi_n^j; n \geq 1\}$ represents a sequence of independent Bernoulli random variables with parameter $p^j(t)$. Hence, the probability measures $\{\pi_t^j, t \in \tau\}$ are Bernoulli distributed as

$$(5.3) \quad \pi_t^j = \mathcal{B}e(p^j(t)) \quad \forall n \geq 1.$$

Analogously, we can define the CDF for any $y \in \mathbb{R}$ and the QF for any $v \in (0, 1)$:

$$(5.4) \quad F_t^j(y) = (1 - p^j(t))\mathbb{1}_{\{y \geq 0\}} + p^j(t)\mathbb{1}_{\{y \geq 1\}}, \quad Q_t^j(v) = \mathbb{1}_{\{v \geq 1 - p^j(t)\}}.$$

5.1. *Convergence to target functions.* Since π_t^j is a discrete distribution, the functional replacement matrix D_n can be expressed as in (A.3) in the Supplementary Material [4]: conditionally on $\{T_n = s\}$, $\{\bar{X}_n^k = 1\}$ and $\{\xi_n^k = y\}$, for any $t \in \tau$ we have that

$$D_n^{ij}(t) = \begin{cases} (1 - \hat{p}^k(s))^{-1} \cdot \int_0^{1-\hat{p}^k(s)} \hat{u}^{ij}(\mathbb{1}_{\{v \geq 1-\hat{p}^j(t)\}}) dv & \text{if } y = 0; \\ (\hat{p}^k(s))^{-1} \cdot \int_{1-\hat{p}^k(s)}^1 \hat{u}^{ij}(\mathbb{1}_{\{v \geq 1-\hat{p}^j(t)\}}) dv & \text{if } y = 1. \end{cases}$$

The consistency of the estimators ensures that Assumption (A2) is satisfied and the asymptotic behavior of the urn process is determined by the limiting generating matrix H defined as

$$H^{ij}(t) = \mathbb{E}[\check{D}_n^{ij}] = \mathbb{E}[u^{ij}(Q_t^j(V_n))] = \mathbb{E}[u^{ij}(\mathbb{1}_{\{V_n \geq 1-p^j(t)\}})],$$

where $V_n \sim U(0, 1)$. Thus, from Theorem 3.1 we have that $\mathbf{Z}_n(t) \xrightarrow{\text{a.s.}} \mathbf{v}(t)$, where \mathbf{v} is such that $w(\mathbf{v}) = 1$ and $H\mathbf{v} = \mathbf{v}$. Hence, the functions u_t^{ij} , $i, j \in \{1, \dots, d\}$, can be chosen such that $\mathbf{v}(t)$ coincides with the desired target function $(\rho^1(t), \dots, \rho^d(t))^\top$.

For instance, consider play-the-winner design for binary responses proposed in [43] and [44]. In the multi-treatments play-the-winner design, when treatment $j \in \{1, \dots, d\}$ is assigned, we replace in the urn a ball of color j if the response is a success or $(d - 1)^{-1}$ balls of each other color if the response is a failure. Thus, the play-the-winner rule can be implemented in our framework by setting $u_t^{ij}(y) = y\delta_{ij} + (1 - y)(1 - \delta_{ij})(d - 1)^{-1}$, for any $y \in \{0; 1\}$, where δ_{ij} is the delta of Kronecker. Note that this choice of u_t^{ij} guarantees the constant balance of the urn required in (A1). Then each element $D_n^{ij}(t)$ of the replacement matrix can be explicitly expressed as follows:

$$\begin{cases} \left(\frac{(\max\{\hat{p}^j(t); \hat{p}^k(s)\} - \hat{p}^k(s))}{1 - \hat{p}^k(s)} \right) \delta_{ij} \\ \quad + \left(\frac{1 - \max\{\hat{p}^j(t); \hat{p}^k(s)\}}{1 - \hat{p}^k(s)} \right) (1 - \delta_{ij})(d - 1)^{-1} & \text{if } y = 0; \\ \left(\frac{(\min\{\hat{p}^j(t); \hat{p}^k(s)\})}{\hat{p}^k(s)} \right) \delta_{ij} \\ \quad + \left(\frac{\hat{p}^k(s) - \min\{\hat{p}^j(t); \hat{p}^k(s)\}}{\hat{p}^k(s)} \right) (1 - \delta_{ij})(d - 1)^{-1} & \text{if } y = 1. \end{cases}$$

