



Covariate-Adjusted Response-Adaptive Design with Delayed Outcomes

Xinwei Ma, Jingshen Wang & Waverly Wei

To cite this article: Xinwei Ma, Jingshen Wang & Waverly Wei (21 May 2026): Covariate-Adjusted Response-Adaptive Design with Delayed Outcomes, Journal of the American Statistical Association, DOI: [10.1080/01621459.2025.2604314](https://doi.org/10.1080/01621459.2025.2604314)

To link to this article: <https://doi.org/10.1080/01621459.2025.2604314>



View supplementary material [↗](#)



Published online: 21 May 2026.



Submit your article to this journal [↗](#)



Article views: 331



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 2 View citing articles [↗](#)



Covariate-Adjusted Response-Adaptive Design with Delayed Outcomes

Xinwei Ma^a , Jingshen Wang^b , and Waverly Wei^c 

^aDepartment of Economics, University of California San Diego, La Jolla, CA; ^bDivision of Biostatistics, University of California Berkeley, Berkeley, CA;

^cDepartment of Data Sciences and Operations, University of Southern California, Los Angeles, CA

ABSTRACT

Covariate-adjusted response-adaptive (CARA) designs have gained widespread adoption for their clear benefits in enhancing experimental efficiency and participant welfare. These designs dynamically adjust treatment allocations during interim analyses based on participant responses and covariates collected during the experiment. However, delayed responses can significantly compromise the effectiveness of CARA designs, as they hinder timely adjustments to treatment assignments when certain participant outcomes are not immediately observed. In this article, we propose a fully forward-looking CARA design that dynamically updates treatment assignments throughout the experiment as response delay mechanisms are progressively estimated. Our design strategy is informed by novel semiparametric efficiency calculations that explicitly account for outcome delays in a multi-stage setting. Through both theoretical investigations and simulation studies, we demonstrate that our proposed design offers a robust solution for handling delayed outcomes in CARA designs, yielding significant improvements in both statistical power and participant welfare. Supplementary materials for this article are available online, including a standardized description of the materials available for reproducing the work.

ARTICLE HISTORY

Received January 2025

Accepted December 2025

KEYWORDS

Delayed outcomes;
Frequentist adaptive
experimental design;
Response adaptive designs.

1. Introduction

1.1. Motivation and Contribution

In clinical trials, social science field experiments, and A/B tests, participants often enroll sequentially, and responses to treatments can vary due to individual-specific characteristics. In these contexts, covariate-adjusted response-adaptive (CARA) designs have gained increasing popularity: by dynamically adjusting treatment assignments based on accumulated covariates and outcome information during the experiment, CARA designs can be adopted to optimize treatment allocation to improve statistical power (Hahn, Hirano, and Karlan 2011; Hu, Zhu, and Hu 2015; Blackwell, Pashley, and Valentino 2023), reduce participant exposure to less effective treatments and enhance overall welfare (Hu and Rosenberger 2006; Hu, Zhu, and Hu 2015; Wei, Ma, and Wang 2023, 2025), and maintain statistical validity for analyzing experimental data (Zhang et al. 2007; Hu and Hu 2012; Bugni, Canay, and Shaikh 2018; Offer-Westort, Coppock, and Green 2021; Bai, Romano, and Shaikh 2022; Robertson et al. 2023; Zhao 2023; Wei, Ma, and Wang 2025; Bibaut and Kallus 2025).

Nevertheless, because CARA designs adjust treatment assignments based on observed participant responses, delays in outcome observation can significantly hinder their effectiveness. Intuitively, if primary outcomes for some participants remain unobserved at the time of interim analysis, the experimental designer may lack information needed to optimize treatment assignments for future participants with similar covariate

profiles. These delays often depend on both the treatment arm and participant covariates, as is common in sequentially enrolled randomized experiments.

As an example, we revisit the study by Fahey et al. (2020), a randomized experiment conducted in the Shinyanga region of Tanzania between April 24 and December 14, 2018. The study aimed to assess the impact of cash incentives on retention in care and viral suppression among people living with HIV (PLHIV). A total of 530 PLHIV aged 18 or older were sequentially enrolled and randomly assigned to either a treatment group receiving cash incentives or a control group without any incentive. Viral load was measured six months after treatment assignment through blood draws. However, delays occurred, as the lab test required a separate clinical visit and some participants missed their scheduled appointments. We calculated the number of days delayed as the difference between the date of the viral load test and the six-month follow-up date, separately for two covariate strata defined by biological sex. From Figure 1, it is clear that the distribution of outcome delays varies significantly depending on treatment assignment and gender. Specifically, the treatment arm exhibits less delay, and within the treatment arm, the delayed response issue seems less pronounced among female participants. Overall, the empirical observation highlights the importance of accounting for delayed outcomes that are both arm- and covariate-specific.

To our knowledge, although there has been prior work on developing valid statistical inference procedures in the presence of delayed responses in CARA (Bai, Hu, and Rosenberger 2002;




Figure 1. Delays in viral load blood test across treatment arms and covariate strata (biological sex).

Hu et al. 2008; Robertson et al. 2023), no CARA designs has incorporated the delay mechanism when optimizing experimental objectives. In this article, we address the challenges that response delays pose to covariate-adjusted response-adaptive (CARA) designs. We begin by examining whether existing optimal allocation strategies remain valid under delayed outcomes and, when they do not, derive the optimal allocation in such settings. We then develop methods to adapt CARA designs accordingly. Finally, we propose an efficient estimator that accounts for delayed responses at the end of the experiment. We elaborate on these contributions in more detail below.

First, we demonstrate that the classical allocation strategy used in existing CARA designs no longer provides optimal treatment assignment when accounting for arm- and covariate-dependent delays. Specifically, we first derive a semiparametric efficiency bound for treatment effect estimation in the presence of outcome delay (Theorem 1), which differ substantially from existing results in the literature (Hahn 1998; Hirano, Imbens, and Ridder 2003; Cattaneo 2010; Zhu and Zhu 2023; Wu et al. 2025). See van der Vaart (2000) and Tsiatis (2006) for general discussions on semiparametric efficiency. This result allows us to show that the optimal delay mechanism involves a unique trajectory of treatment assignment probabilities that vary across different stages, resulting in treatment allocations with lower variance for power maximization, and enhanced failure reduction in welfare improvement designs. To our knowledge, this is the first design to explicitly explore delay mechanisms to optimize trial objectives, while existing literature focuses on conducting sequential tests (Hampson and Jennison 2013), designing Bayesian adaptive designs where delay is independent of the arm and covariates (Lin, Thall, and Yuan 2020), or conducting statistical inference after experiments have concluded (Bai, Hu, and Rosenberger 2002; Hu et al. 2008).

Second, since the delay mechanism is unknown a priori and remains only partially estimable during the experiment, we propose a fully forward-looking CARA design that sequentially updates treatment assignment probabilities to target a trial objective at the conclusion of the experiment. In this process, our algorithm offers several extrapolation strategies when estimating delay mechanism, allowing better alignment with the designer’s prior knowledge. We further discuss in detail two design objectives frequently adopted in current CARA literature: maximizing experiment power (Tymofeyev, Rosenberger, and Hu 2007; Zhao 2023) and improving overall participant welfare (or failure reduction) while maintaining statistical power constraints (Rosenberger et al. 2001; Robertson et al. 2023). We justify the benefits of our proposed forward-looking CARA design through theoretical investigation (Section 4) and a synthetic case study (Section 5). Under the design objective of maximizing experiment power, our proposed forward-looking CARA design attains higher estimation efficiency compared to the Neyman allocation and complete randomization designs in the presence of delayed responses. Under the design objective of improving overall participant welfare, we demonstrate that our proposed forward-looking CARA design achieves additional failure reduction, leading to greater welfare improvement compared to both the ethical design and the complete randomization design while maintaining statistical power. For both design objectives, our proposed design can provide valid statistical inference, which is justified in Theorem 3.

Third, we make theoretical contributions to the literature on CARA and adaptive experimental designs by establishing general conditions under which estimated treatment effects are consistent and asymptotically normally distributed (Theorem 2). Specifically, we do not restrict the potential outcomes or the

delay mechanism to any parametric family of distributions, thereby alleviating the burden of choosing a particular set of parametric assumptions. Our theoretical development builds on martingale methods (Hall and Heyde 1980) and a high-level condition (Condition 1) that the optimized allocations converge in large samples, an assumption satisfied by many adaptive experimental designs. As a result, the general consistency and asymptotic normality findings may be of independent interest. We then specialize in our proposed design and verify the design consistency condition (Section 4.3), providing a suite of tools for designing and analyzing adaptive experiments with delayed outcomes, along with rigorous statistical guarantees.

