



Drawing valid targeted inference when covariate-adjusted response-adaptive RCT meets data-adaptive loss-based estimation, with an application to the LASSO

Wenjing Zheng, Antoine Chambaz, Mark van Der Laan

► To cite this version:

Wenjing Zheng, Antoine Chambaz, Mark van Der Laan. Drawing valid targeted inference when covariate-adjusted response-adaptive RCT meets data-adaptive loss-based estimation, with an application to the LASSO. 2015. hal-01180719

HAL Id: hal-01180719

<https://hal.science/hal-01180719>

Preprint submitted on 28 Jul 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Drawing valid targeted inference when covariate-adjusted response-adaptive RCT meets data-adaptive loss-based estimation, with an application to the LASSO

Wenjing Zheng^{1,2}, Antoine Chambaz³, and Mark van der Laan¹

¹Division of Biostatistics, University of California, Berkeley

²Center for AIDS Prevention Studies,
University of California, San Francisco

³Modal'X, Université Paris Ouest Nanterre

July 8, 2015

Abstract

Adaptive clinical trial design methods have garnered growing attention in the recent years, in large part due to their greater flexibility over their traditional counterparts. One such design is the so-called covariate-adjusted, response-adaptive (CARA) randomized controlled trial (RCT). In a CARA RCT, the treatment randomization schemes are allowed to depend on the patient's pre-treatment covariates, and the investigators have the opportunity to adjust these schemes during the course of the trial based on accruing information (including previous responses), in order to meet a pre-specified optimality criterion, while preserving the validity of the trial in learning its primary study parameter.

In this article, we propose a new group-sequential CARA RCT design and corresponding analytical procedure that admits the use of flexible data-adaptive techniques. The proposed design framework can target general adaption optimality criteria that may not have a closed-form solution, thanks to a loss-based approach in defining and estimating the unknown optimal randomization scheme. Both in predicting the conditional response and in constructing the treatment randomization schemes, this framework uses loss-based data-adaptive estimation over general classes of functions (which may change with sample size). Since the randomization adaptation is response-adaptive, this innovative flexibility potentially translates into more effective adaptation towards the optimality criterion. To target the primary study parameter, the proposed analytical method provides robust inference of the parameter, despite arbitrarily mis-specified response models, under the most general settings.

Specifically, we establish that, under appropriate entropy conditions on the classes of functions, the resulting sequence of randomization schemes converges to a fixed scheme, and the proposed treatment effect estimator is consistent (even under a mis-specified response model), asymptotically Gaussian, and gives rise to valid confidence intervals of given asymptotic levels. Moreover, the limiting randomization scheme coincides with the unknown optimal randomization scheme when, simultaneously, the response model is correctly specified and the optimal scheme belongs to the limit of the user-supplied classes of randomization schemes. We illustrate the applicability of these general theoretical results with a LASSO-based CARA RCT. In this example, both the response model and the optimal treatment randomization are estimated using a sequence of LASSO logistic models that may increase with sample size. It follows immediately from our general theorems that this LASSO-based CARA RCT converges to a fixed design and yields consistent and asymptotically Gaussian effect estimates, under minimal conditions on the smoothness of the basis functions in the LASSO logistic models. We exemplify the proposed methods with a simulation study.

1 Introduction

1.1 Covariate-adjusted, response-adaptive randomized clinical trials

Adaptive clinical trial design methods have garnered growing attention in recent years. In a fixed trial design, all key aspects of the trials are set before the start of the data collection, usually based on assumptions that are yet unsure at the design stage. By contrast, an adaptive trial design allows pre-specified modifications of the ongoing trial based on accruing data, while preserving the validity and integrity of the trial. This flexibility potentially translates into more tailored studies. The study could be more efficient, *e.g.*, have shorter duration, or involve fewer subjects. The study could have greater chance to answer the clinical questions of interest, *e.g.*, detect a treatment effect if one exists, or gather broader dose-response information.

Once they have defined the *primary study objective* of the trial (*e.g.*, testing the effect of a treatment), the investigators may wish to accommodate additional *design objectives* (*e.g.*, minimizing sample size or exposure of patients to inferior treatment) without compromising the trial. To do this, they may use an adaptive randomization trial design. We focus here on the study of the so-called covariate-adjusted response-adaptive (CARA) randomized controlled trials (RCTs). In a CARA RCT, the randomization schemes are allowed to depend on the patient's pre-treatment covariates, and the investigators can adjust the randomization schemes during the course of the trial based on accruing information, including previous responses, in order to meet the pre-specified design objectives. Such adjustments take place at interim time points given by sequential inclusion of blocks of c patients, where $c \geq 1$ is a pre-specified integer. We consider the case of $c = 1$ for simplicity of exposition, though the discussions generalize to any $c > 1$.

The trial protocol pre-specifies the observed data structure, scientific parameters of interest (the primary study objective), analysis methods, and a criterion characterizing an optimal randomization scheme (the design objective). Here, baseline covariates and a primary response of interest are measured on each patient. The primary study objective is the marginal treatment effect. The design objective is captured by an optimality criterion that is a function of the unknown conditional response.

Contrary to a fixed design RCT, a CARA RCT produces non-independent and non-identically distributed observations, therein lie the subtleties in its theoretical study. Traditionally, covariate-adjusted analysis of a fixed design RCT is carried out using a parametric model for the conditional response (or distribution) given treatment and covariates. Under correct specification, the maximum likelihood estimator for this model is consistent and asymptotically Gaussian. The extension of this *non-robust inference* to CARA RCTs has been established and discussed in (Zhang, Hu, Cheung, and Chan, 2007) and (Hu and Rosenberger, 2006). A recent development in the analysis of fixed design RCTs is the use of doubly robust methods like the targeted minimum loss estimation (TMLE, van der Laan and Rubin 2006) to obtain consistent, asymptotically Gaussian estimators under arbitrarily mis-specified models (Moore and van der Laan, 2009, Rosenblum, 2011). A first extension of this *robust inference* to CARA RCTs has been proposed by Chambaz and van der Laan (2013). They showed that, when the treatment assignment is conditioned only on a discrete summary measure of the covariates, it is possible to derive a consistent and asymptotically Gaussian estimator of the study parameter which is robust to mis-specification of an arbitrary parametric response model.

Despite the above developments, several gaps remain to be addressed to fully realize the promise of CARA RCTs. We focus on two of them. Firstly, because the robust inference provided by Chambaz and van der Laan (2013) relies on assigning treatment based on discrete covariate summaries, its application is perhaps limited in real-life RCTs where response to treatment may be correlated with a large number of a patient's baseline characteristics, some of which being continuous. Secondly, even though under robust inference, the choice of the response model does not compromise consistent estimation of the study parameter, it may still affect the estimation of the optimal randomization scheme. Specifically, since the randomization scheme is response-adaptive, a more data-adaptive estimator of the conditional response model can more effectively steer the randomization schemes towards the unknown optimal randomization scheme. Moreover, since a patient's primary response is often correlated with many individual characteristics, greater latitude in adjusting for these baseline covariates, in both treatment assignment and conditional response estimation, allows the investigators to better account for heterogeneity in the patients population. Traditional

parametric regression techniques are often too restrictive in such a high-dimensional scenario. While the use of data-adaptive techniques is very common in the independent and identically distributed (i.i.d.) context, its applicability in an adaptive RCT remains rather uncharted.

In this article, we aim to bridge the two aforementioned gaps in the study of CARA RCTs. Firstly, we achieve robust inference of the study parameter without restrictions on the covariate measures used in the treatment randomization. Secondly, we adopt the use of loss-based data-adaptive estimation over general classes of functions (which may change with sample size) in constructing the treatment randomization schemes and in predicting the unknown conditional response. This allows one to target general randomization optimality criteria that may not have a closed form solution, and it may potentially improve the estimation of the unknown optimal randomization schemes. We establish that, under appropriate entropy conditions on the classes of functions, the resulting sequence of randomization schemes converges to a fixed randomization scheme, and the proposed estimator is consistent (even under a mis-specified response model), asymptotically Gaussian, and gives rise to valid confidence intervals of given asymptotic levels. Moreover, the limiting randomization scheme coincides with the unknown optimal randomization scheme when, simultaneously, the response model is correctly specified and the optimal randomization scheme belongs to the limit of the user-supplied classes of randomization schemes. Our theoretical results benefit from recent advances in maximal inequalities for martingales by van Handel (2011).

For concreteness, our parameter of interest here is the marginal risk difference, ψ_0 , and our design objective is to maximize the efficiency of the study (*i.e.*, to reach a valid result using as few blocks of patients as possible). As we shall see, the optimal randomization scheme is, in this case, the so-called covariate-adjusted Neyman design, which minimizes the Cramér-Rao lower bound on the asymptotic variances of a large class of estimators of ψ_0 . We emphasize that the results presented here are not limited to the marginal risk difference or the Neyman design, and can be easily modified to other study objectives/effect measures and other design objectives/optimality criteria.

To illustrate the proposed framework, we consider the LASSO to estimate the conditional response given treatment and baseline covariates and to target the unknown optimal randomization scheme. This example essentially encompasses the parametric approach in (Chambaz and van der Laan, 2013) as a special case. The asymptotic results ensue under minimal conditions on the smoothness of the LASSO basis functions. The performance of the procedure is evaluated in a simulation study.

In the next section, we give a bird’s eye view of the literature on CARA RCTs and put our contribution in context.

1.2 Literature review

Adaptive randomization has a long history that can be traced back to the 1930s. We refer to (Rosenberger, 1996, Rosenberger, Sverdlov, and Hu, 2012), (Hu and Rosenberger, 2006, Section 1.2) and (Jennison and Turnbull, 2000, Section 17.4) for a comprehensive historical perspective. Many articles are devoted to the study of response-adaptive randomizations, which select current treatment probabilities based on responses of previous patients (but not on the covariates of the current patients). We summarize some representative works below, but refer to (Hu and Rosenberger, 2006, Chambaz and van der Laan, 2011b, Zhang and Rosenberger, 2012, Rosenberger et al., 2012) for a bibliography on that topic. The first methods are based on urn-models (e.g. Wei and Durham (1978), Ivanova (2003)). There, treatment allocation is represented by drawing balls of different colors from an urn, and the urn composition is updated based on accruing responses, with the ethical goal of assigning most patients to the superior treatment arm. Since there is no formal criterion governing how skewed the treatment allocation should be, significant loss of power can arise when the effect size between treatment arms is large (Rosenberger and Hu, 2004). A formal “optimal allocation approach” was proposed by Hu and Rosenberger (2003). There, an optimal allocation is defined as a solution to a possibly constrained optimization problem, such as minimizing sample size (yielding the so-called Neyman allocation), or minimizing failure while preserving power. This optimal allocation is a function of unknown parameters of the conditional response, which are estimated using a parametric model based on available responses. Consistency and asymptotic normality of the maximum likelihood estimator

for this model were established in Hu and Rosenberger (2006).

In a heterogeneous population where response is often correlated with the patient’s individual characteristics, it is sensible to take into account covariates in treatment assignment. CARA randomization extends response-adaptive randomization to tackle heterogeneity by dynamically calculating the allocation probabilities based on previous responses and current and past values of certain covariates. Compared to the broader literature on response-adaptive randomization, the advances in CARA randomization are relatively recent, but growing steadily. Among the first approaches, (Rosenberger, Vidyashankar, and Agarwal, 2001, Bandyopadhyay and Biswas, 2001) considered allocations that are proportional to the covariate-adjusted treatment difference, which is estimated using generalized linear models for the conditional response. Though these procedures are not defined based on formal optimality criteria, their general goal is to allocate more patients to their corresponding superior treatment arm. Atkinson and Biswas (2005) presented a biased-coin design with skewed allocation, which is determined by sequentially maximizing a function that combines the variance of the parameter estimator, based on a Gaussian linear model for the conditional response, and the conditional treatment effect given covariates. Up till here, very little work had been devoted to asymptotic properties of CARA designs. Subsequently, Zhang et al. (2007) established the asymptotic theory for CARA designs converging to a given target covariate-adjusted allocation function when the conditional responses follow a parametric model. Zhang and Hu (2009) proposed a covariate-adjusted doubly-adaptive biased coin design whose asymptotic variance achieves the efficiency bound. In these optimal allocation approaches, the challenge remains that the explicit form of the target covariate-adjusted allocation function is not known. To overcome this, it has often been derived as a covariate-adjusted version of the optimal allocation from a framework with no covariates (Rosenberger et al., 2012). Chang and Park (2013) proposed a sequential estimation of CARA designs under generalized linear models for the conditional response. This procedure allocates treatment based on the patients’ baseline covariates, accruing information and sequential estimates of the treatment effect. It uses a stopping rule that depends on the observed Fisher information. With regard to hypotheses testing, Shao, Yu, and Zhong (2010), Shao and Yu (2013) provided asymptotic results for valid tests under generalized linear models for the conditional responses in the context of covariate-adaptive randomization. Most recently, progress has also been made in CARA designs in the longitudinal settings, see for example (Biswas, Bhattacharya, and Park, 2014, Huang, Liu, and Hu, 2013, Sverdlov, Rosenberger, and Ryznik, 2013).

The above contributions have established the validity of statistical inference for CARA RCTs under a correctly specified model, thus extending many of the classical non-robust inference methods from the fixed design setting into the CARA setting. Doubly robust approaches like TMLE allow to go beyond correctly specified models by leveraging the known treatment randomization to provide the necessary bias reduction over the mis-specified response model. Moore and van der Laan (2009), Rosenblum (2011) address the fixed design setting and Chambaz and van der Laan (2013) provide the first extension to the adaptive design setting.

Finiteness conditions were at the core of (Zhang et al., 2007) (correctly specified parametric response model) and (Chambaz and van der Laan, 2013) (arbitrary parametric response model and treatment assignment based on discrete covariates). They were instrumental in the asymptotic study based on Taylor approximations. Recent advances by van Handel (2011) on maximal deviation bounds for martingales allow us to apply more general empirical processes techniques, thus opening the door for the use of data-adaptive estimators to target the optimal randomization scheme while preserving valid inference. More specifically, we extend the robust inference framework of Chambaz and van der Laan (2013) to allow for the use of general classes of conditional response estimators and randomization schemes. Moreover, we adopt a loss-based approach to defining and targeting the optimal randomization scheme, thereby also extending applicability of CARA RCT to optimal randomization criteria that may not have a closed form solution.

1.3 Organization of the article

The remainder of this article is organized as follows. Section 2 presents the statistical challenges, and describes the proposed targeted, adaptive sampling scheme and inference method to address them. The

section also states our main assumptions and principal result. Section 3 provides contextual comments for content of Section 2. Section 4 presents the building blocks of the main result, therefore shedding light on the inner mechanism of the procedure. Section 5 illustrates the procedure using the LASSO methodology both to target the optimal randomization scheme and to estimate the conditional response given baseline covariates and treatment. The performance of the LASSO-based CARA RCT is assessed in a simulation study in Section 6. The article closes on a discussion in Section 7. All proofs and some useful, technical results are gathered in the appendix.

2 Targeted inference based on data adaptively drawn from a CARA RCT using loss-based estimation

At sample size n , we will have observed the ordered vector $\mathbf{O}_n \equiv (O_1, \dots, O_n)$, with convention $O_0 \equiv \emptyset$. For every $1 \leq i \leq n$, the data structure O_i writes as $O_i \equiv (W_i, A_i, Y_i)$. Here, $W_i \in \mathcal{W}$ consists of the baseline covariates (some of which may be continuous) of the i th patient, $A_i \in \mathcal{A} \equiv \{0, 1\}$ is the binary treatment of interest assigned to her, and $Y_i \in \mathcal{Y}$ is her primary response of interest. We assume that the outcome space $\mathcal{O} \equiv \mathcal{W} \times \mathcal{A} \times \mathcal{Y}$ is bounded. Without loss of generality, we may then assume that $\mathcal{Y} \equiv (0, 1)$, *i.e.*, that the responses are between and bounded away from 0 and 1.

Section 2.1 presents the target statistical parameter and optimal randomization scheme. It also lays out the foundations to describe the proposed CARA RCT. The description is completed in Sections 2.3 and 2.4, where we present our adaptive sampling scheme and targeted minimum loss estimator. Section 2.6 states our main assumptions and result.

2.1 Likelihood, model, statistical parameter, optimal randomization scheme

Let μ_W be a measure on \mathcal{W} equipped with a σ -field, $\mu_A = \text{Dirac}(0) + \text{Dirac}(1)$ be a measure on \mathcal{A} equipped with its σ -field, and μ_Y be the Lebesgue measure on \mathcal{Y} equipped with the Borel σ -field. Define $\mu \equiv \mu_W \otimes \mu_A \otimes \mu_Y$, a measure on \mathcal{O} equipped with the product of the above σ -fields. In an RCT, the unknown, true likelihood of \mathbf{O}_n with respect to (wrt) $\mu^{\otimes n}$ is given by the following factorization of the density of \mathbf{O}_n wrt $\mu^{\otimes n}$:

$$\begin{aligned} \mathcal{L}_{\mathbf{f}_0, \mathbf{g}_n}(\mathbf{O}_n) &\equiv \prod_{i=1}^n \mathbf{f}_{W,0}(W_i) \times (A_i g_i(1|W_i) + (1 - A_i) g_i(0|W_i)) \times \mathbf{f}_{Y,0}(Y_i|A_i, W_i) \\ &= \prod_{i=1}^n \mathbf{f}_{W,0}(W_i) \times g_i(A_i|W_i) \times \mathbf{f}_{Y,0}(Y_i|A_i, W_i), \end{aligned} \quad (1)$$

where (i) $w \mapsto \mathbf{f}_{W,0}(w)$ is the density wrt μ_W of a true, unknown law $Q_{W,0}$ on \mathcal{W} (that we assume being dominated by μ_W), (ii) $\{y \mapsto \mathbf{f}_{Y,0}(y|a, w) : (a, w) \in \mathcal{A} \times \mathcal{W}\}$ is the collection of the conditional densities $y \mapsto \mathbf{f}_{Y,0}(y|a, w)$ wrt μ_Y of true, unknown laws on \mathcal{Y} indexed by (a, w) (that we assume being all dominated by μ_Y), (iii) $g_i(1|W_i)$ is the known (given by user) conditional probability that $A_i = 1$ given W_i , and (iv) $\mathbf{g}_n \equiv (g_1, \dots, g_n)$, the ordered vector of the n first randomization schemes. One reads in (1) (i) that W_1, \dots, W_n are independently sampled from $Q_{W,0}$, (ii) that Y_1, \dots, Y_n are conditionally sampled from $\mathbf{f}_{Y,0}(\cdot|A_1, W_1)\mu_Y, \dots, \mathbf{f}_{Y,0}(\cdot|A_n, W_n)\mu_Y$, respectively, and (iii) that each A_i is drawn conditionally on W_i from the Bernoulli distribution with known parameter $g_i(1|W_i)$.