In this case, since $\mathbb{1}_{\{V_n \geq 1-p^j(t)\}} \sim \text{Be}(p^j(t))$ when $V_n \sim U(0, 1)$, we have

$$\begin{aligned} H_{n-1}^{ij}(t) &= \mathbb{E}[\mathbb{1}_{\{V_n \geq 1-p^j(t)\}} \delta_{ij} + (1 - \mathbb{1}_{\{V_n \geq 1-p^j(t)\}})(1 - \delta_{ij})(d - 1)^{-1}] \\ &= p^j(t)\delta_{ij} + (1 - p^j(t))(1 - \delta_{ij})(d - 1)^{-1}. \end{aligned}$$

Hence, the first right eigenvector $\mathbf{v}(t)$ of $H(t)$ associated with the eigenvalue $\lambda = 1$ is

$$(5.5) \quad v^j(t) = \rho^j(p^1(t), \dots, p^d(t)) = \frac{(1 - p^j(t))^{-1}}{\sum_{i=1}^d (1 - p^i(t))^{-1}},$$

and from Theorem 3.1 we have that $\mathbf{Z}_n(t) \xrightarrow{\text{a.s.}} (\rho^1(t), \dots, \rho^d(t))^\top$.

5.2. *Inference on conditional response distribution.* We now analyze how to do inference in the model (5.1). Specifically, we consider the problem of testing the equivalence of the success probabilities conditionally on the covariates, that is, $H_0 : p^j = p$ for any $j \in \{1, \dots, d\}$, with p given function with values in $(0, 1)$. Using the same functions defined above, that is, $u_t^{ij}(y) = y\delta_{ij} + (1 - y)(1 - \delta_{ij})(d - 1)^{-1}$, we have that, under H_0 ,

$$\begin{aligned} D^{ij}(t) &= \delta_{ij} \mathbb{1}_{\{\bar{\xi}_n=1\}} + (1 - \delta_{ij}) \mathbb{1}_{\{\bar{\xi}_n=0\}}(d - 1)^{-1} \\ &\sim \delta_{ij}W + (1 - \delta_{ij})(1 - W)(d - 1)^{-1}, \end{aligned}$$

where $W \sim \text{Be}(p)$. This implies $H^{jj}(t) = p(t)$ and $H^{ij}(t) = (1 - p(t))(d - 1)^{-1}$ for $i \neq j$, and hence $\mathbf{v}(t) = d^{-1}\mathbf{1}$ and $\lambda_H^*(t) = (dp(t) - 1)/(d - 1)$ for any $t \in \tau$. Then, if $\max_{t \in \tau} p(t) < (d + 1)/(2d)$ we have $\max_{t \in \tau} \lambda_H^*(t) < 1/2$, and hence we can apply the CLT established in Theorem 3.3 to construct inferential procedures to test the null hypothesis. It is worth seeing how the dynamics of the functional urn model changes when H_0 does not hold. In particular, under $H_1 : \{p^1 = p + \Delta\}$, for some Δ with values in $(0, 1)$, we have $D^{i1}(t) \sim \delta_{i1}W + (1 - \delta_{i1})(1 - W)(d - 1)^{-1}$, with $W \sim \text{Be}(p + \Delta)$.