1.2. Setup: CARA with Arm- and Covariate-Dependent Delays

In the remainder of this introduction, we lay out the setup of CARA when responses are subject to delay. Formal statements of our assumptions are collected in Section 4. The experiment is conducted over T stages, labeled $t = 1, 2, \dots, T$. In each stage t , n_t participants are enrolled. Participants are indexed by $i = 1, 2, \dots, N$, with $N = \sum_{t=1}^T n_t$ being the total sample size. We use N_t to denote the cumulative sample size at the end of stage t , that is, $N_t = \sum_{s=1}^t n_s$. Therefore, $N = N_T$, where we omit the subscript to save notation. Finally, $S_i = t$ denotes that individual i is enrolled in stage t .

Upon enrollment, baseline covariates $X_i \in \mathcal{X}$ are collected for each participant, and they are then randomized into one of the treatment arms, denoted by $A_i \in \mathcal{A} = \{0, 1\}$, which represents their actual treatment status. Following the Neyman-Rubin causal model, the potential outcomes are denoted by $Y_i(0)$ and $Y_i(1)$. In line with existing literature in adaptive design, we are in a scenario where the potential outcomes and the covariates, $(Y_i(0), Y_i(1), X_i)$, are independent and identically distributed across different stages. For future reference, we also introduce notation for mean potential outcomes: $\mu(x, a) = \mathbb{E}[Y_i(a)|X_i = x]$ and $\mu(a) = \mathbb{E}[Y_i(a)]$, and the conditional and unconditional treatment effects are defined as $\tau(x) = \mu(x, 1) - \mu(x, 0)$ and $\tau = \mu(1) - \mu(0)$. Other moments of the potential outcomes are $\mu_s(x, a) = \mathbb{E}[Y_i(a)^s|X_i = x]$ and $\mu_s(a) = \mathbb{E}[Y_i(a)^s]$, but we omit the subscript whenever $s = 1$ for ease of notation.

Since outcome information may not be immediately available after treatment assignment, we use D_i to represent the number of stages after which the experimenter can observe Y_i . This means that the participant outcome Y_i is observed at the end of stage t if and only if $D_i + S_i \leq t$. To give a concrete example, Let $S_i = 1$ so the participant is enrolled in the first stage. Then, their outcome information is available at the end of stage 3 if and only if the delay is no more than two stages, which means $D_i \leq 2$, or equivalently, $S_i + D_i \leq 3$. Also note that when $D_i = 0$, it implies no delay in the outcome for this individual.

The delay distribution is denoted by $\rho(d|x, a) = \mathbb{P}[D_i \leq d|X_i = x, A_i = a]$. As the notation suggests, we allow the delay mechanism D_i to depend on the covariates and the treatment assignment status. We will assume delay is conditionally independent of the potential outcomes: $D_i \perp\!\!\!\perp$

$(Y_i(0), Y_i(1), S_i) \mid (X_i, A_i)$. As the delay distribution does not shift over time, it allows us to sequentially learn the delay mechanism using accrued data and to optimize treatment allocation.

In CARA designs, the treatment assignment is sequentially updated based on accumulated data to achieve a pre-specified design objective. Specifically, at stage t , suppose we have collected historical participant information, denoted as $\mathcal{H}_t = \{(A_i, X_i, Y_i, D_i)\}_{i=1}^{N_t}$, where we recall that the sample indexed by $i = 1, 2, \dots, N_t$ consists of all participants enrolled in the first t stages. The treatment assignment for the next stage, $t + 1$, denoted as $\hat{e}_{t+1}^*(a|x)$, is determined by both this historical information and the covariate information of the newly enrolled participants in Stage $t + 1$. Formally, this is captured by $\mathbb{P}[A_i = a|X_i = x, S_i = t + 1, \mathcal{H}_t] = \hat{e}_{t+1}^*(a|x)$, which means the treatment assignment adapts to $\mathcal{H}_t$.

In this manuscript, we work under the setting that T and n_t are pre-fixed. We also hope to note that, in general adaptive trial designs, determining the appropriate sample size and the number of stages involves balancing statistical rigor, operational feasibility, and ethical considerations. Before an experiment starts, practitioners often pre-specify the desired power and significance level, along with the expected treatment effect size based on prior studies or clinical knowledge, to compute an initial overall sample size. Then, the number of stages is selected by evaluating logistical considerations (e.g., recruitment rate, data availability, cost constraints). It is often the case that fewer stages are used to simplify logistics and reduce operational complexity, while more stages increase flexibility but impose higher management costs. Once the number of stages is decided, sample size allocation per stage can be determined either equally or unequally (e.g., adaptive design with small initial pilot study) based on anticipated interim analyses. It is also common practice to conduct simulations during the planning stage to evaluate various stage/sample size configurations, ensuring feasibility and ethical appropriateness of the chosen adaptive strategy. Our design thus aligns with the case where there is a finite number of stages and each stage contains an equal/unequal number of participants.

2. Optimal Delay-Adjusted Treatment Allocation in CARA

While the design goals of CARA vary across applications, maintaining sufficient statistical power to assess the treatment's effectiveness on the primary outcome remains a central concern. A statistically efficient estimator of the treatment effect also supports a robust adaptive design, as the estimated effects guide the sequential revision of the treatment allocation strategy. However, it remains unclear what constitutes a good estimator in the presence of delayed outcomes. Moreover, the semiparametric efficiency bound for estimating the treatment effect in multi-stage experiments with delayed responses has not yet been established. To provide practical guidance for designing response-adaptive experiments, we begin by presenting a new result on the semiparametric efficiency bound for estimating the average treatment effect:

$$\int_{x \in \mathcal{X}} p(x) \left[\frac{\sigma^2(x, 1)}{\sum_{t=1}^T r_t \rho(T - t|x, 1) e_t(1|x)} + \frac{\sigma^2(x, 0)}{\sum_{t=1}^T r_t \rho(T - t|x, 0) e_t(0|x)} + (\tau(x) - \tau)^2 \right] dx, \quad (1)$$

where $r_t = \mathbb{P}[S_i = t]$ being the fraction of participants enrolled in stage t , $p(x)$ is the probability density function of the covariate X , and $\sigma^2(x, a)$ is the conditional variance of the potential outcome $Y_i(a)$. See Theorem 1 in Section 4 for a rigorous statement.

To intuitively understand (1), it helps to first revisit the classical semiparametric efficiency bound in a static (one-period) setting without delays, which is given by:

$$\int_{x \in \mathcal{X}} p(x) \left[\frac{\sigma^2(x, 1)}{e(1|x)} + \frac{\sigma^2(x, 0)}{e(0|x)} + (\tau(x) - \tau)^2 \right] dx. \quad (2)$$

In this classical formulation, the denominators (i.e., propensity scores $e_t(a|x)$) adjust for covariate-specific treatment assignment. The key difference between (2) and (1) is the introduction of an additional term $\rho(T - t|x, a)$ in the denominators, which accounts for the probability distribution of response delays. Intuitively, delays introduce extra uncertainty or missingness in observing outcomes. As we operate in a multiple-stage experimental setting, we aggregate these delay-related probabilities across all stages, appropriately weighted by each stage's sample size captured by r_t . Hence, the revised efficiency bound appropriately adjusts for both the random assignment of treatments (captured by propensity scores) and the randomness arising from delays in observing responses.

From the above result, it is evident that when the delay mechanism of the response depends on the covariates and treatment status, the optimal treatment allocation rules in classical CARA designs must be modified to account for outcome delays. In what follows, we examine how the optimal treatment allocation rule is affected by the arm- and/or covariate-dependent delayed responses in two common experimental objectives: one aimed at enhancing statistical power (Tymofeyev, Rosenberger, and Hu 2007), and the other focused on minimizing the expected number of failures while maintaining a minimum power requirement for binary outcomes (Rosenberger et al. 2001). In both design goals, the variance of the estimated treatment effects plays an essential role. To simplify the presentation, we focus on two-arm CARA designs with $a \in \{0, 1\}$ in this article. For multi-arm experiments, the design objective can be generalized using non-centrality parameters, as discussed in Lachin (1977) and Tymofeyev, Rosenberger, and Hu (2007).