Let \mathcal{F} be the semiparametric collection of all elements of the form $\mathbf{f} = (\mathbf{f}_W, \mathbf{f}_Y(\cdot|a, w), (a, w) \in \mathcal{A} \times \mathcal{W})$ where \mathbf{f}_W is a density wrt μ_W and each $\mathbf{f}_Y(\cdot|a, w)$ is a density wrt μ_Y . In particular, we define $\mathbf{f}_0 \equiv (\mathbf{f}_{W,0}, \mathbf{f}_{Y,0}(\cdot|a, w), (a, w) \in \mathcal{A} \times \mathcal{W}) \in \mathcal{F}$. In light of (1) define, for every $\mathbf{f} \in \mathcal{F}$, $\mathcal{L}_{\mathbf{f}, \mathbf{g}_n}(\mathbf{O}_n) \equiv \prod_{i=1}^n \mathbf{f}_W(W_i) \times g_i(A_i|W_i) \times \mathbf{f}_Y(Y_i|A_i, W_i)$. The set $\{\mathcal{L}_{\mathbf{f}, \mathbf{g}_n} : \mathbf{f} \in \mathcal{F}\}$ is a semiparametric model for the likelihood of \mathbf{O}_n . It contains the true, unknown likelihood $\mathcal{L}_{\mathbf{f}_0, \mathbf{g}_n}$.

For the sake of illustration, we choose the marginal treatment effect on an additive scale as our parameter of interest. Thus, let $\Upsilon : \mathcal{F} \rightarrow [-1, 1]$ be the mapping such that, for every $\mathbf{f} = (\mathbf{f}_W, \mathbf{f}_Y(\cdot|a, w), (a, w) \in \mathcal{A} \times \mathcal{W})$,

$$\Upsilon(\mathbf{f}) = \int (Q_{Y,\mathbf{f}}(1, w) - Q_{Y,\mathbf{f}}(0, w)) \mathbf{f}_W(w) d\mu_W, \quad (2)$$

where $Q_{Y,\mathbf{f}}(a, w) = \int y \mathbf{f}_Y(y|a, w) d\mu_Y$ is the mean of $\mathbf{f}_Y(\cdot|a, w) \mu_Y$. The true marginal treatment effect on an additive scale is $\psi_0 \equiv \Upsilon(\mathbf{f}_0)$. Of particular interest in medical, epidemiological and social sciences research, it can be interpreted causally under assumptions on the data-generating process (Pearl, 2000).

We have not specified yet what is precisely the sequence of randomization schemes $\mathbf{g}_n \equiv (g_1, \dots, g_n)$. Our CARA sampling scheme “targets” a randomization scheme g_0 which minimizes a user-supplied optimality criterion. By targeting g_0 we mean estimating g_0 based on past observations, and relying on the resulting estimator to collect the next block of data, as seen in (1). For the sake of illustration, we consider the case that g_0 is the following minimizer

$$g_0 \equiv \arg \min_g E_{P_{Q_0,g}} \left(\frac{(Y - Q_{Y,\mathbf{f}_0}(A, W))^2}{g^2(A|W)} \right) \quad (3)$$

across all randomization schemes g . We emphasize that the above definition of g_0 involves the unknown \mathbf{f}_0 , so it is unknown too. We will comment on (3) and motivate our interest in g_0 in Section 3. As we shall see, g_0 minimizes a generalized Cramér-Rao lower bound for ψ_0 . Known in the literature as the *Neyman design* (Hu and Rosenberger, 2006), g_0 actually has a closed-form expression as a function of \mathbf{f}_0 . We *do not* use this closed-form expression in order to illustrate the generality of our framework which allows to target any randomization scheme defined as a minimizer of an optimality criterion.

2.2 Notation

Let $O \equiv (W, A, Y)$ be a generic data-structure. Every distribution of O consists of two components: on the one hand, the marginal distribution of W and the conditional distribution of Y given (A, W) , which correspond to a $\mathbf{f} \in \mathcal{F}$; on the other hand, the conditional distribution of A given W , or randomization scheme. To reflect this dichotomy, we denote the distribution of O as $P_{Q,g}$, where Q equals the couple formed by the marginal distribution of W and the conditional distribution of Y given (A, W) , and g equals the randomization scheme. We denote Q_0 the true couple Q in our population of interest, which corresponds to \mathbf{f}_0 and is unknown to us. For a given Q , we denote Q_W the related marginal distribution of W and Q_Y the related conditional expectation of Y given (A, W) . If $Q = Q_0$, then Q_W and Q_Y are denoted $Q_{W,0}$ and $Q_{Y,0}$, respectively.

We denote \mathcal{G} and \mathcal{Q}_Y the set of all randomization schemes and the set of all conditional expectations of Y given (A, W) , respectively. Thus, for any Q and g , $P_{Q_0,g}$ is the true, partially unknown distribution of O when one relies on g , and $E_{P_{Q,g}}(Y|A, W) = Q_Y(A, W)$, $P_{Q,g}(A = 1|W) = g(1|W) = 1 - g(0|W)$ $P_{Q,g}$ -almost surely.

With this notation, ψ_0 can be rewritten

$$\psi_0 = \int (Q_{Y,0}(1, w) - Q_{Y,0}(0, w)) dQ_{W,0}(w)$$

and satisfies $\psi_0 = E_{P_{Q_0,g}}(Q_{Y,0}(1, W) - Q_{Y,0}(0, W))$ whatever is $g \in \mathcal{G}$.

2.3 Loss functions and working models

Let $g^b \in \mathcal{G}$ be the balanced randomization scheme wherein each arm is assigned with probability 1/2 regardless of baseline covariates. Let $g^r \in \mathcal{G}$ be a randomization scheme, bounded away from 0 and 1 by choice, that serves as a reference. In addition, let L be a loss function for $Q_{Y,0}$ and $\mathcal{Q}_{1,n}$ be a working model for the conditional response

$$\mathcal{Q}_{1,n} \equiv \{Q_{Y,\beta} : \beta \in B_n\} \subset \mathcal{Q}_Y.$$

One choice of L is the quasi negative-log-likelihood loss function L^{kl} . For any $Q_Y \in \mathcal{Q}_Y$ bounded away from 0 and 1, $L^{\text{kl}}(Q_Y)$ satisfies

$$-L^{\text{kl}}(Q_Y)(O) \equiv Y \log(Q_Y(A, W)) + (1 - Y) \log(1 - Q_Y(A, W)). \quad (4)$$

Another interesting loss function L for $Q_{Y,0}$ is the least-square loss function L^{ls} , given by

$$L^{\text{ls}}(Q_Y)(O) \equiv (Y - Q_Y(A, W))^2. \quad (5)$$

Likewise, let L_{Q_Y} be a loss function for g_0 , which may depend on $Q_Y \in \mathcal{Q}_Y$, and let $\mathcal{G}_{1,n} \subset \mathcal{G}$ be a working model for the optimal randomization scheme. In this context, a loss function for g_0 may be given, for any $Q_Y \in \mathcal{Q}_Y$, by

$$L_{Q_Y}(g)(O) \equiv \frac{(Y - Q_Y(A, W))^2}{g(A|W)}. \quad (6)$$

We explain the motivation and justification for this loss function in section 3.3.

As suggested by the notation, the sets $\mathcal{Q}_{1,n}$ and $\mathcal{G}_{1,n}$ may depend on n . In that case, the two sequences of sets must be non-decreasing. Moreover, the specifications must guarantee that the elements of $\mathcal{Q}_1 \equiv \cup_{n \geq 1} \mathcal{Q}_{1,n}$ and those of $\mathcal{G}_1 \equiv \cup_{n \geq 1} \mathcal{G}_{1,n}$ be uniformly bounded away from 0 and 1.

2.4 Targeted adaptive sampling and inference

The estimation of g_0 involves the estimation of $Q_{Y,0}$. At each step, the current estimators of $Q_{Y,0}$ and g_0 are also used to craft a targeted estimator of ψ_0 .

We initialize the sampling scheme by setting $g_1 \equiv g^b$. Consider $1 < i < n$. Since

$$g_0 = \arg \min_{g \in \mathcal{G}} E_{P_{Q_0,g}} \left(\frac{L_{Q_{Y,0}}(g)(O)}{g(A|W)} \right) \quad \text{and} \quad Q_{Y,0} = \arg \min_{Q_Y \in \mathcal{Q}} E_{P_{Q_0,g}} (L(Q_Y)(O)),$$

we define

$$g_i \in \arg \min_{g \in \mathcal{G}_{1,i}} \frac{1}{i-1} \sum_{j=1}^{i-1} \frac{L_{Q_{Y,\beta_i}}(g)(O_j)}{g_j(A_j|W_j)}, \quad (7)$$

where

$$\beta_i \in \arg \min_{\beta \in B_i} \frac{1}{i-1} \sum_{j=1}^{i-1} L(Q_{Y,\beta})(O_j) \frac{g^r(A_j|W_j)}{g_j(A_j|W_j)}. \quad (8)$$

By specifying the sequence of randomization schemes, this completes the definition of the likelihood function, hence the characterization of our sampling scheme.

To estimate ψ_0 based on \mathbf{O}_n , we introduce the following one-dimensional parametric model for $Q_{Y,0}$:

$$\{Q_{Y,\beta_n}(\epsilon) \equiv \text{expit}(\text{logit}(Q_{Y,\beta_n}) + \epsilon H(g_n)) : \epsilon \in \mathcal{E}\}, \quad (9)$$

where $\mathcal{E} \subset \mathbb{R}$ is a closed, bounded interval containing 0 in its interior and $H(g)(O) \equiv (2A - 1)/g(A|W)$. The optimal fluctuation parameter is

$$\epsilon_n \in \arg \min_{\epsilon \in \mathcal{E}} \frac{1}{n} \sum_{i=1}^n L^{\text{kl}}(Q_{Y,\beta_n}(\epsilon))(O_i) \frac{g_n(A_i|W_i)}{g_i(A_i|W_i)}. \quad (10)$$

We set $Q_{Y,\beta_n}^* \equiv Q_{Y,\beta_n}(\epsilon_n)$ and define

$$\psi_n^* \equiv \frac{1}{n} \sum_{i=1}^n Q_{Y,\beta_n}^*(1, W_i) - Q_{Y,\beta_n}^*(0, W_i). \quad (11)$$

We show in Section 4 that ψ_n^* consistently estimates ψ_0 . We also show that $\sqrt{n/\Sigma_n}(\psi_n^* - \psi_0)$ is approximately standard normally distributed, where Σ_n is an explicit estimator (20). This enables the construction of confidence intervals of desired asymptotic level. As for the optimal randomization scheme g_0 , we show that it is targeted indeed, in the sense that g_n converges to the projection of g_0 onto $\cup_{n \geq 1} \mathcal{G}_{1,n}$. The assumptions under which our results hold are typical of loss-based inference. They essentially concern the existence and convergence of projections, as well as the complexity of our working models, expressed in terms of bracketing numbers and integrals. In Section 5, we develop and study a specific example based on the LASSO.

2.5 Further notation

Consider a class \mathcal{F} of real-valued functions and $\phi : \mathbb{R} \rightarrow \mathbb{R}$. If $\phi \circ f$ is well-defined for each $f \in \mathcal{F}$, then we note $\phi(\mathcal{F}) \equiv \{\phi \circ f : f \in \mathcal{F}\}$. Set a semi-norm $\|\cdot\|$ on \mathcal{F} and $\delta > 0$. We denote $N(\delta, \mathcal{F}, \|\cdot\|)$ the δ -bracketing number of \mathcal{F} wrt $\|\cdot\|$ and $J(\delta, \mathcal{F}, \|\cdot\|) \equiv \int_0^\delta \sqrt{\log N(\varepsilon, \mathcal{F}, \|\cdot\|)} d\varepsilon$ the corresponding bracketing integral at δ .

In general, given a known $g \in \mathcal{G}$ and an observation O drawn from $P_{Q_0,g}$, $Z \equiv g(A|W)$ is a deterministic function of g and O . Note that Z should be interpreted as a weight associated with O and will be used as such. Therefore, we can augment O with Z , *i.e.*, substitute (O, Z) for O , while still denoting $(O, Z) \sim P_{Q_0,g}$. In particular, during the course of our trial, conditionally on \mathbf{O}_{i-1} , the randomization scheme g_i is known and we can substitute $(O_i, Z_i) = (O_i, g_i(A_i|W_i)) \sim P_{Q_0,g_i}$ for O_i drawn from P_{Q_0,g_i} . The inverse weights $1/g_i(A_i|W_i)$ are bounded because \mathcal{G}_1 is uniformly bounded away from 0 and 1.

The empirical distribution of \mathbf{O}_n is denoted P_n . For a function $f : \mathcal{O} \times [0, 1] \rightarrow \mathbb{R}^d$, we use the notation $P_n f \equiv n^{-1} \sum_{i=1}^n f(O_i, Z_i)$. Likewise, for any fixed $P_{Q,g} \in \mathcal{M}$, $P_{Q,g} f \equiv E_{P_{Q,g}}(f(O, Z))$ and, for each $i = 1, \dots, n$, $P_{Q_0,g_i} f \equiv E_{P_{Q_0,g_i}}[f(O_i, Z_i)|\mathbf{O}_{i-1}]$, $P_{Q_0,\mathbf{g}_n} f \equiv n^{-1} \sum_{i=1}^n E_{P_{Q_0,g_i}}[f(O_i, Z_i)|\mathbf{O}_{i-1}]$.

The supremum norm of a function $f : \mathcal{O} \times [0, 1] \rightarrow \mathbb{R}^d$ is denoted $\|f\|_\infty$. If $d = 1$ and f is measurable, then the $L^2(P_{Q_0,g^r})$ -norm of f is given by $\|f\|_{2,P_{Q_0,g^r}}^2 \equiv P_{Q_0,g^r} f^2$. If f is only a function of W , then we denote $\|f\|_{2,Q_{W,0}}$ its $L^2(P_{Q_0,g^r})$ -norm, to emphasize that it only depends on the marginal distribution $Q_{W,0}$. With a slight abuse of notation, if f is only a function of (A, W) , then $\|f\|_{2,Q_{W,0}}^2$ is the $L^2(P_{Q_0,g^r})$ -norm of $w \mapsto f(1, w)$. In particular, for $Q_Y, Q'_Y \in \mathcal{Q}_Y$ and $g, g' \in \mathcal{G}$, $\|Q_Y - Q'_Y\|_{2,P_{Q_0,g^r}}^2 = E_{P_{Q_0,g^r}} (Q_Y(A, W) - Q'_Y(A, W))^2$, and $\|g - g'\|_{2,Q_{W,0}}^2 = E_{Q_{W,0}} ((g(1|W) - g'(1|W))^2)$.

2.6 Asymptotics

Our main result rely on the following assumptions.

A1. The conditional distribution of Y given (A, W) under Q_0 is not degenerated.

Existence and convergence of projections.

A2. For each $n \geq 1$, there exists $Q_{Y,\beta_{n,0}} \in \mathcal{Q}_{1,n}$ satisfying

$$P_{Q_0,g^r} L(Q_{Y,\beta_{n,0}}) = \inf_{Q_{Y,\beta} \in \mathcal{Q}_{1,n}} P_{Q_0,g^r} L(Q_{Y,\beta}).$$

There also exists $Q_{Y,\beta_0} \in \mathcal{Q}_1$ such that, for all $\delta > 0$,

$$P_{Q_0,g^r} L(Q_{Y,\beta_0}) < \inf_{\{Q_Y \in \mathcal{Q}_1 : \|Q_Y - Q_{Y,\beta_0}\|_{2,P_{Q_0,g^r}} \geq \delta\}} P_{Q_0,g^r} L(Q_Y).$$

A3. For each $n \geq 1$, there exists $g_{n,0} \in \mathcal{G}_{1,n}$ satisfying

$$P_{Q_0,g^r} L_{Q_{Y,\beta_0}}(g_{n,0})/g^r = \inf_{g \in \mathcal{G}_{1,n}} P_{Q_0,g^r} L_{Q_{Y,\beta_0}}(g)/g^r. \quad (12)$$

There also exists $g_0^* \in \mathcal{G}_1$ such that, for all $\delta > 0$,

$$P_{Q_0, g^r} L_{Q_{Y, \beta_0}}(g_0^*)/g^r < \inf_{\left\{g \in \mathcal{G}_1 : \|g - g_0^*\|_{2, Q_{W, 0}} \geq \delta\right\}} P_{Q_0, g^r} L_{Q_{Y, \beta_0}}(g)/g^r. \quad (13)$$

A4. Assume that Q_{Y, β_0} from **A2** and g_0^* from **A3** exist. For each $\epsilon \in \mathcal{E}$, introduce

$$Q_{Y, \beta_0}(\epsilon) \equiv \text{expit}(\text{logit}(Q_{Y, \beta_0}) + \epsilon H(g_0^*)), \quad (14)$$

where $H(g_0^*)(O) \equiv (2A - 1)/g_0^*(A|W)$. Then, there is a unique $\epsilon_0 \in \mathcal{E}$ such that

$$\epsilon_0 \in \arg \min_{\epsilon \in \mathcal{E}} P_{Q_0, g_0^*} L^{\text{kl}}(Q_{Y, \beta_0}(\epsilon)). \quad (15)$$

Reasoned complexity.

A5. The two entropy conditions $J(1, \mathcal{Q}_{1, n}, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(\sqrt{n})$ and $J(1, L(\mathcal{Q}_{1, n}), \|\cdot\|_{2, P_{Q_0, g^r}}) = o(\sqrt{n})$ hold.

A5*. If $\{\delta_n\}_{n \geq 1}$ is a sequence of positive numbers, then $\delta_n = o(1)$ implies $J(\delta_n, \mathcal{Q}_{1, n}, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$.

A6. The entropy condition $J(1, \mathcal{G}_{1, n}, \|\cdot\|_{2, Q_{W, 0}}) = o(\sqrt{n})$ holds.

A6*. If $\{\delta_n\}_{n \geq 1}$ is a sequence of positive numbers, then $\delta_n = o(1)$ implies $J(\delta_n, \mathcal{G}_n, \|\cdot\|_{2, Q_{W, 0}}) = o(1)$.

Theorem 1 (Asymptotic study of the targeted CARA RCT). *Assume that **A2**, **A3**, **A5** and **A6** are met. Then, the targeted CARA design converges in the sense that $\|g_n - g_0^*\|_{2, Q_{W, 0}} \rightarrow 0$ in probability as $n \rightarrow \infty$. If, in addition, **A4** holds, then the TMLE ψ_n^* consistently estimates ψ_0 . Moreover, if **A1**, **A5*** and **A6*** also hold, then $\sqrt{n/\Sigma_n}(\psi_n^* - \psi_0)$ is approximately standard normally distributed, where Σ_n is the explicit estimator given in (20).*

The last statement in the above theorem underpins the statistical analysis of the proposed targeted CARA RCT. In particular, denoting $\xi_{1-\alpha/2}$ the $(1 - \alpha/2)$ -quantile of the standard normal distribution, the interval $[\psi_n^* \pm \xi_{1-\alpha/2} \sqrt{\Sigma_n/n}]$ is a confidence interval of asymptotic level $(1 - \alpha)$.