5.3. *Example.* We provide a very simple example from a clinical trial of external cooling in patients with septic shock (the Sepsicool trial [37]). The trial found little difference in the primary outcome in the entire clinical trials population, but 14-day mortality, a different endpoint, was significantly lower among patients given external cooling in the subgroup with a lower baseline vasopressor dose (i.e., those patients who had less severe illness at baseline). We take the approach of [41] by redesigning the study using our methodology. Using their parameter values, obtained from the results presented in [37], we determine that the underlying probability of survival in the “no cooling” group is 0.657 regardless of severity; the probability of survival in the cooling group is 0.842 for those with low severity, and 0.406 for those with high or moderate severity. As in [41], we assume that 225 patients have low severity and 225 patients have high or moderate severity. We now redesign the trial using our methodology. For the underlying parameter set $\mathbf{p} = (p^1(1), p^1(2), p^2(1), p^2(2))$ is given by $(p^1(1) = p^1(2) = 0.657, p^2(1) = 0.842, p^2(2) = 0.406)$. Computing the asymptotic target allocation using (5.5), we can compute the expected number of deaths

in 450 patients using our methodology is 146.5. If we ignore the binary covariate in the urn process, the resulting expected number of deaths is 161.4.

We can conduct hypothesis testing on $p^j(t)$ using Theorem 3.2. For the simple case of two treatments and a single binary covariate, the joint asymptotic distribution of observed proportions $\hat{\boldsymbol{p}} = (\hat{p}^1(1), \hat{p}^1(2), \hat{p}^2(1), \hat{p}^2(2))$ can be expressed as

$$\sqrt{n}(\hat{\boldsymbol{p}}_n - \boldsymbol{p}) \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \text{diag}(d^1(1), d^1(2), d^2(1), d^2(2))),$$

$$\text{where } d^j(t) = \frac{2}{v^j(t)} p^j(t)(1 - p^j(t)).$$

Note that, when there are no covariates, this is the same asymptotic result from the generalized Pólya urn model [33].

REMARK 5.1. Several recent papers on CARA designs show that results on parameter estimation in the binary case are analogous to those presented in Example 5.3, either for the parameters of the logistic regression β^j (see, e.g., [19, 48, 50]) and for the parameters of the success probability $p^j(t)$ (see, e.g., [24]). The connection between these two approaches is also highlighted in [24], Section 2.2. Naturally, the same results presented in Example 5.3 can be obtained with the unified family of CARA designs recently proposed in [24]. Indeed, this follows by combining [24], Theorem 1 and [24], Example 2, setting $e_k(\cdot) = (1 - p_k(\cdot))^{-1}$ and $\gamma = 0$.

6. Conclusions. This paper proposes a general class of CARA designs that randomly assigns subjects to the treatment groups with a probability that depends on their own covariate profiles and on the previous patients' covariates, treatment assignments and responses. This procedure can be considered as a general methodology to incorporate both the responses and the covariates in the randomized treatment allocation scheme, as a possible alternative to the class of designs presented in [48]. The generality of the proposed framework includes several different ways to model the relationship between covariates and treatment responses, allowing and facilitating the implementation of these designs to a wide range of applications. In particular, in the paper we have discussed the properties with generalized linear models for different types of responses and covariates (discrete and continuous) and for more than two treatments. Moreover, this class of designs does not require that the probability distribution of the covariates be the same for all patients, which is a standard assumption in CARA designs (e.g., see [48]). In fact, we allow this distribution to be adaptively modified by the experimenter using the information collected during the trial. This improvement opens the possibility to apply the theory of optimal designs (see [12]) within the CARA framework for future research.

The class of CARA designs presented here is based on a new functional urn model that extends the classical theory of urn models adopted for response-adaptive designs, in which the covariate information is not considered in the randomization process (e.g., [20, 23, 30, 47]). The urn is represented by a multivariate function of the covariates, and each patient is assigned by sampling from the urn evaluated at his own covariate profile. After any allocation, the entire functional urn composition is updated, even if only the response associated with the patient's covariate profile has been observed, and this allows the incorporation of general covariate spaces in the design. In the context of personalized medicine, this feature can allow the investigation to determine optimal treatments based on the covariate model, even when there is insufficient information about a particular covariate profile.

In response-adaptive randomization, desirable ethical or inferential properties of the design are achieved by targeting an optimal allocation proportion, which is typically determined by some optimality criterion based on the response distribution. Analogously, when covariate information is considered in the trial, the purpose becomes to target an optimal allocation for each fixed value of the covariates. In this paper, we achieve this goal by establishing, in Theorem 3.1, the convergence to any optimal allocation. This result is obtained by allowing non-parametric or semiparametric estimates of the response distribution conditioned on the covariates. This extends the class of designs proposed in [48] in which the target allocation proportion depends on a finite number of parameters.