Design objective 1: Optimal treatment allocation for power maximization. The objective is to sequentially assign treatments to minimize the variance of the estimated average treatment effect (ATE), thereby maximizing statistical power. The optimal allocation is the solution to the optimization problem:

$$\begin{aligned} \min_{e_1(\cdot), \dots, e_T(\cdot)} V(e_1(\cdot), \dots, e_T(\cdot)) \quad \text{subject to} \\ \delta \leq e_t(\cdot) \leq 1 - \delta \quad \text{for } t = 1, \dots, T, \end{aligned} \quad (3)$$

where δ ensures that the assignment probabilities are bounded away from 0 and 1, maintaining sufficient randomness in treatment assignments (Ma and Wang 2020; Heiler and Kazak 2021; Sasaki and Ura 2022; Ma, Sasaki, and Wang 2025; Dorn 2025). The objective function $V(e_1(\cdot), \dots, e_T(\cdot))$ is a component of the semiparametric efficiency bound related to the propensity scores, and is defined as

$$V(e_1(\cdot), \dots, e_T(\cdot)) = \sum_{x \in \mathcal{X}} p(x) \left[\frac{\sigma_t^2(x, 1)}{\sum_{t=1}^T r_t \rho(T - t|x, 1) e_t(x)} + \frac{\sigma_t^2(x, 0)}{\sum_{t=1}^T r_t \rho(T - t|x, 0) (1 - e_t(x))} \right].$$

In line with existing literature on CARA designs, we explicitly consider discrete covariates in the objective function, as most design applications will first group participants based on their covariate information. Therefore, the objective function takes a summation form with respect to the probability mass function.

Design objective 2: Optimal treatment allocation for failure reduction subject to a power constraint. This design aims to minimize the expected number of failures while satisfying a minimum power requirement for binary responses (Rosenberger et al. 2001). In this context, a treatment is considered to fail when the potential outcome equals one. The optimal allocation thus seeks to assign treatments based on the solution to the following optimization problem:

$$\begin{aligned} \min_{e_1(\cdot), \dots, e_T(\cdot)} \sum_{t=1}^T r_t \sum_{x \in \mathcal{X}} p(x) \\ \left[e_t(x) (1 - \mu(x, 1)) + (1 - e_t(x)) (1 - \mu(x, 0)) \right] \quad (4) \\ \text{subject to} \quad \delta \leq e_t(\cdot) \leq 1 - \delta, \\ V(e_1(\cdot), \dots, e_T(\cdot)) + \sum_{x \in \mathcal{X}} p(x) (\tau(x) - \tau)^2 \leq C. \end{aligned}$$

In this formulation, the objective function minimizes the total expected number of failures during the entire experiment. The term inside the summation accounts for the expected failures when assigning treatment 1 with probability $e_t(x)$ and treatment 0 with probability $1 - e_t(x)$. Here, $\mu(x, a)$ represent the expected success probabilities under treatments a , given covariates x . The first feasibility constraints again ensure that the treatment assignment probabilities remain within the range $[\delta, 1 - \delta]$. This design balances the ethical imperative to reduce adverse outcomes with statistical consideration to maintain sufficient power for detecting treatment effects.

Due to the delay mechanism's influence on the variance lower bound, the optimal treatment allocation rules differ substantially from those derived in the existing literature. We shall show that the solutions to the optimization problems in (3) and (4), denoted as $e_1^*(\cdot), \dots, e_T^*(\cdot)$, form a sequence of probabilities that vary across different stages. In comparison, the classical Neyman allocation rule for power maximization:

$$e_t^{\text{Neyman}}(a|x) = \frac{\sigma(x, a)}{\sigma(x, 1) + \sigma(x, 0)}, \quad \text{for } t = 1, \dots, T, \quad (5)$$




Figure 2. Optimal oracle allocation (accounting for outcome delays) and the Neyman allocation for the power maximization objective.

and the failure reduction rule of Rosenberger et al. (2001):

$$e_t^{\text{FR}}(a|x) = \frac{\sqrt{\mu(x, a)}}{\sqrt{\mu(x, 1)} + \sqrt{\mu(x, 0)}}, \quad \text{for } t = 1, \dots, T, \quad (6)$$

are both fixed probabilities that do not differ across stages. Both allocation methods ignore the delay mechanism. Therefore, when response delay plays a significant role, assigning treatments based on existing optimal rules no longer guarantees achieving the desired experimental goals.

We illustrate this point more concretely by computing the optimal oracle treatment allocation in the example introduced in Section 1, and compare it with classical treatment assignment methods such as the Neyman allocation. The covariate strata we consider are formed by two covariates, the biological sex and the WHO clinical stages of HIV. See Section 5 for additional details on the model parameters. In Figure 2, we present the oracle allocations for each stage and each covariate stratum corresponding to the power maximization objective defined in (3), as well as the Neyman allocation. Note that, since the classical Neyman allocation ignores the delayed outcome issue, it results in a constant treatment probability across the experimental stages, whereas the oracle allocation varies significantly across both stages and covariate strata. This highlights the importance of allowing the delay mechanism to be both arm- and covariate-specific.

We also report the objective function value for each design method in the figure. As can be seen, the efficiency gain in the oracle allocation can be substantial: more than 10% in the first covariate stratum and about 8% in the last case. Once again, this numerical exercise demonstrates the potential gains from leveraging the delay mechanism in a CARA framework.

3. Fully Forward-looking CARA with Delayed Outcomes

The results provided in the previous section assume prior knowledge of the delay mechanism, and therefore, the discussed optimal treatment allocation corresponds to an infeasible oracle setting. In practice, the delay mechanism is rarely known and must be estimated as data are sequentially accumulated to optimize the experimental objective. In this section, we propose two delay-informed practical CARA designs tailored to two

previously discussed experimental goals: improving statistical power (Tymofeyev, Rosenberger, and Hu 2007) and minimizing the expected number of failures while meeting a minimum power requirement for binary responses (Rosenberger et al. 2001). We also provide tools for valid statistical inference after the experiment is finished.

As with any CARA design, we use a pre-specified and fixed treatment assignment rule for the first stage due to the lack of a priori knowledge about the covariate and outcome distributions as well as the delay mechanism. In subsequent stages, our proposed fully forward-looking CARA design incorporates two key components: (i) estimating aspects of the delay mechanism from accrued data, while employing extrapolation to assess the tail behavior of the delay distribution; and (ii) using these estimates to formulate a feasible, fully forward-looking optimization problem that determines the optimal treatment assignments for the next stage.

We first discuss the estimation of the delay mechanism. Assuming stage t of the experiment has concluded, we first update our delay distribution estimator via

$$\begin{aligned} \hat{\rho}_t(d|x, a) &= \sum_{\ell=0}^d \hat{\mathbb{P}}[D_i = \ell | X_i = x, A_i = a] \\ &= \sum_{\ell=0}^d \frac{\sum_{i=1}^N \mathbb{1}_{(X_i=x, A_i=a, D_i=\ell, S_i \leq t-\ell)}}{\sum_{i=1}^N \mathbb{1}_{(X_i=x, A_i=a, S_i \leq t-\ell)}}, \quad d \leq t-1, \quad (7) \end{aligned}$$

for $a = 0, 1$. Notice that the above estimator is only defined for $d \leq t-1$. To gain some intuition, consider a concrete example for $t = 3$. At the end of the third stage, data collected from participants enrolled in the first stage will be informative about the delay mechanism $\rho(0|\cdot)$, $\rho(1|\cdot)$, and $\rho(2|\cdot)$. Among those whose outcome information is still missing at the end of stage 3, it is not possible to tell whether such information will become available until the end of stage 4. Following a similar reasoning, participants enrolled in the second stage can help estimate $\rho(0|\cdot)$ and $\rho(1|\cdot)$. The main message is that at any stage t , the delay mechanism is only estimable up to $d \leq t-1$, and the remaining “tail features” are not estimable until more data is collected.

To anticipate these future delay distributions, we propose the following approaches that reflect the experimenter’s expecta-

tions about future delays. For a *conservative* experimenter, we consider a simple extrapolation approach:

$$\hat{\rho}_t(d|x, a) = \hat{\rho}_t(t-1|x, a), \quad d \geq t. \quad (8)$$

This extrapolation essentially assumes that all future delay probabilities will remain at the last estimated value. As we discuss below in Section 4.3, the conservative extrapolation can be understood from a minimax perspective, as we will be sequentially optimizing the objective function under the worst-case scenario of the delay mechanism. For this reason, we adopt this approach in most of our numerical experiments and theoretical investigations.