3 Comments on Section 2

3.1 A closer look at the parameter of interest and the optimal randomization scheme

Central to our approach is formulating ψ_0 as the value at \mathbf{f}_0 of the mapping $\Upsilon : \mathcal{F} \rightarrow [-1, 1]$ given by (2). Let \mathcal{M} denote the set of all possible distributions of O . Because we slightly changed perspective and now think in terms of distributions $P_{Q, g} \in \mathcal{M}$ instead of $\mathbf{f} \in \mathcal{F}$, it is convenient to introduce the mapping $\Psi : \mathcal{M} \rightarrow [-1, 1]$ characterized by

$$\Psi(P_{Q, g}) \equiv \int (Q_Y(1, w) - Q_Y(0, w)) dQ_W(w) = E_{P_{Q, g}}(Q_Y(1, W) - Q_Y(0, W)).$$

Since Ψ only depends on $P_{Q, g}$ through Q , we will now on systematically write $\Psi(Q)$ in place of $\Psi(P_{Q, g})$ to alleviate notation.

The mapping Ψ is pathwise differentiable. Its efficient influence curve sheds light on the asymptotic properties of all regular and asymptotically linear estimators of $\psi_0 = \Psi(Q_0)$. The latter statement is formalized in the following lemma—we refer the reader to (Bickel, Klaassen, Ritov, and Wellner, 1998, van der Laan and Robins, 2003, van der Vaart, 1998) for definitions and proofs.

Lemma 1. *The mapping $\Psi : \mathcal{M} \rightarrow [-1, 1]$ is pathwise differentiable at every $P_{Q,g} \in \mathcal{M}$ wrt the maximal tangent space. Its efficient influence curve at $P_{Q,g}$, denoted $D(P_{Q,g})$, orthogonally decomposes as $D(P_{Q,g})(O) = D_W(Q)(W) + D_Y(Q_Y, g)(O)$ with*

$$\begin{aligned} D_W(Q)(W) &\equiv Q_Y(1, W) - Q_Y(0, W) - \Psi(Q), \\ D_Y(Q_Y, g)(O) &\equiv \frac{2A - 1}{g(A|W)} (Y - Q_Y(A, W)). \end{aligned}$$

The variance $\text{Var}_{P_{Q,g}}(D(P)(O))$ is a generalized Cramér-Rao lower bound for the asymptotic variance of any regular and asymptotically linear estimator of $\Psi(Q)$ when sampling independently from $P_{Q,g}$. Moreover, if either $Q_Y = Q'_Y$ or $g = g'$ then $E_{P_{Q,g}}(D(P_{Q',g'})(O)) = 0$ implies $\Psi(Q) = \Psi(Q')$.

The last statement of Lemma 1, often referred to as a “double-robustness” property, assures that D can be deployed to safeguard against model mis-specifications when estimating ψ_0 . This is especially relevant in an RCT setting, since the randomization scheme g is known whenever one samples an observation from $P_{Q,g}$.

By Lemma 1, the asymptotic variance of a regular, asymptotically linear estimator under independent sampling from $P_{Q_0,g}$ is lower-bounded by

$$\min_{g \in \mathcal{G}} \text{Var}_{P_{Q_0,g}}(D(P_{Q_0,g})(O)) = \min_{g \in \mathcal{G}} E_{P_{Q_0,g}} \left(\frac{(Y - Q_{Y,0}(A, W))^2}{g^2(A|W)} \right).$$

In this light, targeting g_0 defined by (3) means that the goal of adaptation is to reach a randomization scheme of higher efficiency, *i.e.*, to obtain a valid estimate of ψ_0 using as few blocks of patients as possible. As mentioned in section 2.1, though not used in our approach, g_0 actually has a closed form expression $g_0(1|W) = \sigma_0(1, W)/(\sigma_0(1, W) + \sigma_0(0, W))$, where $\sigma_0^2(A, W)$ is the conditional variance of Y given (A, W) under Q_0 . Under this randomization scheme, the treatment arm with higher probability for a patient with baseline covariates W is the one for which the conditional variance of the response is higher.

3.2 On the data-adaptive loss-based estimation of $Q_{Y,0}$

The reference randomization scheme g^r offers the opportunity to differentially weight each observation in (8). This action impacts the convergence of Q_{Y,β_n} and thus that of g_n , as seen in Sections 2.6 and 2.4 (the limit g_0^* depends on g^r).

As we already emphasized, the working model $\mathcal{Q}_{1,n}$ may depend on sample size n . If it does, then the sequence of working models must be non-decreasing and $\mathcal{Q}_1 \equiv \cup_{n \geq 1} \mathcal{Q}_{1,n}$ can be interpreted as the limiting working model for $Q_{Y,0}$. We would typically recommend to start with $\mathcal{Q}_1 = \dots = \mathcal{Q}_{n_0}$ all equal to a small set, with a user-supplied, deterministic n_0 , then to let the complexity grow with n . It is known, however, that such a growth must remain tethered. Assumptions **A5** and **A5*** provide appropriate conditions on the complexity of $\mathcal{Q}_{1,n}$. We refer to Section 3.5 for a discussion of their meaning.

The combined choice of loss function L and working model $\mathcal{Q}_{1,n}$ determines the technique used to estimate $Q_{Y,0}$. For instance, in the traditional parametric approach, the working model $\mathcal{Q}_{1,n}$ does not depend on n and is indexed by a fixed, finite-dimensional parameter set. Under the LASSO methodology, which we carefully describe and study in Section 5, $\text{logit}(\mathcal{Q}_{1,n})$ is the linear span of a given basis, with constraints on the linear combinations imposed through the definition of B_n .

3.3 On the data-adaptive loss-based estimation of g_0

The optimal randomization scheme g_0 is defined as a minimizer of a certain criterion over the class \mathcal{G} of all randomization schemes, see (3). Thus, our loss-based estimation of g_0 based on \mathbf{O}_n consists of defining g_{n+1} as the minimizer in g of an estimator of the optimality criterion over the user-supplied class of randomization schemes $\mathcal{G}_{1,n}$, see (7) and the next paragraph. This approach is applicable in the largest

generality. Alternatively, if W is discrete, then g_0 takes finitely many values and g_{n+1} can be defined explicitly based on Q_{Y,β_n} and \mathbf{O}_n . This is also the case if one is willing to assign treatment only based on a discrete summary measure V of W . In this context, g_0 is defined as in (3), where the argmin is over the subset of \mathcal{G} consisting of those randomization schemes which depend on W only through V . We refer the readers to (Chambaz and van der Laan, 2013) for details. Note that assigning treatment based on such summary measures is perhaps too restrictive in real-life RCTs where response to treatment may be correlated with a large number of a patient's baseline characteristics, some of which being continuous.

We now turn to the joint justification of (6) and (7). The key point is the following equality, valid for every $g' \in \mathcal{G}$:

$$g_0 = \arg \min_{g \in \mathcal{G}} E_{P_{Q_0, g'}} \left(\frac{(Y - Q_{Y,0}(A, W))^2}{g(A|W)g'(A|W)} \right). \quad (16)$$

Equality (16) tells us that g_0 can be estimated using observations drawn from $P_{Q_0, g'}$ based on the loss function L_{Q_Y} provided it is weighted by $1/g'$. Our observations O_1, \dots, O_n are drawn from $P_{Q_0, g_1}, \dots, P_{Q_0, g_n}$, respectively. In this light, (16) also validates (7).

Like $\mathcal{Q}_{1,n}$, the working model $\mathcal{G}_{1,n} \subset \mathcal{G}$ may depend on sample size n . If it does, then the sequence must be non-decreasing and $\mathcal{G}_1 \equiv \cup_{n \geq 1} \mathcal{G}_{1,n}$ can be interpreted as the limiting working model for g_0 . The additional constraint that \mathcal{G}_1 be uniformly bounded away from 0 and 1 is important. It implies the following pivotal property: no matter how $g^r \in \mathcal{G}$ is chosen in Section 2.3, there exists some constant $\kappa > 0$, such that $\max(\|g^r/g\|_\infty, \|g/g^r\|_\infty) \leq \kappa$, for all $g \in \mathcal{G}_1$. For ease of reference, we call it the *dominated ratio property* of \mathcal{G}_1 .

Using a fixed working model for g_0 , *i.e.*, setting $\mathcal{G}_1 = \mathcal{G}_{1,n}$ for every $n \geq 1$, is a valuable option. However, in some situations, *e.g.* if the population is very heterogeneous, using a fixed, large working model \mathcal{G}_1 may delay, sample size-wise, the adaptation, thereby depriving the trial of the advantages of an adaptive design. By allowing $\mathcal{G}_{1,n}$ to depend on n , one gains the flexibility to enrich the working model for g_0 according to the modesty or generosity of the sample size. Similar to what we suggested for $\mathcal{Q}_{1,n}$ in Section 3.2, we would recommend to start with $\mathcal{G}_1 = \dots = \mathcal{G}_{n_0}$ all equal to a small set, typically the singleton $\{g^b\}$, then to let the complexity of $\mathcal{G}_{1,n}$ augment with n , though not too abruptly. Assumptions **A6** and **A6*** provide appropriate conditions on the complexity of $\mathcal{G}_{1,n}$. We refer to Section 3.5 for a discussion of their meaning. The assumptions are mild, and allow us to use the LASSO to target the optimal randomization scheme g_0 , just like we can use the LASSO to estimate $Q_{Y,0}$, see Section 5.

3.4 On targeted minimum loss estimation

The conception of ψ_n^* defined in (11) follows the paradigm of targeted minimum loss estimation. In the setting of a covariate-adjusted RCT with a fixed design and a fixed working model \mathcal{Q}_1 , a TMLE estimator is unbiased and asymptotically Gaussian regardless of the specification of \mathcal{Q}_1 . Chambaz and van der Laan (2013) show that unbiasedness and asymptotic normality still hold in a framework very similar to that of the present article when the randomization schemes depend on W only through a summary measure taking finitely many values and when \mathcal{Q}_1 is a simple linear model. Such a configuration can be obtained as a particular case of the example developed in Section 5.

Although using a mis-specified parametric working model \mathcal{Q}_1 for $Q_{Y,0}$ does not hinder the consistency of the estimator of ψ_0 , it may affect its efficiency and the convergence of the CARA design to the targeted optimal design. By relying on more flexible randomization schemes and on more adaptive estimators of $Q_{Y,0}$, we may better adapt to the optimal randomization scheme g_0 through better variable adjustments and the targeted construction of the instrumental loss function L_{Q_Y} . Because g_0 is the Neyman design, our approach yields greater efficiency through better variable adjustments and more accurate estimation of the variance of the estimator.

Consider now (9) and (11). The model (9) goes through Q_{Y,β_n} at $\epsilon = 0$ and satisfies the score condition $\frac{\partial}{\partial \epsilon} L^{\text{kl}}(Q_{Y,\beta_n}(\epsilon))|_{\epsilon=0} = D_Y(Q_{Y,\beta_n}, g_n)$. If we set $Q_{\beta_n}^* \equiv (Q_{W,n}, Q_{Y,\beta_n}^*)$, where $Q_{W,n}$ is the empirical marginal distribution of W , then $\psi_n^* = \Psi(Q_{\beta_n}^*)$, assuring that ψ_n^* is indeed a substitution estimator of $\psi_0 = \Psi(Q_0)$.

3.5 On the assumptions

Assumption **A2** stipulates the existence of a projection $Q_{Y,\beta_{n,0}}$ of $Q_{Y,0}$ onto every working model $\mathcal{Q}_{1,n}$. It may depend on the user-supplied reference randomization scheme g^r . If $Q_{Y,0} \in \mathcal{Q}_1$, *i.e.*, if \mathcal{Q}_0 is well-specified, then the existence of $Q_{Y,\beta_0} = Q_{Y,0}$ is granted. If $Q_{Y,0} \notin \mathcal{Q}_1$, *i.e.*, if \mathcal{Q}_1 is mis-specified, then **A2** also stipulates the existence of a projection Q_{Y,β_0} of $Q_{Y,0}$ onto \mathcal{Q}_1 . It may also depend on g^r .

Similar comments apply to **A3**. Note that each $g_{n,0}$ and the limiting randomization scheme g_0^* depend on g^r only through Q_{Y,β_0} : replacing g^r with any arbitrarily chosen $g \in \mathcal{G}$ in (12) or (13) does not alter the values of $g_{n,0}$ and g_0^* . Furthermore, (6) and (13) yield that

$$g_0^* = \arg \min_{g \in \mathcal{G}_1} \left\{ \text{Var}_{P_{Q_0,g}}(D_Y(Q_{Y,0}, g)(O)) + P_{Q_0,g} \frac{(Q_{Y,0} - Q_{Y,\beta_0})^2}{g^2} \right\}.$$

This shows that if $Q_{Y,\beta_0} = Q_{Y,0}$ and $g_0 \in \mathcal{G}_1$, then $g_0^* = g_0$, the optimal randomization scheme. In general, g_0^* minimizes an objective function which is the sum of the Cramér-Rao lower bound and a second-order residual. This underscores the motivation for using a flexible estimator in estimating $Q_{Y,0}$: by minimizing the second-order residual, we get closer to adapting towards the desired optimal randomization criterion.

Recall that Q_{Y,β_n} is characterized by (8) and that $Q_{W,n}$ is the empirical marginal distribution of W . Heuristically, if the equality $Q_{Y,\beta_0} = Q_{Y,0}$ holds then one should be able to prove that $\Psi((Q_{W,n}, Q_{Y,\beta_n}))$ is a consistent estimator of ψ_0 . Since $Q_{Y,\beta_0} = Q_{Y,0}$ also yields that $\epsilon_0 = 0$ is the unique solution to (15) in **A4**, one understands that updating Q_{Y,β_n} to $Q_{Y,\beta_n}^* \equiv Q_{Y,\beta_n}(\epsilon_n)$ and $\Psi((Q_{W,n}, Q_{Y,\beta_n}))$ to ψ_n^* as described in (10) and (11) should preserve the consistency in the initially well-specified framework. In the more likely situation where \mathcal{Q}_1 is mis-specified, hence $Q_{Y,\beta_0} \neq Q_{Y,0}$ and $\epsilon_0 \neq 0$, there is no reason to believe that $\Psi((Q_{W,n}, Q_{Y,\beta_n}))$ should be a consistent estimator of ψ_0 . In this light, the updating procedure bends the inconsistent initial estimator into a consistent one by drawing advantage from the double-robustness of D that we presented in Lemma 1.

In Sections 3.2 and 3.3, we commented on the interest of letting the working models $\mathcal{Q}_{1,n}$ and $\mathcal{G}_{1,n}$ depend on sample size n . Assumptions **A5** and **A6** put very mild constraints on how the complexities of the working models should evolve with n to guarantee the convergence of g_n and consistency of ψ_n^* . The constraints are expressed in terms of bracketing integral. We refer the reader to (van der Vaart, 1998, Examples 19.7-19.11, Lemma 19.15) for typical examples. They include “well-behaved” parametric and Vapnik-Cervonenkis (VC) classes. Assumptions **A5*** and **A6*** should be interpreted as more stringent conditions imposed upon $\mathcal{Q}_{1,n}$ and $\mathcal{G}_{1,n}$. Indeed, for instance,

$$J(1, \mathcal{G}_{1,n}, \|\cdot\|_{2,Q_{W,0}})/\sqrt{n} \leq J(1/\sqrt{n}, \mathcal{G}_{1,n}, \|\cdot\|_{2,Q_{W,0}})$$

because the entropy with bracketing is non-increasing, so that **A6*** does imply **A6** (take $\delta_n = 1/\sqrt{n}$). The need for more stringent conditions arises when studying the convergence in law of ψ_n^* .

4 Building blocks of Theorem 1

We now carry out the theoretical study of the targeted CARA design and its corresponding estimator described in Section 2. All proofs are relegated to Section A.1.

We first focus on the convergence of the estimators Q_{Y,β_n} . The counterpart to this result in the i.i.d. setting is well established (Pollard, 1984, van der Vaart, 1998, among others). The following proposition revises those results for the current statistical setting.

Proposition 1 (convergence of Q_{Y,β_n}). *Under **A2**, **A5**, $\|Q_{Y,\beta_n} - Q_{Y,\beta_0}\|_{2,P_{Q_0,g^r}} \rightarrow 0$ in probability as $n \rightarrow \infty$.*

We now turn to the convergence of the sequence of randomization schemes.

Proposition 2 (convergence of the targeted CARA design). *Under **A2**, **A3**, **A5** and **A6**, it holds that $\|g_n - g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ in probability as $n \rightarrow \infty$.*

The following corollary of Proposition 2 will also prove useful.

Corollary 1. *In the setting of Proposition 2, it also holds that $\|g_n - g_0^*\|_{2, Q_{W,0}}$, $\|1/g_n - 1/g_0^*\|_{2, Q_{W,0}}$, $\|n^{-1} \sum_{i=1}^n g_i - g_0^*\|_{2, Q_{W,0}}$, and $\|n^{-1} \sum_{i=1}^n 1/g_i - 1/g_0^*\|_{2, Q_{W,0}}$ converge to 0 in probability and in L^1 as $n \rightarrow \infty$.*

At this stage, the consistency of ψ_n^* can be established. The proof relies on the convergence of Q_{Y, β_n}^* to a limiting conditional distribution, which is a fluctuation of the limit Q_{Y, β_0} of Q_{Y, β_n} , see Proposition 1.

Proposition 3 (consistency of ψ_n^*). *Suppose that **A2**, **A3**, **A4**, **A5** and **A6** are met. Define*

$$Q_{Y, \beta_0}^* \equiv \text{expit}(\text{logit}(Q_{Y, \beta_0}) + \varepsilon_0 H(g_0^*)), \quad (17)$$

with $H(g_0^)(O) \equiv (2A - 1)/g_0^*(A|W)$ and $Q_{\beta_0}^* \equiv (Q_{W,0}, Q_{Y, \beta_0}^*)$. It holds that $\|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} \rightarrow 0$ in probability as $n \rightarrow \infty$. Moreover, $\Psi(Q_{\beta_0}^*) = \psi_0$ and ψ_n^* consistently estimates ψ_0 .*

We need further notation to state our last building block. For both $\beta = \beta_0$ and $\beta = \beta_n$, introduce $d_{Y, \beta}^*$ given by

$$d_{Y, \beta}^*(O, Z) \equiv \frac{2A - 1}{Z} (Y - Q_{Y, \beta}^*(A, W)). \quad (18)$$

Define also

$$\Sigma_0 \equiv P_{Q_0, g_0^*} (d_{Y, \beta_0}^* + D_W(Q_{\beta_0}^*))^2 = P_{Q_0, g_0^*} \left(D(P_{Q_{\beta_0}^*, g_0^*}) \right)^2, \quad (19)$$

$$\Sigma_n \equiv \frac{1}{n} \sum_{i=1}^n (d_{Y, \beta_n}^*(O_i, Z_i) + D_W(Q_{\beta_n}^*)(W_i))^2, \quad (20)$$

where we recall that $Q_{\beta_n}^* \equiv (Q_{W, n}, Q_{Y, \beta_n}^*)$.