In addition, statistical inference on the treatment effects requires the establishment of the joint distribution of sufficient statistics that, in an adaptive setting, are represented by both allocation proportion and adaptive estimators. This is typically a hard task in the framework of CARA designs (see [35]). Theorem 3.2 and Theorem 3.3 provide the theoretical results which allow us to construct inferential procedures based on two different approaches: stratified and covariate-adjusted estimators. The study of their power under different types of alternative hypotheses is essential to investigate the performances of these procedures and to conduct comparisons with other existing CARA designs. In this paper, we have provided the general framework for investigations of this type under different models.

SUPPLEMENTARY MATERIAL

Online supplementary materials: Nonparametric covariate-adjusted response-adaptive design based on a functional urn model (DOI: [10.1214/17-AOS1677SUPP](https://doi.org/10.1214/17-AOS1677SUPP); .pdf). This supplement gives the analytic expressions used in the paper and the proofs of the theorems.

REFERENCES

- [1] ALETTI, G., CRIMALDI, I. and GHIGLIETTI, A. (2017). Synchronization of reinforced stochastic processes with a network-based interaction. *Ann. Appl. Probab.* **27** 3787–3844. MR3737938

- [2] ALETTI, G. and GHIGLIETTI, A. (2017). Interacting generalized Friedman's urn systems. *Stochastic Process. Appl.* **127** 2650–2678. [MR3660886](#)
- [3] ALETTI, G., GHIGLIETTI, A. and PAGANONI, A. M. (2013). Randomly reinforced urn designs with prespecified allocations. *J. Appl. Probab.* **50** 486–498. [MR3102495](#)
- [4] ALETTI, G., GHIGLIETTI, A. and ROSENBERGER, W. F. (2018). Supplement to “Nonparametric covariate-adjusted response-adaptive design based on a functional urn model.” DOI:10.1214/17-AOS1677SUPP.
- [5] ATHREYA, K. B. and KARLIN, S. (1970). Branching processes with random environments. *Bull. Amer. Math. Soc.* **76** 865–870.
- [6] ATKINSON, A. C. (1982). Optimum biased coin designs for sequential clinical trials with prognostic factors. *Biometrika* **69** 61–67.
- [7] ATKINSON, A. C. and BISWAS, A. (2005). Adaptive biased-coin designs for skewing the allocation proportion in clinical trials with normal responses. *Stat. Med.* **24** 2477–2492.
- [8] ATKINSON, A. C. and BISWAS, A. (2013). *Randomised Response-Adaptive Designs in Clinical Trials*. CRC Press, New York.
- [9] BAI, Z.-D. and HU, F. (2005). Asymptotics in randomized urn models. *Ann. Appl. Probab.* **15** 914–940. [MR2114994](#)
- [10] BAI, Z.-D., HU, F. and ZHANG, L.-X. (2002). Gaussian approximation theorems for urn models and their applications. *Ann. Appl. Probab.* **12** 1149–1173. [MR1936587](#)
- [11] BAI, Z. D. and HU, F. (1999). Asymptotic theorems for urn models with nonhomogeneous generating matrices. *Stochastic Process. Appl.* **80** 87–101. [MR1670107](#)
- [12] BALDI ANTOGNINI, A. and ZAGORAIOU, M. (2012). Multi-objective optimal designs in comparative clinical trials with covariates: The reinforced doubly adaptive biased coin design. *Ann. Statist.* **40** 1315–1345. [MR3015027](#)
- [13] BALDI ANTOGNINI, A. and ZAGORAIOU, M. (2017). Estimation accuracy under covariate-adaptive randomization procedures. *Electron. J. Stat.* **11** 1180–1206.
- [14] BANDYOPADHYAY, U. and BISWAS, A. (1999). Allocation by randomized play-the-winner rule in the presence of prognostic factors. *Sankhyā Ser. B* **61** 397–412.
- [15] BANDYOPADHYAY, U. and BISWAS, A. (2001). Adaptive designs for normal responses with prognostic factors. *Biometrika* **88** 409–419.
- [16] BANDYOPADHYAY, U., BISWAS, A. and BHATTACHARYA, R. (2007). A covariate adjusted two-stage allocation design for binary responses in randomized clinical trials. *Stat. Med.* **26** 4386–4399.
- [17] BEGGS, A. W. (2005). On the convergence of reinforcement learning. *J. Econom. Theory* **122** 1–36.
- [18] BENAÏM, M., SCHREIBER, S. J. and TARRÈS, P. (2004). Generalized URN models of evolutionary processes. *Ann. Appl. Probab.* **14** 1455–1478. [MR2071430](#)
- [19] CHEUNG, S. H., ZHANG, L.-X., HU, F. and CHAN, W. S. (2014). Covariate-adjusted response-adaptive designs for generalized linear models. *J. Statist. Plann. Inference* **149** 152–161.
- [20] DURHAM, S. D., FLOURNOY, N. and LI, W. (1998). A sequential design for maximizing the probability of a favourable response. *Canad. J. Statist.* **26** 479–495.
- [21] EGGENBERGER, F. and PÓLYA, G. (1923). Über die Statistik verketteter Vorgänge. *ZAMM Z. Angew. Math. Mech.* **3** 279–289.
- [22] GHIGLIETTI, A., VIDYASHANKAR, A. N. and ROSENBERGER, W. F. (2017). Central limit theorem for an adaptive randomly reinforced urn model. *Ann. Appl. Probab.* **27** 2956–3003. [MR3719951](#)
- [23] HU, F. and ROSENBERGER, W. F. (2006). *The Theory of Response-Adaptive Randomization in Clinical Trials*. Wiley, New York.
- [24] HU, J., ZHU, H. and HU, F. (2015). A unified family of covariate-adjusted response-adaptive designs based on efficiency and ethics. *J. Amer. Statist. Assoc.* **110** 357–367. [MR3338508](#)