For an *optimistic* experimenter, we consider an extrapolation that postulates all missing outcomes will be available after one more experimental stage: $\hat{\rho}_t(d|x, a) = 1, d \geq t$. Lastly, a *neutral* experimenter might gradually adjust the estimated delay probabilities from the last observed value to immediate response by interpolating between $\hat{\rho}_t(t-1|x, a)$ and 1 for $d = t, \dots, T-1$. These strategies enable experimenters to incorporate future delay probabilities into our design, accommodating different perspectives on how delays may evolve. Although it is generally impossible to rank the three extrapolation methods, we show through Monte Carlo experiments that they all deliver valid inference for the desired causal effects and more effectively achieves the design objectives compared to classical approaches (such as complete randomization).

Next, to optimize treatment allocation for stage $t+1$ while accounting for response delays, we introduce a fully forward-looking estimation of the variance term $V(e_1(\cdot), \dots, e_T(\cdot))$ at the current stage t . Our proposed estimate incorporates past allocations up to Stage t , future allocations from stages $t+1$ to T , and the estimated delay mechanisms. Our method is designed to anticipate the impact of future allocations and outcome delays on the variance of the final treatment effect estimate after the experiment concludes:

$$\begin{aligned} \hat{V}_t(e_{t+1}(\cdot), \dots, e_T(\cdot)) &= \sum_{x \in \mathcal{X}} \hat{p}_t(x) \left[\frac{\hat{\sigma}_t^2(x, 1)}{\sum_{\ell=1}^t \hat{r}_\ell \hat{\rho}_t(T-\ell|x, 1) \hat{e}_\ell^*(1|x) + \sum_{\ell=t+1}^T \hat{r}_\ell \hat{\rho}_t(T-\ell|x, 1) e_\ell(x)} \right. \\ &\quad \left. + \frac{\hat{\sigma}_t^2(x, 0)}{\sum_{\ell=1}^t \hat{r}_\ell \hat{\rho}_t(T-\ell|x, 0) \hat{e}_\ell^*(0|x) + \sum_{\ell=t+1}^T \hat{r}_\ell \hat{\rho}_t(T-\ell|x, 0) (1 - e_\ell(x))} \right]. \end{aligned}$$

Here, $\hat{p}_t(x) = N_t^{-1} \sum_{i=1}^N \mathbb{1}_{(X_i=x, S_i \leq t)}$, which is the estimated proportion of participants with covariate x up to stage t . The term $\hat{r}_\ell = n_\ell/N$ denotes the proportion of participants enrolled at stage ℓ . The estimated conditional mean and variance of the outcome $Y_i(a)$ are

$$\begin{aligned} \hat{\sigma}_t^2(x, a) &= \frac{\sum_{i=1}^N \mathbb{1}_{(X_i=x, A_i=a, D_i+S_i \leq t)} (Y_i - \hat{\mu}_t(x, a))^2}{\sum_{i=1}^N \mathbb{1}_{(X_i=x, A_i=a, D_i+S_i \leq t)}}, \\ \hat{\mu}_t(x, a) &= \frac{\sum_{i=1}^N \mathbb{1}_{(X_i=x, A_i=a, D_i+S_i \leq t)} Y_i}{\sum_{i=1}^N \mathbb{1}_{(X_i=x, A_i=a, D_i+S_i \leq t)}}. \end{aligned} \quad (9)$$

While the expressions may seem complicated, they are simply the sample mean and sample variance of the outcome, restricted to the subsample for group x receiving treatment a . We remark that while the true underlying conditional means $\mu(x, a)$ and conditional variances $\sigma^2(x, a)$ do not depend on a specific stage t , their estimate are stage-specific. This is due to the sequential nature of the experiment: we always update previous estimates whenever new data become available.

Using this fully forward-looking variance estimate, we can then determine an effective treatment assignment for stage $t+1$ that accounts for how future treatment allocations and response delays from stages $t+1$ onward affect the design objectives at the end of the experiment. For the first design objective of maximizing power, we solve the following optimization problem to find the treatment allocation probability for stage $t+1$:

$$(\tilde{e}_{t+1}(\cdot), \dots, \tilde{e}_T(\cdot)) = \underset{e_{t+1}(\cdot), \dots, e_T(\cdot)}{\operatorname{argmin}} \quad \hat{V}_t(e_{t+1}(\cdot), \dots, e_T(\cdot)), \quad (10)$$

subject to the constraint that $\delta \leq e_\ell(\cdot) \leq 1 - \delta$ for $\ell \geq t+1$. We then set $\hat{e}_{t+1}^*(1|\cdot) = 1 - \hat{e}_{t+1}^*(0|\cdot) = \tilde{e}_{t+1}(\cdot)$.

For the second design objective of reducing failures under a power constraint, we define a similar fully forward-looking estimate of the expected proportion of total failures as

$$\begin{aligned} \hat{P}_t(e_{t+1}(\cdot), \dots, e_T(\cdot)) &= \sum_{\ell=1}^t \hat{r}_\ell \sum_{x \in \mathcal{X}} \hat{p}_t(x) \\ &\quad \left[\hat{e}_\ell^*(1|x) (1 - \hat{\mu}_t(x, 1)) + \hat{e}_\ell^*(0|x) (1 - \hat{\mu}_t(x, 0)) \right] \\ &\quad + \sum_{\ell=t+1}^T \hat{r}_\ell \sum_{x \in \mathcal{X}} \hat{p}_t(x) \\ &\quad \left[e_\ell(x) (1 - \hat{\mu}_t(x, 1)) + (1 - e_\ell(x)) (1 - \hat{\mu}_t(x, 0)) \right]. \end{aligned}$$

We denote the solution to the following optimization problem as $(\tilde{e}_{t+1}(\cdot), \dots, \tilde{e}_T(\cdot))$:

$$\begin{aligned} \min_{e_{t+1}(\cdot), \dots, e_T(\cdot)} \quad & \hat{P}_t(e_{t+1}(\cdot), \dots, e_T(\cdot)) \\ \text{s.t.} \quad & \hat{V}_t(e_{t+1}(\cdot), \dots, e_T(\cdot)) \\ & + \sum_{x \in \mathcal{X}} \hat{p}_t(x) (\hat{\tau}_t(x) - \hat{\tau}_t)^2 \leq C, \end{aligned} \quad (11)$$

where $\hat{\tau}_t(x) = \hat{\mu}_t(x, 1) - \hat{\mu}_t(x, 0)$ and $\hat{\tau}_t = \sum_{x \in \mathcal{X}} \hat{p}_t(x) \hat{\tau}_t(x)$. We then again set $\hat{e}_{t+1}^*(1|\cdot) = 1 - \hat{e}_{t+1}^*(0|\cdot) = \tilde{e}_{t+1}(\cdot)$.

As both optimization problems in (10) and (11) are fully forward-looking, our proposed design strategies aim to proactively adjust the allocation strategies in future stages, ensuring that the data collected accounts for the future impacts due to outcome delays. This approach contrasts with classical CARA designs that optimize trial objectives solely using historical information. At the end of the last stage T , the average treatment effect estimator is constructed with

$$\hat{\tau}_T = \sum_{x \in \mathcal{X}} \hat{P}_T(x) \hat{\tau}_T(x) = \sum_{x \in \mathcal{X}} \hat{P}_T(x) (\hat{\mu}_T(x, 1) - \hat{\mu}_T(x, 0)).$$

We then provide a variance estimator, thus enabling valid statistical inference:

$$\hat{V}_T = \sum_{x \in \mathcal{X}} \hat{p}_T(x) \left[\frac{\hat{\sigma}_T^2(x, 1)}{\hat{e}_T(x, 1)} + \frac{\hat{\sigma}_T^2(x, 0)}{\hat{e}_T(x, 0)} + (\hat{\tau}_T(x) - \hat{\tau})^2 \right],$$

$$\hat{e}_T(x, a) = \frac{\sum_{i=1}^N \mathbb{1}_{(X_i=x, A_i=a, D_i+S_i \leq T)}}{\sum_{i=1}^N \mathbb{1}_{(X_i=x)}}.$$

We summarize the analysis protocol below, where $z_{1-\frac{\alpha}{2}}$ denotes the $1 - \frac{\alpha}{2}$ percentile of the standard normal distribution.

Protocol for interim analysis and statistical inference

Initial stage 1:

Enroll n_1 participants, and assign treatments with $\hat{e}_1^*(1|x) = \frac{1}{2}$.