Proposition 4 (asymptotic linearity and central limit theorem for ψ_n^*). *Assume that **A1**–**A6**^{*} are met. Then $\Sigma_n = \Sigma_0 + o_P(1)$ with $\Sigma_0 > 0$, and*

$$\psi_n^* - \psi_0 = (P_n - P_{Q_0, \mathbf{g}_n}) (d_{Y, \beta_0}^* + D_W(Q_{\beta_0}^*)) + o_P(1/\sqrt{n}). \quad (21)$$

Moreover, $\sqrt{\Sigma_n/n}(\psi_n^ - \psi_0)$ converges in law to the standard normal distribution.*

Equality (21) is an asymptotic linear expansion of ψ_n^* under our targeted, adaptive sampling scheme. It is the key to the central limit theorem for $\sqrt{n}(\psi_n^* - \psi_0)$.

5 Example: targeted LASSO-based CARA RCT

In Sections 2, 3 and 4, we have presented a general framework for constructing and analyzing CARA RCTs using data-adaptive loss-based estimators for the nuisance parameters, coupled with the TMLE methodology to estimate the study parameter of interest. As described in Section 1, high-dimensional settings are increasingly common in clinical trials working with heterogeneous populations. A popular device in high-dimensional statistics, due to its computational feasibility and amenability to theoretical study, is the LASSO methodology. In a nutshell, the LASSO is a shrinkage and selection method for generalized regression models that optimizes a loss function of the regression coefficients subject to constraint on the L^1 norm. It was introduced by Tibshirani (1996) for obtaining estimators with fewer nonzero parameter values, thus effectively reducing the number of variables upon which the given solution is dependent. In this section, we illustrate the application of the proposed framework using the LASSO to estimate the conditional response

and the optimal randomization scheme. The methodology introduced in Chambaz and van der Laan (2013) is a special case of this targeted LASSO-based CARA RCT.

For simplicity, we assume that all components of W are continuous. With a little extra work, discrete components could be handled as A is handled in (23).

Let $\ell^1 \equiv \{\beta \in \mathbb{R}^{\mathbb{N}} : \sum_{j \in \mathbb{N}} |\beta^j| < \infty\}$. Consider $\{b_n\}_{n \geq 1}$, $\{b'_n\}_{n \geq 1}$, $\{d_n\}_{n \geq 1}$, and $\{d'_n\}_{n \geq 1}$ four non-decreasing, possibly unbounded sequences over \mathbb{R}_+ and, for some $M, M' > 0$ and every $n \geq 1$, introduce the sets

$$B_n \equiv \{\beta \in \ell^1 : \|\beta\|_1 \leq \min(b_n, M) \text{ and } \forall j > d_n, \beta^j = 0\} \quad (22)$$

and B'_n defined like B_n with b'_n , d'_n and M' substituted for b_n , d_n , and M , respectively. Let $\{\phi_j : j \in \mathbb{N}\}$ be a uniformly bounded set of functions from \mathcal{W} to \mathbb{R} . Without loss of generality, we may assume that $\|\phi_j\|_\infty = 1$ for all $j \in \mathbb{N}$. By choice, the functions ϕ_j ($j \in \mathbb{N}$) share a common bounded support \mathcal{W} , and all belong to the class of sufficiently smooth functions, in the sense that there exists $\alpha > \dim(\mathcal{W})/2$ such that all partial derivatives up to order α of all ϕ_j exist and are uniformly bounded (see van der Vaart, 1998, Example 19.9).

For each $\beta \in \ell^1$, we denote $Q_{Y,\beta} : \mathcal{A} \times \mathcal{W} \rightarrow \mathbb{R}$ and $\gamma_\beta : \mathcal{A} \times \mathcal{W} \rightarrow \mathbb{R}$ the functions characterized by

$$\begin{aligned} Q_{Y,\beta}(A, W) &\equiv \text{expit} \left(\sum_{j \in \mathbb{N}} (\beta^{2j} A + \beta^{2j+1} (1 - A)) \phi_j(W) \right), \\ \gamma_\beta(1|W) = 1 - \gamma_\beta(0|W) &\equiv \text{expit} \left(\sum_{j \in \mathbb{N}} \beta^j \phi_j(W) \right). \end{aligned} \quad (23)$$

The LASSO-based CARA RCT design corresponds to a special choice of working models $\{\mathcal{Q}_{1,n}\}_{n \geq 1}$, $\{\mathcal{G}_{1,n}\}_{n \geq 1}$, and loss function L for $Q_{Y,0}$. We take $\mathcal{Q}_{1,n} \equiv \{Q_{Y,\beta} : \beta \in B_n\}$ with M a deterministic upper-bound on $|\logit(Y)|$ and the quasi negative-log-likelihood loss function $L = L^{\text{kl}}$ (4). Note that the elements of $\mathcal{Q}_1 \equiv \cup_{n \geq 1} \mathcal{Q}_{1,n}$ are uniformly bounded away from 0 and 1. We also take $\mathcal{G}_{1,n} \equiv \{\gamma_\beta : \beta \in B'_n\}$. The elements of $\mathcal{G}_1 \equiv \cup_{n \geq 1} \mathcal{G}_{1,n}$ are randomization schemes uniformly bounded away from 0 and 1 by $\text{expit}(-M')$ and $\text{expit}(M')$, respectively ($M' \simeq 4.6$ provides the lower- and upper-bounds 0.01 and 0.99).

Based on \mathbf{O}_n , we estimate $Q_{Y,0}$ with $Q_{Y,\beta_{n+1}}$, where β_{n+1} is given in (8) (set $i = n + 1$ in the formula). Then we target g_0 with g_{n+1} given in (7) (set $i = n + 1$ in the formula), *i.e.*

$$g_{n+1} \in \arg \min_{\beta \in B'_n} \frac{1}{n} \sum_{j=1}^n \frac{L_{Q_{Y,\beta_n}}(\gamma_\beta)(O_j)}{g_j(A_j|W_j)}. \quad (24)$$

The minimization (8) with the constraint $\|\beta\|_1 \leq \min(b_n, M)$, see (22), can be rewritten as a minimization free of the latter constraint by adding a term of the form $\lambda_n \|\beta\|_1$ to the empirical criterion, where λ_n depends on b_n . Note that when d_n or d'_n is held constant and M or M' is infinite by choice, then (8) or (24) should be interpreted as a standard parametric procedure rather than as a LASSO.

Theorem 1 has the following corollary.

Corollary 2 (asymptotic study of the targeted LASSO-based CARA RCT). *Assume that **A1**, **A2**, **A3**, and **A4** are met. Then, the targeted LASSO-based CARA design converges in the sense that $\|g_n - g_0^*\|_{2,Q_{W,0}} \rightarrow 0$ in probability as $n \rightarrow \infty$. Moreover, the TMLE ψ_n^* consistently estimates ψ_0 , and $\sqrt{n}/\Sigma_n(\psi_n^* - \psi_0)$ is approximately standard normally distributed, where Σ_n is the explicit estimator given in (20).*

This corollary teaches us with minimal conditions on the smoothness of the basis functions, the targeted LASSO-based CARA RCT produces a convergent design and a consistent and asymptotically Gaussian estimator for the study parameter.

6 Simulation study

In this section, we exemplify the theoretical results from the previous sections with a brief simulation study. Specifically, we wish to (i) illustrate the robustness of the proposed TMLE estimator for the study parameter ψ_0 , under possibly grossly mis-specified conditional response models, (ii) show the use of data-adaptive LASSO estimators to learn the conditional response in the construction and analysis of the targeted CARA RCT, and (iii) evaluate the performances of the different strategies. The simulation study is conducted using R (R Core Team, 2014).

6.1 Data-generating distribution

Under Q_0 , $W = (U, V, Z_1, \dots, Z_{20})$ consists of 22 independent random variables, where U, Z_1, \dots, Z_{20} are all uniformly distributed on $[0, 1]$, and $V \in \{1, 2, 3\}$ is such that $V = 1$, $V = 2$ and $V = 3$ with probabilities $1/2$, $1/3$, and $1/6$, respectively. Moreover, under Q_0 and conditionally on (A, W) , Y is drawn from the Gamma distribution with conditional mean

$$Q_{Y,0}(A, W) \equiv 2AV + (1 - A)V/2$$

and conditional variance

$$\sigma_0^2(A, W) \equiv \left(\frac{AV}{3(1 + Z_1)} + \frac{4(1 - A)}{3(1 + Z_1)} \right)^2.$$

It is easy to check that $\psi_0 = 2.5$ and that the optimal randomization scheme g_0 is given by $g_0(1|W) \equiv V/(4 + V)$.

6.2 Loss functions and working models

To simplify the language, we refer to a model that accounts for the relevant covariates as a correctly specified model, even though the functional form may not be correct.

Estimation of the conditional response

Because Y is continuous and unbounded, we perform a linear transformation before the estimation procedures to scale Y within $(0, 1)$, then apply the reverse transformation to the final TMLE estimate of ψ_0 and the corresponding variance estimates. We use the quasi negative-log-likelihood loss function L^{kl} given by (4).

At sample size n , we consider two working models $\mathcal{Q}_{1,n}$ for the conditional response. One is the following mis-specified logistic regression model:

$$\mathcal{Q}_{1,n}^p \equiv \left\{ Q_{Y,\beta}^p(A, W) \equiv \text{expit}(\beta_1 A + \beta_2 U) : \beta \in \mathbb{R}^2 \right\}.$$

Contrary to what the notation suggests, it does not change as the sample size grows. It is fitted using the `glm` function in R with the weights as given in (8). Note that the model fails to take into account the covariate V which drives the response in the underlying data-generating process.

The second one, denoted $\mathcal{Q}_{1,n}^\ell$, is a LASSO logistic working model. Let $d_n \equiv \min(20, \lfloor \sqrt{n}/4 \rfloor)$. If n is such that $d_n \leq 5$, then $\mathcal{Q}_{1,n}^\ell$ consists of

$$Q_{Y,\beta}^\ell(A, W) \equiv \text{expit} \left(\beta(A, U, Z_1, \dots, Z_{d_n}, AU, AZ_1, \dots, AZ_{d_n})^\top \right) \quad (\text{all } \beta \in B_n \equiv \mathbb{R}^{2d_n+3}).$$

If n is such that $d_n > 5$, then $\mathcal{Q}_{1,n}^\ell$ consists of

$$Q_{Y,\beta}^\ell(A, W) \equiv \text{expit} \left(\beta(A, U, V, Z_1, \dots, Z_{d_n}, AU, AV, AZ_1, \dots, AZ_{d_n})^\top \right) \quad (\text{all } \beta \in B_n \equiv \mathbb{R}^{2d_n+5}).$$

The resulting sequence of working models is non-decreasing in sample size. The models is fitted using the `cv.glmnet` function from the package `glmnet` (Friedman, Hastie, and Tibshirani, 2010), with weights given in (8) and the option `"lambda.1se"`.

Estimation of the optimal randomization scheme g_0

We also consider two working models $\mathcal{G}_{1,n} = \mathcal{G}_1$ for the optimal randomization scheme. The first one, denoted \mathcal{G}_1^m , is a mis-specified logistic model given by

$$g_\beta^m(A = 1 \mid W) \equiv \text{expit}(\beta_0 + \beta_1 U) \quad (\text{all } \beta \in \mathbb{R}^2).$$

The second one, denoted \mathcal{G}_1^c , is a correctly specified logistic model given by

$$g_\beta^c(A = 1 \mid W) \equiv \text{expit}(\beta_0 + \beta_1 U + \beta_2 V) \quad (\text{all } \beta \in \mathbb{R}^3).$$

The models are fitted using numerical methods to optimize the user-chosen adaptation criterion in (6). We implement this fitting using the `optim` function with a quasi-Newton method (`method="BFGS"`). To satisfy the boundedness conditions, the resulting probability estimates are pre-specified to be truncated to lie within $[0.05, 0.95]$. However, in the actual simulation runs, all estimates lying comfortably within this interval, and hence no truncation took place.

6.3 Study designs

For each pair of working models for the conditional response and for the optimal randomization scheme, we construct a CARA RCT by initializing at a sample size of $n = 300$, and then sequentially recruiting patients in blocks of size 200, up to $n = 3100$. For the initial sample of $n = 300$, treatment is randomly assigned based on the balanced randomization scheme g^b . Subsequently, given n observations, we estimate the conditional response and use this to construct the treatment randomization scheme g_{n+1} used for the next block of 200 patients. We also use this conditional response estimate and the sequence of randomization schemes used so far to obtain a TMLE estimate ψ_n^* of ψ_0 .

In addition to these CARA RCTs, we also consider a fixed design RCT with treatment randomly assigned based on the balanced randomization scheme g^b . We obtain the corresponding TMLE estimates by fluctuating the initial conditional response estimates based on the logistic model $\{Q_{Y,\beta}^p : \beta \in \mathbb{R}^2\}$.

6.4 Results

For each trial design proposed in Section 6.3, we run 500 independent simulated trials. Three figures summarize the results of the simulation study. Each of them consists of two similar graphics, the LHS graphic corresponding to the simulated trials based on the mis-specified model \mathcal{G}_1^m for the optimal randomization scheme, and the RHS graphic to the simulated trials based on the correctly specified model \mathcal{G}_1^c . The subtitles “A~U” and “A~U+V” are the R formulas that encode for \mathcal{G}_1^m and \mathcal{G}_1^c , respectively.

Figure 1 depicts the performance of ψ_n^* in terms of bias (first row), sample variance (second row) and mean squared error (MSE, third row). We note that, despite the mis-specified response models, all TMLE estimators are consistent for the treatment effect parameter ψ_0 . It appears that the LASSO-based estimator may converge at a faster rate. This may be due to its increased efficiency (*i.e.*, smaller sample variance) and more aggressive bias reduction. Recall that the optimality criterion for our adaptive randomization aims at maximizing efficiency of the trial through the minimization of the asymptotic variance of the estimators. The increased efficiency of the LASSO-based CARA RCT, despite a larger working model for the conditional response (increasing with sample size), suggests that a flexible data-adaptive response model coupled with CARA design could indeed better achieve the optimality criterion, compared to a CARA design based on a parametric response model, at least in situations where the parametric model fails to account for important confounding variables. We also note that, under the data-generating process described in Section 6.1, the working model for the optimal randomization scheme has little effect on the efficiency of the TMLE estimators. Yet, comparing the LHS and RHS graphics in Figure 1 suggests that \mathcal{G}_1^m , the smaller, mis-specified model for the optimal randomization scheme allows for slightly more aggressive bias reduction at smaller sample sizes than \mathcal{G}_1^c , its larger, correctly specified counterpart.

Let us turn now to the coverage of our CLT-based, 95%-confidence intervals (CIs). The empirical coverage probabilities are depicted in Figure 2. On the one hand, we see that the empirical coverages are often below the nominal coverage when using the mis-specified working model \mathcal{G}_1^m for the optimal randomization scheme and either $\mathcal{Q}_{1,n}^p$ or $\mathcal{Q}_{1,n}^\ell$ as working models for the conditional response (LHS graphic in Figure 2). On the other hand, the coverage improves drastically when using the correctly specified working model \mathcal{G}_1^c for the optimal randomization scheme and either $\mathcal{Q}_{1,n}^p$ or $\mathcal{Q}_{1,n}^\ell$ as working models for the conditional response (RHS graphic in Figure 2). For a more precise assessment, we frame the coverage evaluation in terms of hypotheses testing. For a given design (and its resulting CLT-based CIs) and at each intermediate sample size n , let C be the number of times in the 500 simulations when the CI covers the parameter of interest ψ_0 . The random variable C is distributed from the Binomial distribution with parameter $(500, \pi)$. For a given significance level $0 < \alpha < 1$, introduce the null hypotheses $H_0^{1-\alpha} : \pi \geq 1 - \alpha$ and its one-sided alternative $H_1^{1-\alpha} : \pi < 1 - \alpha$. We perform one-sided tests of $H_0^{1-\alpha}$ against $H_1^{1-\alpha}$ and display the p -values for $\alpha = 5\%$ (Figure 3, first row) and $\alpha = 6\%$ (Figure 3, second row). On the one hand, the LHS graphic in Figure 3 reveals that 95% coverage is often not guaranteed when using the mis-specified working model $\mathcal{G}_{1,n}^m$ for the optimal randomization scheme and either $\mathcal{Q}_{1,n}^p$ or $\mathcal{Q}_{1,n}^\ell$ as working models for the conditional response, but also that 94% coverage cannot be ruled out. On the other hand, the RHS graphic in Figure 3 suggests that 95% coverage cannot be ruled out when using the correctly specified model $\mathcal{G}_{1,n}^c$ for the optimal randomization scheme and either $\mathcal{Q}_{1,n}^p$ or $\mathcal{Q}_{1,n}^\ell$ as working models for the conditional response.

7 Discussion

We have presented in this article a new group-sequential CARA RCT design and corresponding analytical procedure that admits the use of flexible data-adaptive techniques. The proposed method extends the work of Chambaz and van der Laan (2013) by providing robust inference of the study parameter under the most general settings. Our framework adopts a loss-based approach in estimating the optimal randomization scheme, and hence can target general optimality criteria that may not have a closed-form solutions. Moreover, our use of loss-based data-adaptive estimation over general classes of functions (which may change with sample size), both in constructing the treatment randomization schemes and in predicting the unknown conditional response, may potentially improve the randomization adaptation towards the optimality criterion.

We established that, under appropriate entropy conditions on the classes of functions, the resulting sequence of randomization schemes converges to a fixed scheme, and the proposed treatment effect estimator is consistent (even under a mis-specified response model), asymptotically Gaussian, giving rise to valid confidence intervals of given asymptotic levels. Moreover, the limiting randomization scheme coincides with the unknown optimal randomization scheme when, simultaneously, the response model is correctly specified and the optimal randomization scheme belongs to the limit of the user-supplied classes of randomization schemes. We illustrated the applicability of these general theoretical results with a LASSO-based CARA RCT. In this example, both the response model and the optimal treatment randomization are estimated using a sequence of LASSO logistic models that may increase with sample size. It follows immediately from our general theorems that this LASSO-based CARA RCT converges to a fixed randomization scheme and yields consistent and asymptotically Gaussian effect estimates, under minimal conditions on the smoothness of the basis functions in the LASSO logistic models.

We conducted a simulation study to evaluate the performance of the proposed methods. It confirmed the robustness of the TMLE estimators under mis-specified response models. Coverage of the CLT-based confidence intervals are assessed through by hypotheses testing. Overall there is no evidence (across 500 independent simulations) that the 95%-confidence intervals would have coverages that are less than 94%. In addition, we do observe improved coverage when using the correct working model for the optimal randomization scheme. In this simulation study, the increased efficiency of CARA design with a LASSO-based response model, compared to the CARA (or balanced) design with a parametric response model, demonstrates that the use of data-adaptive response models can indeed more effectively steer the adaptation towards the optimality criterion (which was chosen to be efficiency in our example). More comprehensive empirical studies

are needed to generalize these facts to other simulation scenarios.