- [25] HU, Y. and HU, F. (2012). Asymptotic properties of covariate-adaptive randomization. *Ann. Statist.* **40** 1794–1815. [MR3015044](#)
- [26] JANSON, S. (2004). Functional limit theorems for multitype branching processes and generalized Pólya urns. *Stochastic Process. Appl.* **110** 177–245.
- [27] KADANE, J. B. (1996). *Bayesian Methods and Ethics in a Clinical Trial Design*. Wiley, New York.
- [28] KALISH, L. A. and HARRINGTON, D. P. (1988). Efficiency of balanced treatment allocation for survival analysis. *Biometrics* **44** 815–821.
- [29] LARUELLE, S. and PAGÈS, G. (2013). Randomized urn models revisited using stochastic approximation. *Ann. Appl. Probab.* **23** 1409–1436. [MR3098437](#)
- [30] MAY, C. and FLOURNOY, N. (2009). Asymptotics in response-adaptive designs generated by a two-color, randomly reinforced urn. *Ann. Statist.* **37** 1058–1078. [MR2502661](#)
- [31] POCOCK, S. J. and SIMON, R. (1975). Sequential treatment assignment with balancing for prognostic factors in the controlled clinical trial. *Biometrics* **31** 103–115.
- [32] ROSENBERGER, W. F. (1996). New directions in adaptive designs. *Statist. Sci.* **11** 137–149. [MR1477667](#)
- [33] ROSENBERGER, W. F., FLOURNOY, N. and DURHAM, S. D. (1997). Asymptotic normality of maximum likelihood estimators from multiparameter response-driven designs. *J. Statist. Plann. Inference* **60** 69–76.
- [34] ROSENBERGER, W. F. and LACHIN, J. M. (2016). *Randomization in Clinical Trials*, 2nd ed. Wiley, New York.
- [35] ROSENBERGER, W. F. and SVERDLOV, O. (2008). Handling covariates in the design of clinical trials. *Statist. Sci.* **23** 404–419. [MR2483911](#)
- [36] ROSENBERGER, W. F., VIDYASHANKAR, A. N. and AGARWAL, D. K. (2001). Covariate-adjusted response-adaptive designs for binary response. *J. Biopharm. Statist.* **11** 227–236.
- [37] SCHORTGEN, F., CLABAULT, K., KATSAHIAN, S., DEVAQUET, J., MERCAT, A., DEYE, N., DELLAMONICA, J., BOUADMA, L., COOK, F., BEJI, O., BRUN-BUISSON, C., LEMAIRE, F. and BROCHARD, L. (2012). Fever control using external cooling in septic shock. *Am. J. Respir. Crit. Care Med.* **185** 1088–1095.
- [38] SMYTHE, R. T. (1996). Central limit theorems for urn models. *Stochastic Process. Appl.* **65** 115–137.
- [39] TAVES, D. R. (1974). Minimization: A new method of assigning patients to treatment and control groups. *Clin. Pharmacol. Ther.* **15** 443–453.
- [40] TOULIS, P., RENNIE, J. and AIROLDI, E. M. (2014). Statistical analysis of stochastic gradient methods for generalized linear models. In *31st International Conference on Machine Learning, ICML 2014* 2128–2136. JMLR: W&CP, Beijing.
- [41] VILLAR, S. and ROSENBERGER, W. F. (2018). Covariate-adjusted response-adaptive randomization for multi-arm clinical trials using a modified forward looking gittins index rule. *Biometrics* **74** 49–57.
- [42] WEI, L. J. (1978). An application of an urn model to the design of sequential controlled clinical trials. *J. Amer. Statist. Assoc.* **73** 559–563.
- [43] WEI, L. J. and DURHAM, S. (1978). The randomized play-the-winner rule in medical trials. *J. Amer. Statist. Assoc.* **73** 840–843.
- [44] ZELEN, M. (1969). Play the winner rule and the controlled clinical trial. *J. Amer. Statist. Assoc.* **64** 131–146.
- [45] ZHANG, L. and ROSENBERGER, W. F. (2006). Response-adaptive randomization for clinical trials with continuous outcomes. *Biometrics* **62** 562–569.
- [46] ZHANG, L.-X. (2016). Central limit theorems of a recursive stochastic algorithm with applications to adaptive designs. *Ann. Appl. Probab.* **26** 3630–3658. [MR3582813](#)
- [47] ZHANG, L.-X., HU, F. and CHEUNG, S. H. (2006). Asymptotic theorems of sequential estimation-adjusted urn models. *Ann. Appl. Probab.* **16** 340–369. [MR2209345](#)