Interim analysis after stage t has concluded:

Available data: $(S_i, X_i, A_i, D_i, Y_i)$ for $i = 1, 2, \dots, N_t$;

Following (7) and (8), construct the delay distribution estimate $\hat{\rho}_t(d|x, a)$;

Following (9), construct the outcome distribution estimates $\hat{\sigma}_t(x, a)$ and $\hat{\mu}_t(x, a)$;

From either (10) or (11), solve the allocation for stage $t+1$: $\hat{e}_{t+1}^*(1|x)$.

Treatment assignment for stage $t+1$:

Enroll n_{t+1} participants, and assign treatments with $\hat{e}_{t+1}^*(1|x)$.

Statistical inference after the final stage T has concluded:

Available data: $(S_i, X_i, A_i, D_i, Y_i)$ for $i = 1, 2, \dots, N$;

Construct the treatment effect estimator $\hat{\tau}_T$ and the variance estimator $\hat{V}_T$;

Report the $(1 - \alpha)\%$ confidence interval: $\hat{\tau}_T \pm z_{1-\frac{\alpha}{2}} \sqrt{\hat{V}_T/N}$.

4. Theoretical Investigation

This section provides theoretical justifications for the proposed CARA design. To begin, we present the semiparametric efficiency bound for treatment effect estimation in a multi-stage setting with delayed outcomes. We then develop a general result on the consistency and asymptotic normality of the estimated average treatment effect. Our theoretical investigation then proceeds by examining the specific setting considered in Section 3, where we extrapolate the estimated delay mechanism in each stage using a conservative rule. Finally, we observe that in an experiment with only finitely many stages, it is generally not possible to achieve the oracle allocation, as this would require perfect ex ante knowledge of the delay mechanism. Therefore, the last part of our theoretical investigation addresses the efficiency loss resulting from extrapolating the delay mechanism in our design method.

4.1. Notation, Assumptions, and the Efficiency Bound

We begin by recalling our adaptive experiment setting: the experiment consists of T stages, and in each stage $t = 1, 2, \dots, T$, n_t participants are enrolled. Let $N_t = \sum_{s=1}^t n_s$ represent the cumulative sample size at the end of stage t , with $N = N_T$ being the total sample size. We use the notation $S_i = t$ to indicate that individual i is enrolled in stage t . For each individual, we observe some baseline covariates $X_i \in \mathcal{X}$. The

observed outcome variable is denoted by Y_i , and it is related to the potential outcomes through $Y_i = A_i Y_i(1) + (1 - A_i) Y_i(0)$, where $A_i \in \mathcal{A}$ denotes the actual treatment assigned and $\mathcal{A} = \{0, 1\}$. Moments of the potential outcomes are $\mu_s(x, a) = \mathbb{E}[Y_i(a)^s | X_i = x]$ and $\mu_s(a) = \mathbb{E}[Y_i(a)^s]$; we omit the subscript whenever $s = 1$.

Our first assumption imposes standard regularity conditions on the covariates and the potential outcomes.

Assumption 1 (Covariates and potential outcomes). (i) The covariates and potential outcomes, $(X_i, Y_i(0), Y_i(1))$, are independent and identically distributed across $i = 1, 2, \dots, N$. (ii) The covariates have finite support (i.e., $|\mathcal{X}| < \infty$). Let $p(x) = \mathbb{P}[X_i = x]$, then $\min_{x \in \mathcal{X}} p(x) > 0$. (iii) The potential outcomes have finite fourth moments: $\max_{x \in \mathcal{X}} \max_{a \in \mathcal{A}} \mu_4(x, a) < \infty$.

The next assumption addresses the delay mechanism, where we adopt the notation D_i to indicate the number of stages after which the experimenter can observe Y_i . By our definition, Y_i is observed at the end of stage t if and only if $D_i + S_i \leq t$. Also recall that the delay distribution is represented by $\rho(d|x, a) = \mathbb{P}[D_i \leq d | X_i = x, A_i = a]$.

Assumption 2 (Delay mechanism). (i) Outcome delay is independent of the potential outcomes after conditioning on the covariates and the treatment assignment: $D_i \perp\!\!\!\perp \{Y_i(a) : a \in \mathcal{A}\} | X_i, A_i, S_i$. (ii) $\min_{x \in \mathcal{X}} \min_{a \in \mathcal{A}} \rho(0|x, a) > 0$.

Finally, we consider multi-stage experiment settings.

Assumption 3 (Asymptotic regime). The total number of stages, T , is fixed, and for all t , $n_t/N \rightarrow r_t > 0$.

We are now ready to present the semiparametric efficiency bound for average treatment effect estimation in the oracle setting, where one has perfect knowledge of the delay mechanism and treatments are randomly assigned conditional on the covariates. This efficiency bound serves as the foundation and starting point for our proposed adaptive experimental design, as our algorithm minimizes a feasible version of the efficiency bound constructed using accrued information.

Theorem 1 (Semiparametric efficiency bound). Let Assumptions 1–3 hold. In addition, assume (i) S_i are independent and identically distributed with $\mathbb{P}[S_i = t] = r_t$; (ii) the treatment assignment probabilities, $e_t(a|x) = \mathbb{P}[A_i = a | X_i = x, S_i = t]$, are bounded away from 0 and 1. Then the efficient influence function for estimating $\mu(a)$ is

$$\begin{aligned} \psi(X_i, A_i, D_i, Y_i, S_i | a) &= \frac{\mathbb{1}_{(A_i=a, D_i+S_i \leq T)}}{\sum_{t=1}^T r_t \rho(T-t | X_i, a) e_t(a | X_i)} (Y_i(a) - \mu(X_i, a)) \\ &\quad + \mu(X_i, a) - \mu(a), \end{aligned}$$

where $a \in \mathcal{A} = \{0, 1\}$. In addition, the efficient influence function for estimating τ is

$$\begin{aligned} \psi(X_i, A_i, D_i, Y_i, S_i) &= \psi(X_i, A_i, D_i, Y_i, S_i | 1) \\ &\quad - \psi(X_i, A_i, D_i, Y_i, S_i | 0), \end{aligned}$$

leading to the semiparametric efficiency bound:

$$V = \sum_{x \in \mathcal{X}} p(x) \left[\frac{\sigma^2(x, 1)}{\sum_{t=1}^T r_t \rho(T - t|x, 1) e_t(1|x)} + \frac{\sigma^2(x, 0)}{\sum_{t=1}^T r_t \rho(T - t|x, 0) e_t(0|x)} + (\tau(x) - \tau)^2 \right].$$

We note that while the previous theorem is developed in the context of discrete covariates, it can be generalized to accommodate continuously distributed covariates (see, (1)).

4.2. Asymptotic Properties of Treatment Effects Estimates

In this subsection, we demonstrate that the treatment effect estimator is consistent and admits an asymptotic normal distribution. The main conclusions of this section rely on a high-level condition regarding the convergence of the optimized treatment allocation rule, making these conclusions applicable to a wide range of CARA and a broad class of frequentist adaptive experimental design settings (Hu and Zhang 2004; Antognini and Zagoraiou 2015; Hu, Zhu, and Hu 2015). We begin with this high-level condition, which will be verified in the next subsection for our proposed forward-looking CARA design.

Condition 1 (Convergence of the optimized treatment allocation). There exists nonrandom $e_t(\cdot|\cdot)$, such that $\hat{e}_t^*(a|x) = e_t(a|x) + o_p(1)$ for all $a \in \mathcal{A}$ and all $x \in \mathcal{X}$.

The next theorem characterizes the asymptotic normality of both the estimated subgroup treatment effects and the estimated average treatment effect.

Theorem 2 (Asymptotic normality of estimated treatment effects). Let Assumptions 1–3 hold. Then under Condition 1,

$$\begin{aligned} & \sqrt{N}(\hat{\tau}_T(x) - \tau(x)) \\ & \xrightarrow{d} \mathcal{N}\left(0, \frac{1}{p(x)} \left(\frac{\sigma^2(x, 1)}{\sum_{t=1}^T r_t \rho(T - t|x, 1) e_t(1|x)} + \frac{\sigma^2(x, 0)}{\sum_{t=1}^T r_t \rho(T - t|x, 0) e_t(0|x)} \right) \right), \end{aligned}$$

and $\sqrt{N}(\hat{\tau}_T - \tau) \xrightarrow{d} \mathcal{N}(0, V)$, where V is defined in Theorem 1.