We will soon make available a R package to allow interested readers to test the procedure. In the future, we will also consider alternative strategies to randomly assign successive patients to the treatment arms in such a way that the overall empirical conditional distribution of treatment given baseline covariates be as close as possible to the current best estimator of the targeted optimal randomization scheme. This will require both new theoretical developments and simulation studies.

A Appendix

The expression “ $a \lesssim b$ ” means that there exists a universal, positive constant c such that $a \leq c \times b$. We use $\mathbf{1}\{\mathcal{C}\}$ to denote the indicator function of the set \mathcal{C} . We denote the uniform norm of a real-valued operator Π on \mathcal{F} as $\|\Pi\|_{\mathcal{F}} \equiv \sup_{f \in \mathcal{F}} |\Pi(f)|$. Given two measurable functions f, λ of (O, Z) and the random variable $\Lambda = \lambda(O, Z)$, we find it convenient to use shorthand notation $P_{Q_0, g} f \Lambda \equiv E_{P_{Q_0, g}}(f(O, Z)\Lambda)$ and $P_n f \Lambda \equiv E_{P_n}(f(O, Z)\Lambda) = n^{-1} \sum_{i=1}^n f(O_i, Z_i) \lambda(O_i, Z_i)$. From here onward, the uncountable supremum is interpreted as the essential supremum.

Section A.1 presents the proofs of Propositions 1, 2, 3, 4, Corollary 1, Theorem 1 and Corollary 2. Technical results underpinning the proofs of Section A.1 are gathered in Section A.2.

A.1 Main proofs

Proof of Proposition 1. We apply Lemma 4 with $\Theta \equiv \mathcal{Q}_1$, $\Theta_n \equiv \mathcal{Q}_{1,n}$, d the distance induced on Θ by the norm $\|\cdot\|_{2, P_{Q_0, g^r}}$, \mathbf{M} and \mathbf{M}_n characterized over Θ by $\mathbf{M}(Q_Y) \equiv P_{Q_0, g^r} L(Q_Y)$ and $\mathbf{M}_n(Q_Y) \equiv P_n L(Q_Y) g^r / Z = n^{-1} \sum_{i=1}^n L(Q_Y)(O_i) g^r(A_i | W_i) / Z_i$. Assumption **A2** implies that (a) and (b) from Lemma 4 are met. It remains to prove that (c) also holds or, in other terms, that $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{Q}_{1,n}} = o_P(1)$.

For any $Q_Y \in \Theta$, characterize $\ell(Q_Y)$ by setting $\ell(Q_Y)(O, Z) \equiv L(Q_Y)(O) g^r(A|W) / Z$. Then we can rewrite $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{Q}_{1,n}}$ as follows:

$$\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{Q}_{1,n}} = \|P_n \ell - P_{Q_0, g^r} L\|_{\mathcal{Q}_{1,n}} = \|(P_n - P_{Q_0, \mathbf{g}_n}) \ell\|_{\mathcal{Q}_{1,n}} = \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{Q}_{1,n})}.$$

The dominated ratio property implies that $J(1, \ell(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}}) = O(J(1, L(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}})) = o(\sqrt{n})$, by **A5**. Since $\ell(\mathcal{Q}_1)$ is uniformly bounded by construction, Lemma 8 applies and yields $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{Q}_{1,n})} = o_P(1)$.

Thus, we can apply Lemma 4. It yields that $\|Q_{Y, \beta_n} - Q_{Y, \beta_0}\|_{2, P_{Q_0, g^r}} = o_P(1)$, which is the desired result. \square

The next proof goes along similar lines.

Proof of Proposition 2. We apply Lemma 4 with $\Theta \equiv \mathcal{G}_1$, $\Theta_n \equiv \mathcal{G}_{1,n}$, d the distance induced on Θ by the norm $\|\cdot\|_{2, Q_{W, 0}}$, \mathbf{M} and \mathbf{M}_n characterized over Θ by $\mathbf{M}(g) \equiv P_{Q_0, g^r} L_{Q_{Y, \beta_0}}(g) / g^r$ and $\mathbf{M}_n(g) \equiv P_n L_{Q_{Y, \beta_n}}(g) / Z = n^{-1} \sum_{i=1}^n L_{Q_{Y, \beta_n}}(g)(O_i) / Z_i$. Assumption **A3** implies that (a) and (b) from Lemma 4 are met. It remains to prove that (c) also holds or, in other terms, that $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{G}_{1,n}} = o_P(1)$.

Let ℓ and ℓ_n be characterized over \mathcal{G}_1 by $\ell(g)(O, Z) \equiv L_{Q_{Y, \beta_0}}(g)(O) / Z$ on the one hand and $\ell_n(g)(O, Z) \equiv L_{Q_{Y, \beta_n}}(g)(O) / Z$ on the other hand. A simple decomposition and the triangle inequality yield the following inequality:

$$\begin{aligned} \|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{G}_{1,n}} &= \|(P_n \ell - P_{Q_0, g^r} L_{Q_{Y, \beta_0}} / g^r) + P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} \\ &\leq \|P_n \ell - P_{Q_0, g^r} L_{Q_{Y, \beta_0}} / g^r\|_{\mathcal{G}_{1,n}} + \|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} \\ &= \|(P_n - P_{Q_0, \mathbf{g}_n}) \ell\|_{\mathcal{G}_{1,n}} + \|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} \\ &= \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{G}_{1,n})} + \|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}}. \end{aligned} \tag{25}$$

Consider the first RHS term in (25). Because Y and Q_{Y,β_0} are bounded, and because \mathcal{G}_1 is bounded away from 0 and 1 by construction, it holds that $J(1, \ell(\mathcal{G}_{1,n}), \|\cdot\|_{2,Q_{W,0}}) = O(J(1, 1/\mathcal{G}_{1,n}, \|\cdot\|_{2,Q_{W,0}})) = o(\sqrt{n})$ by **A6**. Since $\ell(\mathcal{G}_1)$ is uniformly bounded, Lemma 8 applies and yields $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{G}_{1,n})} = o_P(1)$.

We now turn to the second RHS term in (25). Note $|L^{\text{ls}}(Q_{Y,\beta_n}) - L^{\text{ls}}(Q_{Y,\beta_0})| \lesssim |Q_{Y,\beta_n} - Q_{Y,\beta_0}|$ because Y is bounded and \mathcal{G}_1 is uniformly bounded. This justifies the second inequality below, the first one being a consequence of the uniform boundedness of $\mathcal{G}_{1,n}$, and the last one a consequence of the fact that g^r is bounded away from 0:

$$\begin{aligned} \|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} &\lesssim P_n |L^{\text{ls}}(Q_{Y,\beta_n}) - L^{\text{ls}}(Q_{Y,\beta_0})|/Z \\ &\lesssim P_n |Q_{Y,\beta_n} - Q_{Y,\beta_0}|/Z \\ &= P_{Q_0, \mathbf{g}_n} |Q_{Y,\beta_n} - Q_{Y,\beta_0}|/Z + (P_n - P_{Q_0, \mathbf{g}_n}) |Q_{Y,\beta_n} - Q_{Y,\beta_0}|/Z \\ &\lesssim P_{Q_0, g^r} |Q_{Y,\beta_n} - Q_{Y,\beta_0}| + (P_n - P_{Q_0, \mathbf{g}_n}) |Q_{Y,\beta_n} - Q_{Y,\beta_0}|/Z. \end{aligned}$$

The Cauchy-Schwarz inequality implies that $P_{Q_0, g^r} |Q_{Y,\beta_n} - Q_{Y,\beta_0}| \leq \|Q_{Y,\beta_n} - Q_{Y,\beta_0}\|_{2, P_{Q_0, g^r}} = o_P(1)$ by Proposition 1, whose assumptions are met. For any $Q_Y \in \mathcal{Q}_1$, introduce $h(Q_Y)$ characterized by $h(Q_Y)(O, Z) \equiv |Q_{Y,\beta_n}(A, W) - Q_{Y,\beta_0}(A, W)|/Z$. Obviously,

$$|(P_n - P_{Q_0, \mathbf{g}_n}) |Q_{Y,\beta_n} - Q_{Y,\beta_0}|/Z| \leq \|(P_n - P_{Q_0, \mathbf{g}_n})h\|_{\mathcal{Q}_{1,n}} = \|P_n - P_{Q_0, \mathbf{g}_n}\|_{h(\mathcal{Q}_{1,n})}.$$

Since \mathcal{Q}_1 and \mathcal{G}_1 are uniformly bounded away from 0 and 1 by construction, it holds that $h(\mathcal{Q}_1)$ is uniformly bounded and that $J(1, h(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}}) = O(J(1, \{|Q_Y - Q_{Y,\beta_0} : Q_Y \in \mathcal{Q}_{1,n}\}, \|\cdot\|_{2, P_{Q_0, g^r}})) = o(\sqrt{n})$ by **A5**. Therefore, Lemma 8 applies and yields $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{h(\mathcal{Q}_{1,n})} = o_P(1)$.

We thus have showed that both $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{G}_{1,n})} = o_P(1)$ and $\|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} = o_P(1)$, hence $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{G}_{1,n}} = o_P(1)$ in light of (25). Consequently, we can apply Lemma 4. It yields that $\|g_n - g_0^*\|_{2, Q_{W,0}} = o_P(1)$, which is the desired result. \square

Proof of Corollary 1. Since \mathcal{G}_1 is uniformly bounded, $\|g_n - g_0^*\|_{2, Q_{W,0}} = o_P(1)$ implies $\|g_n - g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ in L^1 as $n \rightarrow \infty$. Since (i) $1/g_n - 1/g_0^* = (g_0^* - g_n)/g_n g_0^*$, and (ii) \mathcal{G}_1 is uniformly bounded away from 0 and 1, $\|1/g_n - 1/g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ follows from $\|g_n - g_0^*\|_{2, Q_{W,0}} \rightarrow 0$, both in probability and in L^1 as $n \rightarrow \infty$. Consider now the L^1 -convergence of $\|n^{-1} \sum_{i=1}^n g_i - g_0^*\|_{2, Q_{W,0}}$. By convexity,

$$E \left(\left\| \frac{1}{n} \sum_{i=1}^n g_i - g_0^* \right\|_{2, Q_{W,0}} \right) \leq \frac{1}{n} \sum_{i=1}^n E (\|g_i - g_0^*\|_{2, Q_{W,0}}).$$

We already know that $E(\|g_n - g_0^*\|_{2, Q_{W,0}}) = o(1)$. Applying Cesaro's lemma yields that $n^{-1} \sum_{i=1}^n E(\|g_i - g_0^*\|_{2, Q_{W,0}}) = o(1)$, too. From this, we deduce that $\|n^{-1} \sum_{i=1}^n g_i - g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ in L^1 as $n \rightarrow \infty$. This implies that the convergence also holds in probability because \mathcal{G}_1 is uniformly bounded. Likewise,

$$E \left(\left\| \frac{1}{n} \sum_{i=1}^n 1/g_i - 1/g_0^* \right\|_{2, Q_{W,0}} \right) \leq \frac{1}{n} \sum_{i=1}^n E (\|1/g_i - 1/g_0^*\|_{2, Q_{W,0}}),$$

where $E(\|1/g_n - 1/g_0^*\|_{2, Q_{W,0}}) = o(1)$ is already known. Thus, the same argument as above yields that $\|n^{-1} \sum_{i=1}^n 1/g_i - 1/g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ in L^1 and in probability as $n \rightarrow \infty$. This completes the proof. \square

Proof of Proposition 3. This is a three-part proof. First, we show that $|\epsilon_n - \epsilon_0| = o_P(1)$. Second, we prove that $\|Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*\|_{2, P_{Q_0, g^r}} = o_P(1)$. Third, we demonstrate that $\Psi(Q_{\beta_0}^*) = \psi_0$, then that ψ_n^* consistently estimates ψ_0 .

We apply (van der Vaart, 1998, Theorem 5.9) (substituting \mathbf{M}_n and \mathbf{M} for Ψ_n and Ψ) with $\Theta \equiv \mathcal{E}$, d the Euclidean distance, \mathbf{M} and \mathbf{M}_n characterized over Θ by $\mathbf{M}(\epsilon) = P_{Q_0, g_0^*} D_Y(Q_{Y,\beta_0}(\epsilon), g_0^*)$, and $\mathbf{M}_n(\epsilon) = P_n D_Y(Q_{Y,\beta_n}(\epsilon), g_n) g_n / Z$, see (14) and (9) for the definitions of $Q_{Y,\beta_0}(\epsilon)$ and $Q_{Y,\beta_n}(\epsilon)$. From the differentiability of $\epsilon \mapsto L^{\text{kl}}(Q_{Y,\beta}(\epsilon))$, validity of the differentiation under the integral sign, and definition of

ϵ_0 (15), we deduce that $\mathbf{M}(\epsilon_0) = 0$. By definition of ϵ_n (10), $\mathbf{M}_n(\epsilon_n) = 0$ too. Assumption **A4** implies that the second condition of the theorem is met. Therefore it suffices to check that the first one holds too, *i.e.* to prove that $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{E}} = o_P(1)$.

Introduce $\mathcal{F} = \{f_\epsilon : \epsilon \in \mathcal{E}\}$ with $f_\epsilon(O, Z) \equiv (2A - 1)(Y - Q_{Y,\beta_0}(\epsilon)(A, W))/Z$ for each $\epsilon \in \mathcal{E}$. We start with the following derivation:

$$\begin{aligned} \|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{E}} &= \sup_{\epsilon \in \mathcal{E}} \left| P_n \left(f_\epsilon + \frac{2A-1}{Z} (Q_{Y,\beta_0}(\epsilon) - Q_{Y,\beta_n}(\epsilon)) \right) - P_{Q_0, \mathbf{g}_n} f_\epsilon \right| \\ &\leq \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}} + \sup_{\epsilon \in \mathcal{E}} \left| P_n \frac{2A-1}{Z} (Q_{Y,\beta_0}(\epsilon) - Q_{Y,\beta_n}(\epsilon)) \right|. \end{aligned} \quad (26)$$

Consider the first RHS term in (26). Set $\epsilon_1, \epsilon_2 \in \mathcal{E}$. Because the expit function is 1-Lipschitz and \mathcal{G}_1 is uniformly bounded, it holds that $\|f_{\epsilon_1} - f_{\epsilon_2}\|_{\infty} \lesssim |\epsilon_1 - \epsilon_2|$. Since \mathcal{E} is a bounded set by construction, the uniformly bounded, parametric class \mathcal{F} satisfies $J(1, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, \mathbf{g}_r}}) < \infty$ (see van der Vaart, 1998, Example 19.7). Consequently, we can apply Lemma 8 (with a fixed class) and conclude that $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}} = o_P(1)$.

Next, the second term in the RHS of (26) is upper-bounded by $\Delta_n \equiv \sup_{\epsilon \in \mathcal{E}} P_n |Q_{Y,\beta_0}(\epsilon) - Q_{Y,\beta_n}(\epsilon)|/Z$. Since (i) expit is 1-Lipschitz, (ii) $\mathcal{Q}_{1,n}$ is bounded away from 0 and 1, and logit is Lipschitz on any compact subset of $]0, 1[$, it holds that

$$\begin{aligned} \Delta_n &\leq \sup_{\epsilon \in \mathcal{E}} P_n |\text{logit}(Q_{Y,\beta_0}) - \text{logit}(Q_{Y,\beta_n}) + \epsilon(H(g_n) - H(g_0^*))|/Z \\ &\lesssim P_n |Q_{Y,\beta_0} - Q_{Y,\beta_n}|/Z + P_n |1/g_n - 1/g_0^*|/Z \\ &= (P_n - P_{Q_0, \mathbf{g}_n}) |Q_{Y,\beta_0} - Q_{Y,\beta_n}|/Z + P_{Q_0, \mathbf{g}_n} |Q_{Y,\beta_0} - Q_{Y,\beta_n}|/Z \\ &\quad + (P_n - P_{Q_0, \mathbf{g}_n}) |1/g_n - 1/g_0^*|/Z + P_{Q_0, \mathbf{g}_n} |1/g_n - 1/g_0^*|/Z. \end{aligned} \quad (27)$$

While studying the second RHS term of (25) in the proof of Proposition 2, we proved the following facts: $P_{Q_0, \mathbf{g}_n} |Q_{Y,\beta_0} - Q_{Y,\beta_n}|/Z \lesssim P_{Q_0, \mathbf{g}^r} |Q_{Y,\beta_0} - Q_{Y,\beta_n}| = o_P(1)$ and $(P_n - P_{Q_0, \mathbf{g}_n}) |Q_{Y,\beta_0} - Q_{Y,\beta_n}|/Z = o_P(1)$ (the assumptions of Proposition 2 are met here too). Therefore, it only remains to study the two rightmost terms in the RHS of (27). Since \mathcal{G}_1 is uniformly bounded away from 0, $(P_n - P_{Q_0, \mathbf{g}_n}) |1/g_n - 1/g_0^*|/Z = O(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{1/\mathcal{G}_{1,n}})$. Moreover, Lemma 8 applies because $1/\mathcal{G}_1$ is uniformly bounded and **A6** is met, hence $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{1/\mathcal{G}_{1,n}} = o_P(1)$ and $(P_n - P_{Q_0, \mathbf{g}_n}) |1/g_n - 1/g_0^*|/Z = o_P(1)$. Finally,

$$P_{Q_0, \mathbf{g}_n} |1/g_n - 1/g_0^*|/Z \lesssim P_{Q_0, \mathbf{g}^r} |1/g_n - 1/g_0^*| \leq \|1/g_n - 1/g_0^*\|_{2, Q_{W,0}} = o_P(1)$$

by Cauchy-Schwarz and Corollary 1, whose assumptions are met here too. In summary, $\Delta_n = o_P(1)$.

We have show that the RHS expression in (26) converges to 0 in probability as $n \rightarrow \infty$, hence $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{E}} = o_P(1)$. Thus, all assumptions of Lemma 4 hold, from which we deduce that ϵ_n converges to ϵ_0 as $n \rightarrow \infty$. This completes the first part of the proof.