- [48] ZHANG, L.-X., HU, F., CHEUNG, S. H. and CHAN, W. S. (2007). Asymptotic properties of covariate-adjusted response-adaptive designs. *Ann. Statist.* **35** 1166–1182. [MR2341702](#)
- [49] ZHANG, L.-X., HU, F., CHEUNG, S. H. and CHAN, W. S. (2011). Immigrated urn models—Theoretical properties and applications. *Ann. Stat.* **39** 643–671. [MR2797859](#)
- [50] ZHU, H. (2015). Covariate-adjusted response adaptive designs incorporating covariates with and without treatment interactions. *Canad. J. Statist.* **43** 534–553.

G. ALETTI
DEPARTMENT OF ENVIRONMENTAL
SCIENCE AND POLICY
UNIVERSITÀ DEGLI STUDI DI MILANO
VIA SALDINI 50
20133 MILAN
ITALY
E-MAIL: giacomo.aletti@unimi.it

A. GHIGLIETTI
DEPARTMENT OF STATISTICAL SCIENCE
UNIVERSITÀ CATTOLICA DEL SACRO CUORE
VIA LARGO GEMELLI 1
20123 MILAN
ITALY
E-MAIL: andrea.ghiglietti@unicatt.it

W. F. ROSENBERGER
DEPARTMENT OF STATISTICS
GEORGE MASON UNIVERSITY
4400 UNIVERSITY DRIVE MS4A7
FAIRFAX, VIRGINIA 22030
USA
E-MAIL: wrosenbe@gmu.edu