It is helpful to compare this theorem to the semiparametric efficiency bound established in the previous subsection. While the two results, Theorems 1 and 2, yield the same variance formula, they are conceptually very different. The semiparametric efficiency bound is derived under a fixed treatment assignment regime, which we then use to guide our design algorithms—either by minimizing the efficiency bound for power improvement or by maximizing failure reduction subject to a variance upper bound. However, from an ex ante perspective, it may not be immediately clear why this is a useful exercise, as it is not

obvious how the bound could actually be achieved. Encouragingly, Theorem 2 helps close the loop by showing that the asymptotic variance of the estimated treatment effect matches the bound in Theorems 1 under mild regularity conditions, thereby justifying our design objective. That said, it is worth clarifying that, in general, one cannot achieve the optimized oracle efficiency bound, an issue we discuss further below.

Theorem 3 (Statistical inference). Let Assumptions 1–3 hold. Then under Condition 1, $\hat{V}_T = V_T + o_p(1)$.

4.3. Convergence of the Optimized Treatment Allocation

Previously we have shown that the treatment effect estimators are consistent and asymptotically normally distributed. These results build on the high-level condition that the empirical treatment allocation, $\hat{e}_t^*(\cdot)$, will converge in large samples (Condition 1). In this subsection, we further investigate the theoretical properties of the proposed experimental design, specifically, solved from (10), and show that it satisfies Condition 1, thereby tying all loose ends. To give a road map, we will first show that the asymptotic analogs of our design algorithm have unique solutions. In this process, we will carefully distinguish the oracle problem of minimizing the semiparametric efficiency bound (which requires perfect knowledge of the delay mechanism) from the problem of sequential optimization using only accrued knowledge (which requires extrapolating the estimated delay mechanism). Next, we show that the empirically optimized treatment allocation converges to some large-sample limit. To conclude this subsection, we will present a bound on efficiency loss due to extrapolating the delay mechanism.

In our theoretical investigation, we will focus on the conservative method for estimating the delay mechanism. This allows us to provide simple and easy-to-interpret conditions, hence avoiding lengthy discussions. Another appealing property of the conservative approach is that it can be understood from a “minimax” perspective, where we sequentially minimize the asymptotic variance under the worst-case scenario of the delay mechanism. Despite our focus on one particular approach in this subsection, we show in our simulations that all three methods (conservative, optimistic, and neutral) deliver valid inference.

To start, we recall that $V(e_1(\cdot), \dots, e_T(\cdot))$ is the objective function in (3). Its minimizers will be denoted by $e_t^*(a|x)$ for $a \in \{0, 1\}$, $x \in \mathcal{X}$, and $t = 1, 2, \dots, T$. As we have discussed, the oracle optimal solution is not achievable in general, as it builds on perfect knowledge of the delay mechanism. In any realistic experimental setting, however, this information can only be learned sequentially. In fact, even we set $n_t = \infty$, part of the delay mechanism can still remain unknown for any t : to be even more precise, one typically only learn $\rho(d|x, a)$ for $d < t$ using information at the end of stage t . For this reason, we have introduced feasible and concrete ways to extrapolate the delay mechanism. Following our discussion in Section 2, we define $\rho_t^\dagger(d|x, a) = \min \{\rho(d|x, a), \rho(t-1|x, a)\}$.

It is helpful to compare this definition with Assumption 2. There, we assumed that the delay mechanism is time-homogeneous, so that $\rho(d|x, a)$ is not indexed by t . On the other hand, $\rho_t^\dagger(d|x, a)$ arises due to our imperfect knowledge

about the delay mechanism and the conservative extrapolation employed, and as a result it will depend on a specific stage. It is not the true delay mechanism, but it is the asymptotic analogue of $\hat{\rho}_t(d|x, a)$, which we used in our optimization algorithm. To study convergence property of $\hat{e}_t^*(a|x)$, we will make a definition first. Let $e_1^\dagger(0|x) = e_1^\dagger(1|x) = 1/2$. Then for $t = 1, 2, \dots, T$, define recursively that

$$\begin{aligned} e_{t+1}^\dagger(1|\cdot) &= 1 - e_{t+1}^\dagger(0|\cdot) \\ &= \operatorname{argmin}_{e_{t+1}(\cdot)} \min_{e_{t+2}(\cdot), \dots, e_T(\cdot)} V_t^\dagger(e_{t+1}(\cdot), \dots, e_T(\cdot)), \end{aligned}$$

where

$$\begin{aligned} V_t^\dagger(e_{t+1}(\cdot), \dots, e_T(\cdot)) &= \sum_{x \in \mathcal{X}} p(x) \left[\frac{\sigma^2(x, 1)}{\sum_{\ell=1}^t r_\ell \rho_t^\dagger(T - \ell|x, 1) e_\ell^\dagger(1|x) + \sum_{\ell=t+1}^T r_\ell \rho_t^\dagger(T - \ell|x, 1) e_\ell(x)} \right. \\ &\quad \left. + \frac{\sigma^2(x, 0)}{\sum_{\ell=1}^t r_\ell \rho_t^\dagger(T - \ell|x, 0) e_\ell^\dagger(0|x) + \sum_{\ell=t+1}^T r_\ell \rho_t^\dagger(T - \ell|x, 0) (1 - e_\ell(x))} \right]. \end{aligned}$$

Our first main result is to show that both $e_t^*(a|x)$ and $e_t^\dagger(a|x)$ are unique. To this end, we make the following assumption.

Assumption 4 (Variation in delay mechanism). For any $d \neq d'$, $\rho(d|x, 1)/\rho(d|x, 0) \neq \rho(d'|x, 1)/\rho(d'|x, 0)$.

We remark that this assumptions helps establish the uniqueness of $e_t^*(a|x)$. However, it is still possible to have $\rho_t^\dagger(d|x, 1)/\rho_t^\dagger(d|x, 0) = \rho_t^\dagger(d'|x, 1)/\rho_t^\dagger(d'|x, 0)$. Fortunately, this is easily resolved by a tie-breaking rule. See the supplementary materials for a specific recommendation and additional discussions.

Theorem 4 (Consistent treatment allocation). Let Assumptions 1–4 hold. Then for $a \in \{0, 1\}$, $x \in \mathcal{X}$, and $t = 1, 2, \dots, T$, both $e_t^*(a|x)$ and $e_t^\dagger(a|x)$ are unique. In addition, Condition 1 holds with $\hat{e}_t^*(a|x) \xrightarrow{P} e_t^\dagger(a|x)$.

Before closing this section, we provide two insights regarding the variance minimization problem. The first result suggests that as more information about the delay mechanism becomes available in the adaptive experiment, the optimized asymptotic variance always (weakly) decreases. While this result may seem natural, we note that it relies on the use of conservative extrapolation. As discussed earlier, conservative extrapolation can be understood from a minimax perspective, where the asymptotic variance is sequentially minimized under the worst-case scenario of the delay mechanism.

The second result provides a bound on efficiency loss due to imperfect knowledge about the delay mechanism. Collectively, these two results provide theoretical guarantee to our proposed adaptive experimental design.

Theorem 5 (Optimized variance and bound on efficiency loss). Let Assumptions 1–4 hold. Define V^* as the minimized oracle asymptotic variance, and V to be the asymptotic variance of the

estimated average treatment effect. Then (i) for all $1 \leq s < t \leq T - 1$,

$$\begin{aligned} V^* &\leq V \leq \min_{e_{t+1}(\cdot), \dots, e_T(\cdot)} V_t^\dagger(e_{t+1}(\cdot), \dots, e_T(\cdot)) \\ &\leq \min_{e_{s+1}(\cdot), \dots, e_T(\cdot)} V_s^\dagger(e_{s+1}(\cdot), \dots, e_T(\cdot)); \end{aligned}$$

and (ii) for any $\epsilon > 0$ and d_ϵ with $\min_{x \in \mathcal{X}, a \in \mathcal{A}} \rho(d_\epsilon|x, a) \geq 1 - \epsilon$, $V \leq V^* + C(d_\epsilon T^{-1} + \epsilon)$, where C does not depend on T , ϵ or d_ϵ .

5. Synthetic Case Study

To evaluate the performance of the proposed forward-looking CARA design, we conduct simulation studies using a data generation mechanism calibrated from the study of Fahey et al. (2020). For the first set of simulation results, we consider biological sex as the only covariate X_i , leading to two subgroups: $X_i = M$ and $X_i = F$ with frequencies 36% and 64%, respectively. We set $n_t = 100$ for $t = 1, \dots, T$ with $T \in \{4, 6, 8\}$. We report simulation evidence separately for the two design objectives in (3) and (4), which require different transformation of the outcome variable.