Let $\mathcal{Q}_Y \times \mathcal{G} \times \mathcal{E}$ be equipped with the norm characterized by $\|(Q_Y, g, \epsilon)\| = \|Q_Y\|_{2, P_{Q_0, \mathbf{g}^r}} + \|g\|_{2, Q_{W,0}} + |\epsilon|$. Propositions 1, 2 and the first part of the proof imply that $(Q_{Y,\beta_n}, g_n, \epsilon_n)$ converges to $(Q_{Y,\beta_0}, g_0^*, \epsilon_0)$ in probability wrt $\|\cdot\|$ as $n \rightarrow \infty$. Let $f : \mathcal{Q}_Y \times \mathcal{G} \times \mathcal{E} \rightarrow \mathcal{Q}_Y$ be characterized by

$$f(Q_Y, g, \epsilon)(O) \equiv \text{expit}(\text{logit}(Q_Y(A, W)) + \epsilon(2A - 1)/g(A|W)) \quad (28)$$

Set $(Q_{Y,1}, g_1, \epsilon_1), (Q_{Y,2}, g_2, \epsilon_2) \in \mathcal{Q}_Y \times \mathcal{G} \times \mathcal{E}$. Because (i) expit is 1-Lipschitz, (ii) $\mathcal{Q}_{1,n}$ is bounded away from 0 and 1, and logit is Lipschitz on any compact subset of $]0, 1[$, (iii) \mathcal{G}_1 is uniformly bounded away from 0, (iv) \mathcal{E} is a bounded set, it holds that

$$\begin{aligned} &\|f(Q_{Y,1}, g_1, \epsilon_1) - f(Q_{Y,2}, g_2, \epsilon_2)\|_{2, P_{Q_0, \mathbf{g}^r}} \\ &\leq \|\text{logit}(Q_{Y,1}) - \text{logit}(Q_{Y,2})\|_{2, P_{Q_0, \mathbf{g}^r}} + \|\epsilon_2(1/g_1 - 1/g_2)\|_{2, Q_{W,0}} + \|(\epsilon_1 - \epsilon_2)/g_1\|_{2, Q_{W,0}} \\ &\lesssim \|Q_{Y,1} - Q_{Y,2}\|_{2, P_{Q_0, \mathbf{g}^r}} + \|g_1 - g_2\|_{2, Q_{W,0}} + |\epsilon_1 - \epsilon_2| = \|(Q_{Y,1}, g_1, \epsilon_1) - (Q_{Y,2}, g_2, \epsilon_2)\| \end{aligned}$$

(f is Lipschitz). Therefore, the convergence $\|(Q_{Y,\beta_n}, g_n, \epsilon_n) - (Q_{Y,\beta_0}, g_0^*, \epsilon_0)\| = o_P(1)$ and equalities $Q_{Y,\beta_n}^* = f(Q_{Y,\beta_n}, g_n, \epsilon_n)$, $Q_{Y,\beta_0}^* = f(Q_{Y,\beta_0}, g_0^*, \epsilon_0)$, entail $\|Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*\|_{2,P_{Q_0,g^r}} = o_P(1)$, hence our first claim. This completes the second part of the proof.

The second claim follows from the double-robustness of the efficient influence curve D . Indeed, $\mathbf{M}(\epsilon_0) = P_{Q_0,g_0^*} D_Y(Q_{Y,\beta_0}^*, g_0^*) = 0$ from the first part of this proof, and $P_{Q_0,g_0^*} D_W(Q_{\beta_0}^*) = 0$ from the definitions of Ψ and D_W , hence $P_{Q_0,g_0^*} D(P_{Q_{\beta_0}^*,g_0^*}) = 0$. Thus, Lemma 1 guarantees that $\Psi(Q_{\beta_0}^*) = \Psi(Q_0)$ since P_{Q_0,g_0^*} and $P_{Q_{\beta_0}^*,g_0^*}$ share the same g_0^* . We now turn to the third and last claim. For both $\beta = \beta_0$ and $\beta = \beta_n$, introduce $q_{Y,\beta}^*$ characterized by

$$q_{Y,\beta}^*(W) \equiv Q_{Y,\beta}^*(1, W) - Q_{Y,\beta}^*(0, W). \quad (29)$$

Define also $Q_{\beta_n}^{\sim} \equiv (Q_{W,0}, Q_{Y,\beta_n}^*)$ and $\psi_n^{\sim} \equiv \Psi(Q_{\beta_n}^{\sim})$. Since g^r is bounded away from 0, the Cauchy-Schwarz inequality yields

$$\begin{aligned} |\psi_n^{\sim} - \psi_0| &= |\psi_n^{\sim} - \Psi(Q_{\beta_0}^*)| = |P_{Q_0,g^r}(Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*)(2A - 1)/g^r| \\ &\lesssim P_{Q_0,g^r} |Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*| \leq \|Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*\|_{2,P_{Q_0,g^r}} = o_P(1). \end{aligned}$$

Furthermore, $\psi_n^* - \psi_n^{\sim} = (P_n - P_{Q_0,\mathbf{g}_n})(q_{Y,\beta_n}^* - q_{Y,\beta_0}^*) + (P_n - P_{Q_0,\mathbf{g}_n})q_{Y,\beta_0}^*$. By similar arguments as before, we establish that $(P_n - P_{Q_0,\mathbf{g}_n})(q_{Y,\beta_n}^* - q_{Y,\beta_0}^*) = o_P(1)$. In addition, the law of large numbers (for independent, identically distributed random variables, since q_{Y,β_0}^* is a bounded function of W only) guarantees that $(P_n - P_{Q_0,\mathbf{g}_n})q_{Y,\beta_0}^* = o_P(1)$. In summary, $\psi_n^* - \psi_0 = (\psi_n^* - \psi_n^{\sim}) + (\psi_n^{\sim} - \psi_0) = o_P(1)$, as stated. This completes the proof. \square

The asymptotic linear expansion (21) in Proposition 4 is a by-product of the exact linear expansion that we state and prove below. Recall the definitions of $d_{Y,\beta}^*$ and $q_{Y,\beta}^*$ ($\beta = \beta_0$ or $\beta = \beta_n$) given in (18) and (29).

Lemma 2 (exact linear expansion of ψ_n^*). *It follows from the definition of ψ_n^* that*

$$\psi_n^* - \psi_0 = -P_{Q_0,g_0^*} D(P_{Q_{\beta_n}^*,g_0^*}) \quad (30)$$

$$\begin{aligned} &= (P_n - P_{Q_0,\mathbf{g}_n})(d_{Y,\beta_0}^* + D_W(Q_{\beta_0}^*)) \\ &\quad + (P_n - P_{Q_0,\mathbf{g}_n})((d_{Y,\beta_n}^* - d_{Y,\beta_0}^*) + (q_{Y,\beta_n}^* - q_{Y,\beta_0}^*)). \end{aligned} \quad (31)$$

Proof of Lemma 2. Consider (30). By Lemma 1, the efficient influence curve decomposes as $D(P_{Q_{\beta_n}^*,g_0^*}) = D_Y(Q_{Y,\beta_n}^*, g_0^*) + D_W(Q_{\beta_n}^*)$. Define $q_{Y,0}(W) \equiv Q_{Y,0}(1, W) - Q_{Y,0}(0, W)$. Firstly, $P_{Q_0,g_0^*} D_W(Q_{\beta_n}^*) = P_{Q_0,g_0^*} q_{Y,\beta_n}^* - \psi_n^*$. Secondly, $P_{Q_0,g_0^*} D_Y(Q_{Y,\beta_n}^*, g_0^*) = P_{Q_0,g_0^*} (2A - 1)(Y - Q_{Y,\beta_n}^*)/g_0^* = P_{Q_0,g_0^*} (q_{Y,0} - q_{Y,\beta_n}^*)$. Adding these two equalities yields $P_{Q_0,g_0^*} D(P_{Q_{\beta_n}^*,g_0^*}) = P_{Q_0,g_0^*} q_{Y,0} - \psi_n^* = \psi_0 - \psi_n^*$, which is the desired result.

We now turn to (31). Denote P_{n,\mathbf{g}_n} the empirical distribution of \mathbf{O}_n weighted by $g_n(A_i|W_i)/g_i(A_i|W_i)$. By construction of the fluctuation (9) and definition of ϵ_n (10), it holds that $P_{n,\mathbf{g}_n} D_Y(Q_{Y,\beta_n}^*, g_n) = 0$. Moreover, (11) can be rewritten as $P_n D_W(Q_{\beta_n}^*) = 0$. Therefore, (30) is equivalent to

$$\psi_n^* - \psi_0 = (P_n - P_{Q_0,g_0^*}) D_W(Q_{\beta_n}^*) + (P_{n,\mathbf{g}_n} D_Y(Q_{Y,\beta_n}^*, g_n) - P_{Q_0,g_0^*} D_Y(Q_{Y,\beta_n}^*, g_0^*)). \quad (32)$$

Adding and subtracting $(P_n - P_{Q_0,g_0^*}) D_W(Q_{\beta_0}^*)$ to the first term in the RHS expression of (32) implies

$$\begin{aligned} (P_n - P_{Q_0,g_0^*}) D_W(Q_{\beta_n}^*) &= (P_n - P_{Q_0,g_0^*}) D_W(Q_{\beta_0}^*) + (P_n - P_{Q_0,g_0^*}) (D_W(Q_{\beta_n}^*) - D_W(Q_{\beta_0}^*)) \\ &= (P_n - P_{Q_0,g_0^*}) D_W(Q_{\beta_0}^*) + (P_n - P_{Q_0,g_0^*}) (q_{Y,\beta_n}^* - q_{Y,\beta_0}^*) \\ &= (P_n - P_{Q_0,\mathbf{g}_n}) D_W(Q_{\beta_0}^*) + (P_n - P_{Q_0,\mathbf{g}_n}) (q_{Y,\beta_n}^* - q_{Y,\beta_0}^*), \end{aligned} \quad (33)$$

where the last equality is valid because $D_W(Q_{\beta_0}^*)$, q_{Y,β_n}^* , q_{Y,β_0}^* are functions of W only. As for the second

term in the RHS expression of (32), it equals

$$\begin{aligned}
& \frac{1}{n} \sum_{i=1}^n \left(\frac{g_n(A_i|W_i)}{g_i(A_i|W_i)} \frac{2A_i - 1}{g_n(A_i|W_i)} (Y_i - Q_{Y,\beta_n}^*(A_i, W_i)) - P_{Q_0, g_0^*} \frac{2A - 1}{g_0^*(A|W)} (Y - Q_{Y,\beta_n}^*) \right) \\
&= \frac{1}{n} \sum_{i=1}^n \left(\frac{2A_i - 1}{g_i(A_i|W_i)} (Y_i - Q_{Y,\beta_n}^*(A_i, W_i)) - P_{Q_0, g_i} \frac{2A - 1}{g_i(A|W)} (Y - Q_{Y,\beta_n}^*) \right) \\
&= (P_n - P_{Q_0, \mathbf{g}_n}) d_{Y,\beta_n}^* \\
&= (P_n - P_{Q_0, \mathbf{g}_n}) d_{Y,\beta_0}^* + (P_n - P_{Q_0, \mathbf{g}_n}) (d_{Y,\beta_n}^* - d_{Y,\beta_0}^*). \tag{34}
\end{aligned}$$

The equalities (32), (33) and (34) imply (31). \square

It appears that the second term in the RHS expression of (31) is asymptotically negligible at rate \sqrt{n} . Indeed,

Lemma 3. *It holds that $(P_n - P_{Q_0, \mathbf{g}_n}) \left((d_{Y,\beta_n}^* - d_{Y,\beta_0}^*) + (q_{Y,\beta_n}^* - q_{Y,\beta_0}^*) \right) = o_P(1/\sqrt{n})$.*

Proof of Lemma 3. The key to this proof is Lemma 10.

Introduce $\mathcal{Q}_{1,n}^* \equiv \{f(Q_{Y,\beta}, g, \epsilon) : Q_{Y,\beta} \in \mathcal{Q}_{1,n}, g \in \mathcal{G}_{1,n}, \epsilon \in \mathcal{E}\}$, where f is given by (28), and set $\delta > 0$. The elements of $\mathcal{Q}_{1,n}^*$ are uniformly bounded away from 0 and 1. By **A5**, **A6** and Lemma 11, the bracketing numbers $N(\delta, \text{logit}(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}})$ and $N(\delta, 1/\mathcal{Q}_{1,n}, \|\cdot\|_{2, P_{Q_0, g^r}})$ are finite. Obviously, the bracketing number $N(\delta, \mathcal{E}, |\cdot|)$ is finite too. Choose arbitrarily three collections of δ -brackets of smallest possible cardinality that cover $\text{logit}(\mathcal{Q}_{1,n})$, $1/\mathcal{G}_{1,n}$, and \mathcal{E} . Given $f(Q_{Y,\beta}, g, \epsilon) \in \mathcal{Q}_{1,n}^*$, let $[l_Q, u_Q]$, $[l_g, u_g]$ and $[l_\epsilon, u_\epsilon]$ be δ -brackets from these collections and containing $\text{logit}(Q_{Y,\beta})$, $1/g$ and ϵ , respectively. We can assume without loss of generality that the uniform lower- and upper-bounds of $\text{logit}(\mathcal{Q}_{1,n})$ (respectively, $1/\mathcal{G}_{1,n}$) are also lower- and upper-bounds on l_Q , u_Q , (respectively, l_g , u_g). We can also assume that $|l_\epsilon|, |u_\epsilon| \leq \sup_{\epsilon \in \mathcal{E}} |\epsilon|$. Characterize λ and γ by setting $\lambda(O) \equiv A l_\epsilon H(u_g)(O) + (1 - A) u_\epsilon H(l_g)(O)$ and, similarly, $\gamma(O) \equiv A u_\epsilon H(l_g)(O) + (1 - A) l_\epsilon H(u_g)(O)$. Then $[\text{expit}(l_Q + \lambda), \text{expit}(u_Q + \gamma)]$ is a bracket containing $f(Q_{Y,\beta}, g, \epsilon)$. Since expit is 1-Lipschitz, it follows that $(\text{expit}(u_Q + \gamma) - \text{expit}(l_Q + \lambda))^2 \leq ((u_Q - l_Q) + (\gamma - \lambda))^2 \leq 2(u_Q - l_Q)^2 + 2(\gamma - \lambda)^2$ where $(\gamma - \lambda)^2 \lesssim (u_\epsilon - l_\epsilon)^2 + (H(u_g) - H(l_g))^2 \lesssim (u_\epsilon - l_\epsilon)^2 + (u_g - l_g)^2$. Consequently, there exists a universal constant $c \geq 1$ such that $[\text{expit}(l_Q + \lambda), \text{expit}(u_Q + \gamma)]$ be a $c\delta$ -bracket. Thus, $N(\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta/c, \text{logit}(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}}) \times N(\delta/c, 1/\mathcal{G}_{1,n}, \|\cdot\|_{2, Q_{W,0}}) \times N(\delta/c, \mathcal{E}, |\cdot|)$ hence, by Lemma 11, $J(\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}}) \lesssim J(\delta, \mathcal{Q}_{1,n}, \|\cdot\|_{2, P_{Q_0, g^r}}) + J(\delta, \mathcal{G}_{1,n}, \|\cdot\|_{2, Q_{W,0}}) + J(\delta, \mathcal{E}, |\cdot|)$. Therefore, **A5*** and **A6*** imply that if $\delta_n = o(1)$ then $J(\delta_n, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$ as well. Now, we use this to prove the lemma.

For each $Q_Y \in \mathcal{Q}_{1,n}^*$, characterize $d_Y(Q_Y)$ by setting $d_Y(Q_Y)(O, Z) \equiv (2A - 1)(Y - Q_Y(A, W))/Z$. By uniform boundedness of $\cup_{n \geq 1} \mathcal{Q}_{1,n}^*$, Y and Z , the existence of a sequence of envelope functions satisfying (a) in Lemma 10 is granted. Moreover, Lemma 11 yields that there exists $c > 0$ such that $J(\delta, d_Y(\mathcal{Q}_{1,n}^*), \|\cdot\|_{2, P_{Q_0, g^r}}) \leq cJ(\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}})$ for all $\delta > 0$. Thus, $\delta_n = o(1)$ implies $J(\delta_n, d_Y(\mathcal{Q}_{1,n}^*), \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$, and condition (b) in Lemma 10 is met too. Now, the convergence $\|Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*\|_{2, P_{Q_0, g^r}} = o_P(1)$, established in Proposition 3, implies $P_{Q_0, g^r}(d_Y(Q_{Y,\beta_n}^*) - d_Y(Q_{Y,\beta_0}^*))^2 = o_P(1)$ by Cauchy-Schwarz, since $|d_Y(Q_{Y,\beta_n}^*) - d_Y(Q_{Y,\beta_0}^*)| \lesssim |Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*|$. We apply Lemma 10 to obtain $\sqrt{n}(P_n - P_{Q_0, \mathbf{g}_n})(d_{Y,\beta_n}^* - d_{Y,\beta_0}^*) = o_P(1)$.

Now, for each $Q_Y \in \mathcal{Q}_{1,n}^*$, characterize $q_Y(Q_Y)$ by setting $q_Y(Q_Y)(W) \equiv Q_Y(1, W) - Q_Y(0, W)$. Choose a collection of $N(\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}})$ δ -brackets $[l_k, u_k]$ covering $\mathcal{Q}_{1,n}^*$ and set arbitrarily $Q_Y \in \mathcal{Q}_{1,n}^*$. Assume without loss of generality that $Q_Y \in [l_1, u_1]$ and characterize l'_1 and u'_1 by setting $l'_1(W) \equiv l_1(1, W) - u_1(0, W)$ and $u'_1(W) \equiv u_1(1, W) - l_1(0, W)$. It holds that $q_Y(Q_Y) \in [l'_1, u'_1]$ and $P_{Q_0, g^r}(u'_1 - l'_1)^2 \leq 2\delta^2/c$ for $0 < c \equiv \min(\inf g^r, 1 - \sup g^r) < 1$. Thus, $N(\delta, q_Y(\mathcal{Q}_{1,n}^*), \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\sqrt{2/c}\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}})$, hence $J(\delta_n, q_Y(\mathcal{Q}_{1,n}^*), \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$ whenever $\delta_n = o(1)$: condition (b) in Lemma 10 is met. Condition (a) in the same lemma is also met since $\cup_{n \geq 1} q_Y(\mathcal{Q}_{1,n}^*)$ is uniformly bounded. Moreover, $\|Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*\|_{2, P_{Q_0, g^r}} = o_P(1)$ implies $\|q_Y(Q_{Y,\beta_n}^*) - q_Y(Q_{Y,\beta_0}^*)\|_{2, P_{Q_0, g^r}} = o_P(1)$ since $P_{Q_0, g^r}(q_Y(Q_{Y,\beta_n}^*) - q_Y(Q_{Y,\beta_0}^*))^2 \leq 2P_{Q_0, g^r}(Q_{Y,\beta_n}^*(1, W) - Q_{Y,\beta_0}^*(1, W))^2 + 2P_{Q_0, g^r}(Q_{Y,\beta_n}^*(1, W) - Q_{Y,\beta_0}^*(1, W))^2$ and, for both

$a = 0, 1$, $P_{Q_0, g^r}(Q_{Y, \beta_n}^*(a, W) - Q_{Y, \beta_0}^*(a, W))^2 = P_{Q_0, g^r}(Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*)^2 \mathbf{1}\{A = a\}/g^r(a|W) \lesssim P_{Q_0, g^r}(Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*)^2$ because g^r is bounded away from 0 and 1. We apply Lemma 10 to obtain $\sqrt{n}(P_n - P_{Q_0, \mathbf{g}_n})(q_{Y, \beta_n}^* - q_{Y, \beta_0}^*) = o_P(1)$.

This completes the proof. \square

The proof of Proposition 4 is now at hand.

Proof of proposition 4. We first note that (21) follows straightforwardly from Lemmas 2 and 3.