Setup 1: Power maximization. For the power maximization design objective, we compare four design methods: (i) our proposed design; (ii) our design enhanced by the doubly adaptive biased coin design (DBCD) (Hu and Zhang 2004; Tymofeyev, Rosenberger, and Hu 2007); (iii) complete randomization; (iv) Neyman allocation. We note that DBCD is a response-adaptive randomization design that adjusts treatment allocation toward the optimal allocation dynamically. Since our approach involves assigning treatments across multiple covariate strata, we apply the DBCD design separately within each stratum. We also follow the convention in the literature and use the natural logarithm of the viral load as the outcome variable. Following the trial data, the conditional mean and standard deviation of the potential outcomes are $\mu(F, 1) = 2.50$, $\mu(F, 0) = 2.98$, $\sigma(F, 1) = 0.36$, $\sigma(F, 0) = 2.06$ for the female subgroup, and $\mu(M, 1) = 2.47$, $\mu(M, 0) = 2.72$, $\sigma(M, 1) = 0.82$, $\sigma(M, 0) = 0.31$ for the male subgroup. Given the calibrated model primitives, the average treatment effect is $\tau = -0.33$. To evaluate the power of different designs under the power maximization objective, we also vary the magnitude of the ATE in our data-generating process by perturbing $\mu(M, 0)$, such that the perturbed ATE takes values between 0 and 0.66.

Setup 2: Failure reduction. For the failure reduction objective, we again compare across four designs: (i–ii) the proposed method (with and without DBCD enhancement); (iii) the ethical design (Rosenberger et al. 2001); (iv) complete randomization. In this simulation setup, we transform the original viral load measurements into a binary variable, indicating whether the participant has achieved viral suppression; specifically, if a participant's viral load is below 1000 copies per mL, we set $Y_i = 1$ (viral suppression). Conditional means of the potential outcomes are $\mu(F, 1) = 0.78$, $\mu(F, 0) = 0.57$, $\mu(M, 1) = 0.84$ and $\mu(M, 0) = 0.63$. The average treatment effect is $\tau = 0.21$. To evaluate the power of different designs under the failure reduction objective, we vary the magnitude of ATE in our data-generating process by perturbing $\mu(M, 0)$, such that the perturbed ATE takes values between 0 and 0.42.




Figure 3. Optimized variance (panels (A)–(C)) and failure rate (panels (D)–(F)) for experimental horizons $T = 4, 5, 6$. Covariate stratification is based on biological sex.

As we discussed in Section 1 and illustrated in Figure 1, the outcomes in the original data exhibit both arm- and covariate-dependent delays, which we also incorporate in our simulation setup. To be precise, the delay variable is generated conditionally according a multinomial distribution: $D_i | (X_i = x, A_i = a) \sim \text{Multinomial}(1, \rho(x, a))$. Here, we slightly abuse the notation to define $\rho(x, a) = (\mathbb{P}(D_i = 0 | x, a), \dots, \mathbb{P}(D_i = T - 1 | x, a))$, where each element represents the probability of delaying by exactly d stages (instead of cumulative). For $T = 4$, the delay mechanism parameters are $\rho(F, 1) = (0.64, 0.18, 0.07, 0.03)$ and $\rho(F, 0) = (0.63, 0.18, 0.05, 0.02)$ for the female subgroup, and $\rho(M, 1) = (0.55, 0.23, 0.10, 0.02)$ and $\rho(M, 0) = (0.54, 0.11, 0.21, 0.01)$ for the male subgroup.

We show in panels (A)–(C) of Figure 3 the optimized variance under the power maximization objective. For complete randomization, the variances of the treatment effect estimator are 7.75, 8.04, and 8.17, respectively. The results clearly demonstrate that the proposed fully forward-looking CARA design and its DBCD-enhanced version exhibit a smaller deviation from the oracle variance (indicated by the horizontal black dashed line); to compare, complete randomization leads to severe efficiency loss. Interestingly, the Neyman allocation performs reasonably well in this design, although it still leads to slightly larger variances.

Panels (D)–(F) in the figure collects the failure rates for both our proposed method and the ethical design. For complete randomization, the failure rates are 0.31, 0.33, and 0.35. Again, our proposed method and its DBCD-enhanced version tend to

Table 1. Coverage probability comparison.

Design 1: Power maximization		
Delay perspective	Proposed CARA design	DBCD + Proposed CARA design
Conservative	0.94 (0.01)	0.95 (0.01)
Optimistic	0.95 (0.01)	0.94 (0.01)
Neutral	0.95 (0.01)	0.95 (0.01)
Design 2: Failure reduction		
Delay perspective	Proposed CARA design	DBCD + Proposed CARA design
Conservative	0.94 (0.01)	0.95 (0.01)
Optimistic	0.96 (0.01)	0.94 (0.01)
Neutral	0.95 (0.01)	0.94 (0.01)

perform better, and the overall failure rate there is very close to the oracle level (horizontal black dashed line). In contrast, failure rates tend to be considerably higher for complete randomization.

To assess the finite-sample distributional properties of the treatment effect estimator, we present the coverage probabilities in Table 1 for the two design objectives and the three extrapolations methods for the estimated delay mechanism. Encouragingly, the empirical coverage is extremely close to the nominal level. Lastly, we demonstrate and compare the power properties of different design strategies in Figure 4: all methods control the Type I error rate well when the true ATE is zero; with $\tau > 0$, however, our CARA design has a clear advantage as it is able to detect the nonzero treatment effect more frequently.




Figure 4. Power comparison for the two design objectives. Covariate stratification is based on biological sex.




Figure 5. Optimized variance (panels (A)–(C)) and failure rate (panels (D)–(F)) for experimental horizons $T = 4, 5, 6$. Covariate stratification is based on biological sex and WHO clinical stages.

In summary, this synthetic case study showcases that our fully forward-looking CARA design not only provides valid statistical inference in the presence of delayed responses, but it also delivers higher estimation efficiency and enhanced failure reduction for the two design objectives.

Before closing this section, we provide additional simulation evidence where stratification is based on biological sex and the WHO clinical stages of HIV. Specifically, we define four groups: S_1 for male subjects with WHO stages 1 and 2, S_2 for female subjects with stages 1 and 2, S_3 for male with stages 3 and 4, and S_4 for female with stages 3 and 4.

For the first stratum S_1 , the calibrated parameters are $\sigma(S_1, 1) = 1.63$, $\sigma(S_1, 0) = 1.87$, $\rho(S_1, 1) = (0.10, 0.15,$

$0.32, 0.04)$, and $\rho(S_1, 0) = (0.03, 0.10, 0.22, 0.30)$. For the second stratum S_2 , the parameters are $\sigma(S_2, 1) = 1.85$, $\sigma(S_2, 0) = 1.90$, $\rho(S_2, 1) = (0.14, 0.16, 0.21, 0.05)$ and $\rho(S_2, 0) = (0.10, 0.14, 0.17, 0.06)$. For S_3 , we adopt $\sigma(S_3, 1) = 2.37$, $\sigma(S_3, 0) = 1.63$, $\rho(S_3, 1) = (0.10, 0.19, 0.21, 0.14)$ and $\rho(S_3, 0) = (0.11, 0.16, 0.19, 0.09)$. And finally for S_4 , we use $\sigma(S_4, 1) = 1.73$, $\sigma(S_4, 0) = 2.14$, $\rho(S_4, 1) = (0.05, 0.20, 0.17, 0.18)$ and $\rho(S_4, 0) = (0.19, 0.08, 0.24, 0.02)$. The oracle treatment allocations under the proposed design and the Neyman allocation are demonstrated earlier in Figure 2.

We present simulation results on the optimized variances in panels (A)–(C) of Figure 5, and the failure rates in (D)–(F). Although the covariate stratification differs (see Figure 3),




Figure 6. Power comparison for the two design objectives. Covariate stratification is based on biological sex and WHO clinical stages.

a consistent pattern emerges: our proposed design yields lower variance and enhanced failure reduction by explicitly accounting for the delay mechanism. To compare, the variances of the treatment effect estimator employing complete randomization are 43.96, 45.17, and 47.82, and the failure rates are 0.34, 0.32, and 0.34, respectively. To complement our earlier findings in Figure 4, we also report power comparisons in Figure 6.

6. Conclusion

In this article, we introduce fully forward-looking covariate-adjusted response-adaptive (CARA) designs that effectively address the challenge of delayed outcomes in adaptive experiments, allowing the delay mechanism to be both arm- and covariate-dependent. Our approach sequentially estimates the delay mechanism, which is then used to inform optimal treatment allocation. Through a comprehensive synthetic case study, we demonstrate that our proposed design can lead to substantial gains in both experimental efficiency and participant welfare compared to traditional CARA designs. We provide rigorous theoretical foundation by deriving semiparametric efficiency bounds in the presence of delayed responses and establishing the consistency and asymptotic normality of treatment effect estimators.