Set $f_0 \equiv d_{Y, \beta_0}^* + D_W(Q_{\beta_0}^*)$ and $f_n \equiv d_{Y, \beta_n}^* + D_W(Q_{\beta_n}^*)$. With this notation, $\Sigma_0 = P_{Q_0, g_0^*} f_0^2$, $\Sigma_n = P_n f_n^2$. Introduce also $S_n \equiv P_{Q_0, \mathbf{g}_n} f_0^2$. For either $(f, \beta) = (f_0, \beta_0)$ or $(f, \beta) = (f_n, \beta_n)$, it holds that

$$\begin{aligned} P_{Q_0, \mathbf{g}_n} f^2 &= \frac{1}{n} \sum_{i=1}^n P_{Q_0, g_i} f^2 \\ &= P_{Q_0, g_0^*} (D_W(Q_{\beta}^*)^2 + 2D_Y(Q_{Y, \beta}^*, g_0^*)D_W(Q_{\beta}^*)) + \frac{1}{n} \sum_{i=1}^n P_{Q_0, g_0^*} \frac{(Y - Q_{Y, \beta}^*)^2}{g_0^* g_i} \\ &= P_{Q_0, g_0^*} (D_W(Q_{\beta}^*)^2 + 2D_Y(Q_{Y, \beta}^*, g_0^*)D_W(Q_{\beta}^*)) + P_{Q_0, g_0^*} \frac{(Y - Q_{Y, \beta}^*)^2}{g_0^*} \frac{1}{n} \sum_{i=1}^n 1/g_i. \end{aligned}$$

Now, because $(Y - Q_{Y, \beta}^*)^2 \leq 1$ and g_0^* is bounded away from 0 and 1, the Cauchy-Schwarz inequality yields

$$\begin{aligned} |P_{Q_0, \mathbf{g}_n} f^2 - P_{Q_0, g_0^*} f^2| &= \left| P_{Q_0, g_0^*} \frac{(Y - Q_{Y, \beta}^*)^2}{g_0^*} \left(\frac{1}{n} \sum_{i=1}^n 1/g_i - 1/g_0^* \right) \right| \\ &\lesssim P_{Q_0, g_0^*} \left| \frac{1}{n} \sum_{i=1}^n 1/g_i - 1/g_0^* \right| \leq \left\| \frac{1}{n} \sum_{i=1}^n 1/g_i - 1/g_0^* \right\|_{2, Q_{W, 0}}. \quad (35) \end{aligned}$$

Thus, taking $f = f_0$ and applying corollary 1, we obtain $E(S_n) = \Sigma_0 + o(1)$ and $S_n = \Sigma_0 + o_P(1)$ (Note that $\Sigma_0 > 0$ by **A1**). Let us show now that $\Sigma_n = \Sigma_0 + o_P(1)$ by proving $\Sigma_n - S_n = o_P(1)$. We use the following decomposition:

$$\begin{aligned} \Sigma_n - S_n &= (P_n - P_{Q_0, \mathbf{g}_n})(f_n^2 - f_0^2) + (P_n - P_{Q_0, \mathbf{g}_n})f_0^2 + P_{Q_0, \mathbf{g}_n}(f_n^2 - f_0^2) \\ &= (P_n - P_{Q_0, \mathbf{g}_n})(f_n^2 - f_0^2) + (P_n - P_{Q_0, \mathbf{g}_n})f_0^2 + P_{Q_0, g_0^*}(f_n^2 - f_0^2) + o_P(1), \quad (36) \end{aligned}$$

where the second equality holds because $P_{Q_0, \mathbf{g}_n} f^2 = P_{Q_0, g_0^*} f^2 + o_P(1)$ for both $f = f_0$ and $f = f_n$ (by (35) and Corollary 1). Because f_0 and all f_n 's ($n \geq 1$) are uniformly bounded, the first term in the RHS expression of (36) satisfies

$$\begin{aligned} |(P_n - P_{Q_0, \mathbf{g}_n})(f_n^2 - f_0^2)| &\lesssim |(P_n - P_{Q_0, \mathbf{g}_n})(f_n - f_0)| \\ &= |(P_n - P_{Q_0, \mathbf{g}_n})(d_{Y, \beta_n}^* - d_{Y, \beta_0}^*) + (q_{Y, \beta_n}^* - q_{Y, \beta_0}^*)| = o_P(1/\sqrt{n}) \end{aligned}$$

by Lemma 3 (see (29) for the definition of $q_{Y, \beta}^*$). Since f_0 is bounded, the Kolmogorov strong law of large numbers (Sen and Singer, 1993, Theorem 2.4.2) guarantees that the second term in the RHS expression of (36) converges to 0 P -almost-surely, hence $(P_n - P_{Q_0, \mathbf{g}_n})f_0^2 = o_P(1)$. Consider now the third term in the RHS expression of (36). Note that $(f_n - f_0)(O, Z) = (2A - 1)(Q_{Y, \beta_0}^* - Q_{Y, \beta_n}^*)(A, W)/Z + (q_{Y, \beta_n}^* - q_{Y, \beta_0}^*)(W) - (\psi_n^* - \psi_0)$, hence $|f_n - f_0| \lesssim |Q_{Y, \beta_0}^* - Q_{Y, \beta_n}^*| + |q_{Y, \beta_n}^* - q_{Y, \beta_0}^*| + |\psi_n^* - \psi_0|$ because Z is bounded away from 0 and 1. Using again (i) that f_0 and all f_n 's ($n \geq 1$) are uniformly bounded, and (ii) the Cauchy-Schwarz inequality and the dominated ratio property, we get

$$\begin{aligned} |P_{Q_0, g_0^*}(f_n^2 - f_0^2)| &\lesssim |P_{Q_0, g_0^*}(f_n - f_0)| \\ &\lesssim P_{Q_0, g_0^*} |Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*| + P_{Q_0, g_0^*} |q_{Y, \beta_n}^* - q_{Y, \beta_0}^*| + |\psi_n^* - \psi_0| \\ &\lesssim \|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} + \|q_{Y, \beta_n}^* - q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} + |\psi_n^* - \psi_0|. \end{aligned}$$

We know that $\|Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*\|_{2,P_{Q_0,g^r}} = o_P(1)$ by Proposition 1, we showed at the end of the proof of Lemma 3 that this implies $\|q_{Y,\beta_n}^* - q_{Y,\beta_0}^*\|_{2,P_{Q_0,g^r}} = o_P(1)$, and Proposition 3 guarantees that $\psi_n^* - \psi_0 = o_P(1)$. Consequently, $|P_{Q_0,g_0^*}(f_n^2 - f_0^2)| = o_P(1)$. We have thus proven that all terms in the RHS expression of (36) are $o_P(1)$, hence $\Sigma_n - \Sigma_n = o_P(1)$ and $\Sigma_n = \Sigma_0 + o_P(1)$, as we claimed earlier.

We show now that (21), which we rewrite here $\psi_n^* - \psi_0 = (P_n - P_{Q_0,g_n})f_0 + o_P(1/\sqrt{n})$, implies that $\sqrt{n}/\Sigma_0(\psi_n^* - \psi_0)$ converges in law to the standard normal distribution. This is a consequence of (Sen and Singer, 1993, Theorem 3.3.7) because (i) $S_n/E(S_n) - 1 = o_P(1)$, and (ii) for each $\alpha > 0$, $E(P_n f_0^2 \mathbf{1}\{f_0^2 \geq \alpha^2 n E(S_n)\}) = o(E(S_n))$ trivially holds since f_0 is bounded and $E(S_n) = \Sigma_0 + o(1)$ with $\Sigma_0 > 0$. Then Slutsky's lemma and $\Sigma_n = \Sigma_0 + o_P(1)$ yield the convergence in law of $\sqrt{n}/\Sigma_n(\psi_n^* - \psi_0)$ to the same limiting distribution. This completes the proof. \square

The proof of Corollary 2 boils down to (i) showing that **A5**, **A5***, **A6**, **A6*** are met and (ii) applying Theorem 1.

Proof of Corollary 2. We show below that **A5** and **A5*** are met. A parallel argument can be used to show that **A6** and **A6*** hold too. Since **A1**–**A4** are satisfied by assumption, Theorem 1 thus applies and yields the stated result.

Fix $\delta > 0$, a sequence $\{\delta_n\}_{n \geq 1}$ of positive numbers such that $\delta_n = o(1)$, and $n \geq 1$. By construction, the functions ϕ_j ($j \in \mathbb{N}$) all belong to a class \mathcal{C} of smooth functions over the bounded support \mathcal{W} such that all partial derivatives up to order $\alpha > \dim(\mathcal{W})/2$ of all $f \in \mathcal{C}$ exist and are uniformly bounded by a constant $C > 0$. By (van der Vaart, 1998, Example 19.9), it holds that $\log N(\delta, \mathcal{C}, \|\cdot\|_{2,P_{Q_0,g^r}}) \lesssim \delta^{-V}$ for $V \equiv \dim(\mathcal{W})/\alpha < 2$.

Note that $\mathcal{F} \equiv \{\sum_{j \in \mathbb{N}} \beta_j \phi_j = \sum_{j=0}^{d_n} \beta_j \phi_j : \beta \in B_n\}$ is a subset of \mathcal{C} , provided that the constant C in the definition of \mathcal{C} is large enough (if not, it suffices to replace C with MC , with M the constant involved in (22)). We apply three times Lemma 11 to obtain that $J(\delta, \mathcal{F}, \|\cdot\|_{2,P_{Q_0,g^r}}) \gtrsim J(\delta, \text{logit}(\mathcal{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}}) \gtrsim J(\delta, \mathcal{Q}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}) \gtrsim J(\delta, L^{\text{kl}}(\mathcal{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}})$: from left to right, the inequalities follow from (i) the third claim of Lemma 11 with h, h' given by $h(O) \equiv A$ and $h'(O) \equiv (1 - A)$, (ii) from the sixth claim with $\phi \equiv \text{expit}$, which is increasing and 1-Lipschitz, and (iii) from the seventh claim with h given by $h(O) \equiv Y$. Therefore,

$$\begin{aligned} J(\delta_n, L^{\text{kl}}(\mathcal{Q}_{1,n}), \|\cdot\|_{2,P_{Q_0,g^r}}) &\lesssim J(\delta_n, \mathcal{Q}_{1,n}, \|\cdot\|_{2,P_{Q_0,g^r}}) \\ &\lesssim J(\delta_n, \mathcal{C}, \|\cdot\|_{2,P_{Q_0,g^r}}) \lesssim \int_0^{\delta_n} \varepsilon^{-V/2} d\varepsilon = o(1), \end{aligned}$$

and **A5*** is fulfilled. Choosing $\delta_n = 1/\sqrt{n}$ yields that **A5** is also fulfilled. This completes the proof. \square

A.2 Useful technical results

Convergence of M -estimators.

The following lemma is a simple adaptation of (van der Vaart and Wellner, 1996, Corollary 3.2.3).

Lemma 4 (convergence of M -estimators). *Let \mathbf{M}_n be a real-valued, stochastic processes indexed by a metric space (Θ, d) , and let $\mathbf{M} : \Theta \rightarrow \mathbb{R}$ be a real-valued, deterministic function over Θ . Consider a sequence of subsets $\Theta_n \subset \Theta$ and the following assumptions:*

- (a) *There exists $\theta_0 \in \Theta$ such that $\mathbf{M}(\theta_0) < \inf_{\theta \notin T} \mathbf{M}(\theta)$ for every open set $T \subset \Theta$ containing θ_0 .*
- (b) *For each $n \geq 1$, there exists $\theta_n^* \in \Theta_n$ such that $\mathbf{M}(\theta_n^*) = \inf_{\theta \in \Theta_n} \mathbf{M}(\theta)$. Moreover, $\mathbf{M}(\theta_n^*) - \mathbf{M}(\theta_0) = o(1)$.*
- (c) *It holds that $\|\mathbf{M}_n - \mathbf{M}\|_{\Theta_n} = o_P(1)$.*

Under the above three assumptions, if $\theta_n \in \Theta_n$ satisfies $\mathbf{M}_n(\theta_n) - \mathbf{M}_n(\theta_n^*) \leq 0$ for all $n \geq 1$, then $d(\theta_n, \theta_0) = o_P(1)$.

Proof of Lemma 4. Set $n \geq 1$. By **(a)**, it holds that

$$\begin{aligned} 0 &\leq \mathbf{M}(\theta_n) - \mathbf{M}(\theta_0) \\ &= (\mathbf{M}(\theta_n) - \mathbf{M}_n(\theta_n)) + (\mathbf{M}_n(\theta_n) - \mathbf{M}_n(\theta_n^*)) + (\mathbf{M}_n(\theta_n^*) - \mathbf{M}(\theta_n^*)) + (\mathbf{M}(\theta_n^*) - \mathbf{M}(\theta_0)). \end{aligned}$$

The above first and third RHS terms are both upper-bounded by $\|\mathbf{M}_n - \mathbf{M}\|_{\Theta_n}$. The second RHS term is non-positive by definition of θ_n . The fourth RHS terms is $o(1)$ by **(b)**. Thus, it actually holds that $0 \leq \mathbf{M}(\theta_n) - \mathbf{M}(\theta_0) \leq 2\|\mathbf{M}_n - \mathbf{M}\|_{\Theta_n} + o(1) = o_P(1)$ by **(c)**.

Set $\varepsilon > 0$. By **(a)**, there exists $\delta > 0$ such that $d(\theta_n, \theta_0) \geq \varepsilon$ implies $\mathbf{M}(\theta_n) - \mathbf{M}(\theta_0) \geq \delta$. Since we have shown that $\mathbf{M}(\theta_n) - \mathbf{M}(\theta_0) = o_P(1)$, we can therefore conclude that $d(\theta_n, \theta_0) = o_P(1)$ too. \square

Maximal inequalities and convergence of empirical processes.

In this article, we repeatedly exploit uniform laws of large numbers. They are derived from maximal inequalities for martingales by van Handel (2011) that also played an important role in (Chambaz and van der Laan, 2011a,c). For completeness, we now state these results.

Let $\phi : \mathbb{R} \rightarrow \mathbb{R}_+$ be such that $\phi(x) = e^x - x - 1$. Let \mathcal{F} be a class of measurable functions, $n \geq 1$ be an integer, $K > 0$ and $\delta > 0$ be two positive constants. For each $f \in \mathcal{F}$, $n(P_n - P_{Q_0, \mathbf{g}_n})f = \sum_{i=1}^n (f(O_i, Z_i) - P_{Q_0, g_i} f)$ is a discrete martingale sum.

Set $N = N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, the δ -bracketing number of \mathcal{F} wrt $\|\cdot\|_{2, P_{Q_0, g^r}}$. Following van Handel (2011), we define a $(\delta, n, \mathcal{F}, K)$ -bracketing set as a collection $\{(\Lambda_i^j, \Gamma_i^j) : i \leq n\}_{j \leq N}$ of random variables such that (i) for each $f \in \mathcal{F}$, there exists $j \leq N$ satisfying $\Lambda_i^j \leq f(O_i, Z_i) \leq \Gamma_i^j$ for all $i \leq n$, and (ii) for all $j \leq N$, $2K^2 n^{-1} \sum_{i=1}^n P_{Q_0, g_i} \phi(|\Lambda_i^j - \Gamma_i^j|/K) \leq \delta^2$. Let $\mathcal{N}(\delta, n, \mathcal{F}, K)$ denote the cardinality of the smallest $(\delta, n, \mathcal{F}, K)$ -bracketing set. Finally, introduce for each $f \in \mathcal{F}$ the random variable $R_{n, K}(f) = 2K^2 n^{-1} \sum_{i=1}^n P_{Q_0, g_i} \phi(|f|/K)$.

Lemma 5 (Proposition A.2 by van Handel (2011)). *There exists an universal constant $C > 0$ such that, for all $R > 0$,*

$$P \left(\sup_{f \in \mathcal{F}} \mathbf{1} \{R_{n, K}(f) \leq R\} \max_{i \leq n} \frac{i}{n} (P_i - P_{Q_0, \mathbf{g}_i}) f \geq \alpha \right) \leq 2 \exp \left(-\frac{n\alpha^2}{C^2(c_1 + 1)R} \right),$$

for any $\alpha, c_0, c_1 > 0$ such that $c_0^2 \geq C^2(c_1 + 1)$ and

$$\frac{c_0}{\sqrt{n}} \int_0^{\sqrt{R}} \sqrt{\log \mathcal{N}(\varepsilon, n, \mathcal{F}, K)} d\varepsilon \leq \alpha \leq \frac{c_1 R}{K}.$$

Lemma 6 (Corollary A.8 by van Handel (2011)). *Suppose the class \mathcal{F} is finite. For all $R > 0$ and any event C ,*

$$E \left(\max_{f \in \mathcal{F}} \mathbf{1} \{nR_{n, K}(f) \leq R\} \max_{i \leq n} i (P_i - P_{Q_0, \mathbf{g}_i}) f \right) \leq \sqrt{2R \log \left(1 + \frac{|\mathcal{F}|}{P(C)} \right)} + 8K \log \left(1 + \frac{|\mathcal{F}|}{P(C)} \right).$$

If, in addition, $\max_{f \in \mathcal{F}} \|f\|_\infty \leq U$, then K can be replaced with $U/3$ in the second term of the above RHS expression.

Importantly, van Handel (2011)'s proofs of Lemmas 5 and 6 remain valid when the class \mathcal{F} is allowed to depend on n . To use lemmas 5 and 6, it is necessary to get a grip on $\mathcal{N}(\delta, n, \mathcal{F}, K)$ and the random variables $R_{n, K}(f)$, $f \in \mathcal{F}$. The next lemma is helpful in this regard.

Recall that, by the dominated ratio property of \mathcal{G}_1 , $\|g/g^r\|_\infty \leq \kappa$ for all $g \in \mathcal{G}_1$.

Lemma 7 (L^2 -norm version of lemma 7 by Chambaz and van der Laan (2011c)). Assume that $U \equiv \sup_{f \in \mathcal{F}} \|f\|_\infty$ is finite. Then, for all $f \in \mathcal{F}$, $R_{n,4U}(f) \leq 4/3n \sum_{i=1}^n P_{Q_0, g_i} |f|^2$. Moreover, it holds that $\mathcal{N}(\sqrt{2\kappa}\delta, n, \mathcal{F}, 4U) \leq N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$.

Proof of Lemma 7. Set $f \in \mathcal{F}$, $i \leq n$, and $m \geq 2$. It holds that $P_{Q_0, g_i} |f|^m \leq U^{m-2} P_{Q_0, g_i} |f|^2 \leq \frac{m!}{2} U^{m-2} P_{Q_0, g_i} |f|^2$. Therefore, for $K = 4U$,

$$\begin{aligned} 2K^2 P_{Q_0, g_i} \phi(|f|/K) &= 2(4U)^2 \sum_{m \geq 2} \frac{P_{Q_0, g_i} |f|^m}{m! (4U)^m} \\ &\leq 2(4U)^2 \sum_{m \geq 2} \frac{\frac{m!}{2} U^{m-2} P_{Q_0, g_i} |f|^2}{m! (4U)^m} = 16 \sum_{m \geq 2} \frac{P_{Q_0, g_i} |f|^2}{4^m} = 4P_{Q_0, g_i} |f|^2/3. \end{aligned}$$

The monotone convergence theorem guarantees the first equality. Summing up the above inequalities for $i = 1, \dots, n$ yields the first bound.