Supplementary Materials

The supplementary material (Ma, Wang, and Wei 2025) contains general theoretical results, their proofs, and reports additional simulation evidence.

Acknowledgments

We thank the editor, the associate editor, two anonymous reviewers, seminar participants, and conference audiences for their valuable feedback, which significantly improved the article.

Disclosure Statement

The authors report there are no competing interests to declare.

Funding

Xinwei Ma gratefully acknowledges financial support from the National Institutes of Health through the grant R01-AG089512. Jingshen Wang gratefully acknowledges financial support from the National Institutes of

Health through the grant R01-AG089512, and the National Science Foundation through the grants DMS-2239047 and DMS-2220537.

ORCID

Xinwei Ma <http://orcid.org/0000-0001-8827-9146>
 Jingshen Wang <http://orcid.org/0000-0002-1432-3834>
 Waverly Wei <http://orcid.org/0000-0003-3504-4436>

References

- Antognini, A. B., and Zagoraiou, M. (2015), "On the Almost Sure Convergence of Adaptive Allocation Procedures," *Bernoulli*, 21, 881–908. [8]
- Bai, Y., Romano, J. P., and Shaikh, A. M. (2022), "Inference in Experiments with Matched Pairs," *Journal of the American Statistical Association*, 117, 1726–1737. [1]
- Bai, Z.-D., Hu, F., and Rosenberger, W. F. (2002), "Asymptotic Properties of Adaptive Designs for Clinical Trials with Delayed Response," *Annals of Statistics*, 30, 122–139. [1,2]
- Bibaut, A., and Kallus, N. (2025), "Demystifying Inference after Adaptive Experiments," *Annual Review of Statistics and its Application*, 12, 407–423. [1]
- Blackwell, M., Pashley, N. E., and Valentino, D. (2023), "Batch Adaptive Designs to Improve Efficiency in Social Science Experiments," working paper. [1]
- Bugni, F. A., Canay, I. A., and Shaikh, A. M. (2018), "Inference under Covariate-Adaptive Randomization," *Journal of the American Statistical Association*, 113, 1784–1796. [1]
- Cattaneo, M. D. (2010), "Efficient Semiparametric Estimation of Multi-Valued Treatment Effects under Ignorability," *Journal of Econometrics*, 155, 138–154. [2]
- Dorn, J. (2025), "How Much Weak Overlap Can Doubly Robust t-statistics Handle?" arXiv preprint arXiv:2504.13273. [4]
- Fahey, C. A., Njau, P. F., Katabaro, E., Mfaume, R. S., Ulena, N., Mwenda, N., et al. (2020), "Financial Incentives to Promote Retention in Care and Viral Suppression in Adults with HIV Initiating Antiretroviral Therapy in Tanzania: A Three-Arm Randomised Controlled Trial," *The Lancet HIV*, 7, e762–e771. [1,9]
- Hahn, J. (1998), "On the Role of the Propensity Score in Efficient Semiparametric Estimation of Average Treatment Effects," *Econometrica*, 66, 315–331. [2]
- Hahn, J., Hirano, K., and Karlan, D. (2011), "Adaptive Experimental Design Using the Propensity Score," *Journal of Business & Economic Statistics*, 29, 96–108. [1]
- Hall, P., and Heyde, C. C. (1980), *Martingale Limit Theory and its Application*, New York, NY: Academic Press. [3]
- Hampson, L. V., and Jennison, C. (2013), "Group Sequential Tests for Delayed Responses," *Journal of the Royal Statistical Society, Series B*, 75, 3–54. [2]

- Heiler, P., and Kazak, E. (2021), "Valid Inference for Treatment Effect Parameters under Irregular Identification and Many Extreme Propensity Scores," *Journal of Econometrics*, 222, 1083–1108. [4]
- Hirano, K., Imbens, G. W., and Ridder, G., (2003), "Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score," *Econometrica*, 71, 1161–1189. [2]
- Hu, F., and Rosenberger, W. F. (2006), *The Theory of Response-adaptive Randomization in Clinical Trials*, New York, NY: Wiley. [1]
- Hu, F., and Zhang, L.-X. (2004), "Asymptotic Properties of Doubly Adaptive Biased Coin Designs for Multitreatment Clinical Trials," *Annals of Statistics*, 32, 268–301. [8,9]
- Hu, F., Zhang, L.-X., Cheung, S. H., and Chan, W. S. (2008), "Doubly Adaptive Biased Coin Designs with Delayed Responses," *Canadian Journal of Statistics*, 36, 541–559. [2]
- Hu, J., Zhu, H., and Hu, F. (2015), "A Unified Family of Covariate-Adjusted Response-Adaptive Designs based on Efficiency and Ethics," *Journal of the American Statistical Association*, 110, 357–367. [1,8]
- Hu, Y., and Hu, F. (2012), "Asymptotic Properties of Covariate-Adaptive Randomization," *Annals of Statistics*, 40, 1794–1815. [1]
- Lachin, J. M. (1977), "Sample Size Determinations for $r \times c$ Comparative Trials," *Biometrics*, 33, 315–324. [4]
- Lin, R., Thall, P. F., and Yuan, Y. (2020), "An Adaptive Trial Design to Optimize Dose-Schedule Regimes with Delayed Outcomes," *Biometrics*, 76, 304–315. [2]
- Ma, X., Sasaki, Y., and Wang, Y. (2025), "Testing Limited Overlap," *Econometric Theory*, 41, 1129–1162. [4]
- Ma, X., and Wang, J. (2020), "Robust Inference Using Inverse Probability Weighting," *Journal of the American Statistical Association*, 115, 1851–1860. [4]
- Ma, X., Wang, J., and Wei, W. (2025), "Supplementary Material to 'Covariate-Adjusted Response-Adaptive Design with Delayed Outcomes.'" [12]
- Offer-Westort, M., Coppock, A., and Green, D. P. (2021), "Adaptive Experimental Design: Prospects and Applications in Political Science," *American Journal of Political Science*, 65, 826–844. [1]
- Robertson, D. S., Lee, K. M., López-Kolkovska, B. C., and Villar, S. S. (2023), "Response-Adaptive Randomization in Clinical Trials: From Myths to Practical Considerations," *Statistical Science*, 38, 185–208. [1,2]
- Rosenberger, W. F., Stallard, N., Ivanova, A., Harper, C. N., and Ricks, M. L. (2001), "Optimal Adaptive Designs for Binary Response Trials," *Biometrics*, 57, 909–913. [2,4,5,9]
- Sasaki, Y., and Ura, T. (2022), "Estimation and Inference for Moments of Ratios with Robustness Against Large Trimming Bias," *Econometric Theory*, 38, 66–112. [4]
- Tsiatis, A. A. (2006), *Semiparametric Theory and Missing Data*, New York, NY: Springer. [2]
- Tymofeyev, Y., Rosenberger, W. F., and Hu, F. (2007), "Implementing Optimal Allocation in Sequential Binary Response Experiments," *Journal of the American Statistical Association*, 102, 224–234. [2,4,5,9]
- van der Vaart, A. W. (2000), *Asymptotic Statistics*, Cambridge, UK: Cambridge University Press. [2]
- Wei, W., Ma, X., and Wang, J. (2023), "Fair Adaptive Experiments," in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NeurIPS 2023, pp. 19157–19169. [1]
- (2025), "Adaptive Experiments Toward Learning Treatment Effect Heterogeneity," *Journal of the Royal Statistical Society, Series B*, 87, 1055–1084. [1]
- Wu, Y., Zheng, Z., Zhang, G., Zhang, Z., and Wang, C. (2025), "Nonstationary A/B Tests: Optimal Variance Reduction, Bias Correction, and Valid Inference," *Management Science*, 71, 4707–4727. [2]
- Zhang, L.-X., Hu, F., Cheung, S. H., and Chan, W. S. (2007), "Asymptotic Properties of Covariate-Adjusted Response-Adaptive Designs," *Annals of Statistics*, 35, 1166–1182. [1]
- Zhao, J. (2023), "Adaptive Neyman Allocation," arXiv preprint arXiv:2309.08808. [1,2]
- Zhu, H., and Zhu, H. (2023), "Covariate-Adjusted Response-Adaptive Designs based on Semiparametric Approaches," *Biometrics*, 79, 2895–2906. [2]