Let $(\ell^j, u^j)_{j \leq N}$ be a set of $N = N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$ δ -brackets covering \mathcal{F} wrt $\|\cdot\|_{2, P_{Q_0, g^r}}$. Let $\Lambda_i^j = \max(\ell^j(O_i, Z_i), -U)$ and $\Gamma_i^j = \min(u^j(O_i, Z_i), U)$ for all $i \leq n, j \leq N$. Set $f \in \mathcal{F}$ and $j \leq N$ such that $f \in [\ell^j, u^j]$. Then, for all $i \leq n$, (i) $\Lambda_i^j \leq f(O_i, Z_i) \leq \Gamma_i^j$, (ii) $-U \leq \Lambda_i^j \leq \Gamma_i^j \leq U$, and (iii) $\ell^j \leq \Lambda_i^j \leq \Gamma_i^j \leq u^j$. Thus, for all $m \geq 2$,

$$\begin{aligned} P_{Q_0, g_i} |\Lambda_i^j - \Gamma_i^j|^m &\leq (2U)^{m-2} P_{Q_0, g_i} |\Lambda_i^j - \Gamma_i^j|^2 \leq (2U)^{m-2} \kappa P_{Q_0, g^r} |\Lambda_i^j - \Gamma_i^j|^2 \\ &\leq (2U)^{m-2} \kappa P_{Q_0, g^r} |\ell^j - u^j|^2 \leq (2U)^{m-2} \kappa \delta^2 \leq \frac{m!}{2} (2U)^{m-2} \kappa \delta^2. \end{aligned}$$

Consequently, still using $K = 4U$, it holds that

$$2K^2 P_{Q_0, g_i} \phi(|\Lambda_i^j - \Gamma_i^j|/4U) = 2(4U)^2 \sum_{m \geq 2} \frac{P_{Q_0, g_i} |\Lambda_i^j - \Gamma_i^j|^m}{m! (4U)^m} \leq 32U^2 \sum_{m \geq 2} \frac{\frac{m!}{2} (2U)^{m-2} \kappa \delta^2}{m! (4U)^m} = 2\kappa \delta^2.$$

Again, the monotone convergence theorem validates the first equality. Summing up the above inequalities for $i = 1, \dots, n$ yields $2K^2 n^{-1} \sum_{i=1}^n P_{Q_0, g_i} \phi(|\Lambda_i^j - \Gamma_i^j|/K) \leq 2\kappa \delta^2$, hence $\{(\Lambda_i^j, \Gamma_i^j) : i \leq n\}_{j \leq N}$ is a $(\sqrt{2\kappa}\delta, n, \mathcal{F}, K)$ -bracketing set and $\mathcal{N}(\sqrt{2\kappa}\delta, n, \mathcal{F}, 4U) \leq N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$. This completes the proof. \square

By combining Lemmas 5 and 7, we now establish a uniform law of large numbers.

Lemma 8. Let $\{\mathcal{F}_n\}_{n \geq 1}$ be a sequence of sets of measurable functions such that $U \equiv \sup_{f \in \cup_{n \geq 1} \mathcal{F}_n} \|f\|_\infty$ be finite. If $J(\sqrt{2/3\kappa}U, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(\sqrt{n})$, then for all $\alpha > 0$ there exists $c > 0$ and $n_0 \geq 1$ such that, for every $n \geq n_0$,

$$P \left(\sup_{f \in \mathcal{F}_n} (P_n - P_{Q_0, g_n}) f \geq \alpha \right) \leq 2e^{-nc}.$$

Consequently, $\sup_{f \in \mathcal{F}_n} |(P_n - P_{Q_0, g_n})f|$ converges to 0 P -almost surely.

Lemma 8 modifies Theorem 8 in Chambaz and van der Laan (2011c) to use an L^2 -metric and allow the classes of functions to change with n .

Proof of Lemma 8. Set $\alpha > 0$, and let $K = 4U$, $R = 4/3U^2$, $c_1 = \alpha K/R$, $c_0 = C\sqrt{c_1 + 1}$, where C is the universal constant from Lemma 5. Note that $\sqrt{R/2\kappa} = \sqrt{2/3\kappa}U$. By assumption, there exists $n_0 \geq 1$ such that, for all $n \geq n_0$, $J(\sqrt{R/2\kappa}, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \sqrt{n}\alpha/c_0\sqrt{2\kappa}$.

Set $n \geq n_0$. By Lemma 7, $\mathcal{N}(\sqrt{2\kappa}\delta, n, \mathcal{F}_n, 4U) \leq N(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})$. Therefore,

$$\begin{aligned} \frac{c_0}{\sqrt{n}} \int_0^{\sqrt{R}} \sqrt{\log \mathcal{N}(\varepsilon, n, \mathcal{F}_n, 4U)} d\varepsilon &\leq \frac{\sqrt{2\kappa}c_0}{\sqrt{n}} \int_0^{\sqrt{R/2\kappa}} \sqrt{\log N(\varepsilon, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon \\ &= \frac{\sqrt{2\kappa}c_0}{\sqrt{n}} J(\sqrt{R/2\kappa}, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \alpha = \frac{c_1 R}{K}. \end{aligned}$$

Lemma 5 applies and yields here

$$P \left(\sup_{f \in \mathcal{F}_n} (P_n - \mathbf{P}_{Q_0, \mathbf{g}_n}) f \geq \alpha \right) \leq P \left(\sup_{f \in \mathcal{F}_n} \max_{i \leq n} \frac{i}{n} (P_i - P_{Q_0, \mathbf{g}_i}) f \geq \alpha \right) \leq 2e^{-nc},$$

with $c = \alpha^2/c_0^2 R$. This completes the proof. \square

Lemma 5 also allows us to adapt the maximal inequality of (van der Vaart, 1998, Lemma 19.34), valid under independent, identically distributed sampling, to our targeted, adaptive sampling. We state and prove this result in lemma 9. We introduce the function Log given by $\text{Log}(x) \equiv \max(1, \log(x))$ (all $x > 0$).

Lemma 9. *Let \mathcal{F} be a class of measurable, real-valued functions and $\delta > 0$ be such that $P_{Q_0, g^r} f^2 \leq \delta^2$ for every $f \in \mathcal{F}$. Let F be an envelope function of \mathcal{F} . Define $a(\varepsilon) = \varepsilon / \sqrt{\text{Log } N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})}$ for all $\varepsilon > 0$. For each $n \geq 1$, it holds that*

$$\sqrt{n} E(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}}) \lesssim J(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}}) + \sqrt{n} E(P_{Q_0, \mathbf{g}_n} F \mathbf{1}\{F > \sqrt{n} a(\delta)\}) \quad (37)$$

$$\leq J(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}}) + \sqrt{n} \kappa P_{Q_0, g^r} F \mathbf{1}\{F > \sqrt{n} a(\delta)\}. \quad (38)$$

Proof of Lemma 9. The proof parallels that of (van der Vaart, 1998, Lemma 19.34).

Preliminary. Inequality (38) follows readily from (37) because $F \mathbf{1}\{F > \sqrt{n} a(\delta)\}$ is non-negative and \mathcal{G}_1 is endowed with the dominated ratio property. To understand the sum of two terms on the RHS of (37), first note that $E(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}})$ is upper-bounded by

$$E \left(\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n}) f \mathbf{1}\{F \leq \sqrt{n} a(\delta)\}| \right) + E \left(\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n}) f \mathbf{1}\{F > \sqrt{n} a(\delta)\}| \right). \quad (39)$$

Now, for every $f \in \mathcal{F}$,

$$|(P_n - P_{Q_0, \mathbf{g}_n}) f \mathbf{1}\{F > \sqrt{n} a(\delta)\}| \leq (P_n + P_{Q_0, \mathbf{g}_n}) F \mathbf{1}\{F > \sqrt{n} a(\delta)\},$$

hence, by the tower rule, the second term in (39) is smaller than $E((P_n + P_{Q_0, \mathbf{g}_n}) F \mathbf{1}\{F > \sqrt{n} a(\delta)\}) = 2E(P_{Q_0, \mathbf{g}_n} F \mathbf{1}\{F > \sqrt{n} a(\delta)\})$. Thus, to prove (37), it remains to show that \sqrt{n} times the first term in (39) is smaller than $J(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, up to a universal, multiplicative constant.

Before proceeding, note that $N(\varepsilon, \{f \mathbf{1}\{F \leq \sqrt{n} a(\delta)\} : f \in \mathcal{F}\}, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$ for all $\varepsilon > 0$, so that we may assume, without loss of generality, that $\sup_{f \in \mathcal{F}} \|f\|_{\infty} \leq \sqrt{n} a(\delta)$. What follows is based on a chaining technique to replace \mathcal{F} with a finite class.

Chaining. We now define a nested sequence of partitions on \mathcal{F} , then deduce a finite representation of \mathcal{F} from it. Fix q_0 such that $\delta \leq 2^{-q_0} \leq 2\delta$. For each integer $q \geq q_0$, denote $\tilde{N}_q \equiv N(2^{-q}, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$. Since $\varepsilon \mapsto N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$ is non-decreasing, it holds that

$$\sum_{q \geq q_0} 2^{-q} \sqrt{\text{Log } \tilde{N}_q} \lesssim \int_0^\delta \sqrt{\text{Log } N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon. \quad (40)$$

1. For each $q \geq q_0$, cover \mathcal{F} with \tilde{N}_q many brackets $[l_{q,i}, u_{q,i}]_{i \leq \tilde{N}_q}$ such that $P_{Q_0, g^r} \Delta_{q,i}^2 \leq 2^{-2q}$ for all $i \leq \tilde{N}_q$. Note that we may assume, without loss of generality, that $\Delta_{q,i} \equiv u_{q,i} - l_{q,i} \leq 2F \leq 2\sqrt{n} a(\delta)$ for all $i \leq \tilde{N}_q$. Define $\mathcal{F}_{q,1} \equiv [l_{q,1}, u_{q,1}]$ then, recursively, $\mathcal{F}_{q,i} \equiv [l_{q,i}, u_{q,i}] \cap \left(\bigcup_{j < i} [l_{q,j}, u_{q,j}] \right)^c$ for $2 \leq i \leq \tilde{N}_q$. We have our first partition: $\mathcal{F} = \bigcup_{i=1}^{\tilde{N}_q} \mathcal{F}_{q,i}$, which we call partition of \mathcal{F} at level q .

From the sequence of partitions $\{\{\mathcal{F}_{q,i} : i \leq \tilde{N}_q\}\}_{q \geq q_0}$, we derive a nested sequence of partitions as follows. The first partition is $\{\mathcal{F}_{q_0,i} : i \leq \tilde{N}_{q_0}\}$ itself. Then, recursively, at a level q such that $\{\mathcal{F}_{q,i} : i \leq \tilde{N}_q\}$ is not a successful refinement of $\{\mathcal{F}_{(q-1),i} : i \leq \tilde{N}_{q-1}\}$, we replace each partitioning set at level q

by its intersection with all partitioning sets at level $(q-1)$. All partitioning sets derived in this fashion from $\mathcal{F}_{q,i}$ are associated with the same $\Delta_{q,i}$. For a given $q \geq q_0$, the possibly new partition consists of at most $N_q = \prod_{q'=q_0}^q \tilde{N}_{q'}$ partitioning sets. Using the inequality $\sqrt{\text{Log } N_q} \leq \sum_{q'=q_0}^q \sqrt{\text{Log } \tilde{N}_{q'}}$, we see that (40) is preserved in the sense that

$$\begin{aligned} \sum_{q \geq q_0} 2^{-q} \sqrt{\text{Log } N_q} &\leq \sum_{q \geq q_0} 2^{-q} \sum_{q'=q_0}^q \sqrt{\text{Log } \tilde{N}_{q'}} \\ &\lesssim \sum_{q \geq q_0} 2^{-q} \sqrt{\text{Log } \tilde{N}_q} \leq \int_0^\delta \sqrt{\text{Log } N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon. \end{aligned} \quad (41)$$

2. At each level $q \geq q_0$ and for each $\mathcal{F}_{q,i}$ ($i \leq N_q$), fix a representative $f_{q,i} \in \mathcal{F}_{q,i}$. For every $f \in \mathcal{F}$, if $f \in \mathcal{F}_{q,i}$, then we set $\pi_q f \equiv f_{q,i}$ and $\Delta_q f \equiv \Delta_{q,i}$. Introduce $a_{q_0} \equiv 2a(2^{-q_0}) = 2^{-q_0+1}/\sqrt{\text{Log } N_{q_0}}$ and, for each $q > q_0$, $f \in \mathcal{F}$, $B_{q_0} f \equiv 0$,

$$\begin{aligned} a_q &\equiv 2^{-q+1}/\sqrt{\text{Log } N_q}, \\ A_{q-1} f &\equiv \mathbf{1} \{ \Delta_{q_0} f \leq \sqrt{n} a_{q_0}, \dots, \Delta_{q-1} f \leq \sqrt{n} a_{q-1} \}, \\ B_q f &\equiv \mathbf{1} \{ \Delta_{q_0} f \leq \sqrt{n} a_{q_0}, \dots, \Delta_{q-1} f \leq \sqrt{n} a_{q-1}, \Delta_q f > \sqrt{n} a_q \}. \end{aligned}$$

By nestedness of the sequence of partitions, $f \mapsto A_q f$ and $f \mapsto B_q f$ are constant over each $\mathcal{F}_{q,i}$ ($i \leq N_q$). Moreover, $B_q f + A_q f = A_{q-1} f$ for all $q > q_0$ and $f \in \mathcal{F}$. In addition, since $\varepsilon \mapsto a(\varepsilon)$ is non-decreasing, $\Delta_{q_0} f \leq 2F \leq 2\sqrt{n}a(\delta) \leq 2\sqrt{n}a(2^{-q_0}) = \sqrt{n}a_{q_0}$, hence $A_{q_0} f = 1$.

Using these facts, any $f \in \mathcal{F}$ decomposes as

$$f = \pi_{q_0} f + \sum_{q \geq q_0+1} (f - \pi_q f) B_q f + \sum_{q \geq q_0+1} (\pi_q f - \pi_{q-1} f) A_{q-1} f. \quad (42)$$

To see this, note first that either (i) $B_q f = 0$ for all $q \geq q_0$, which implies, by recursion, that $A_q f = 1$ for all $q \geq q_0$, or (ii) there exists $q_1 \geq q_0$ such that $B_{q_1} f = 1$, in which case $B_q f = 0$ for all $q \geq q_0, q \neq q_1$, and $A_q f = 1$ for all $q_0 \leq q < q_1$, $A_q f = 0$ for all $q \geq q_1$. If (i) holds, then we deal with a telescopic sum and (42) boils down to $f = \pi_{q_0} f + \lim_{q \rightarrow \infty} \pi_q f - \pi_{q_0} f$. The above equality is valid because both $\pi_q f$ and f are in the bracket $[l_q, u_q]$, whose size $\|u_q - l_q\|_{2, P_{Q_0, g^r}} \rightarrow 0$ as $n \rightarrow \infty$. If (ii) holds, then $f = \pi_{q_0} f + (f - \pi_{q_1} f) + \sum_{q=q_0+1}^{q_1} (\pi_q f - \pi_{q-1} f)$ is evidently true.

Define for convenience $\mathcal{F}_a = \{\pi_{q_0} f / \sqrt{n} : f \in \mathcal{F}\}$, $\mathcal{F}_b = \{\sum_{q \geq q_0+1} (f - \pi_q f) B_q f / \sqrt{n} : f \in \mathcal{F}\}$, and $\mathcal{F}_c = \{\sum_{q \geq q_0+1} (\pi_q f - \pi_{q-1} f) A_{q-1} f / \sqrt{n} : f \in \mathcal{F}\}$. Each sum in the definition of \mathcal{F}_b consists of at most one single term. Each sum in the definition of \mathcal{F}_c is either finite, or telescopic, with a limit, in which case the dominated convergence theorem guarantees that $P_{Q_0, g_n} \sum_{q \geq q_0+1} (\pi_q f - \pi_{q-1} f) A_{q-1} f = \sum_{q \geq q_0+1} P_{Q_0, g_n} (\pi_q f - \pi_{q-1} f) A_{q-1} f$. Therefore, (42) yields

$$\begin{aligned} E(\|(P_n - P_{Q_0, g_n})\|_{\mathcal{F}}) / \sqrt{n} &\leq E(\|(P_n - P_{Q_0, g_n})\|_{\mathcal{F}_a}) \\ &\quad + E(\|(P_n - P_{Q_0, g_n})\|_{\mathcal{F}_b}) + E(\|(P_n - P_{Q_0, g_n})\|_{\mathcal{F}_c}). \end{aligned} \quad (43)$$

We shall study in turn each term in the RHS expression of (43).

Class \mathcal{F}_a . For every $f \in \mathcal{F}$, (i) $|\pi_{q_0} f| \leq \sqrt{n}a(\delta) \leq \sqrt{n}a(2^{-q_0}) = \sqrt{n}a_{q_0}/2$, hence $\sup_{h \in \mathcal{F}_a} \|h\|_\infty \leq a_{q_0}/2$, and (ii) $P_{Q_0, g^r}(\pi_{q_0} f)^2 \leq \delta^2$ (true by assumption). Apply Lemma 6 with $\mathcal{F} = \mathcal{F}_a$, C the whole probability space, $U = a_{q_0}/2$, $K = 4U$, $R = 4\kappa\delta^2/3$ (an upper-bound on $nR_{n, 4U}(\pi_{q_0} f / \sqrt{n})$ valid uniformly in $f \in \mathcal{F}$ by Lemma 7): it holds that

$$\begin{aligned} nE(\|P_n - P_{Q_0, g_n}\|_{\mathcal{F}_a}) &\lesssim \delta \sqrt{\text{Log } N_{q_0}} + a_{q_0} \text{Log } N_{q_0} \\ &\leq 2^{-q_0} \sqrt{\text{Log } N_{q_0}} + 2^{-q_0+1} \frac{\text{Log } N_{q_0}}{\sqrt{\text{Log } N_{q_0}}} \\ &\leq \sum_{q \geq q_0} 2^{-q} \sqrt{\text{Log } N_q}. \end{aligned} \quad (44)$$

Class \mathcal{F}_b . For every $q > q_0$, $f \in \mathcal{F}$, $|f - \pi_q f| \leq \Delta_q f$ implies

$$|(P_n - P_{Q_0, \mathbf{g}_n})(f - \pi_q f)| \leq (P_n + P_{Q_0, \mathbf{g}_n})\Delta_q f \leq |(P_n - P_{Q_0, \mathbf{g}_n})\Delta_q f| + 2P_{Q_0, \mathbf{g}_n}\Delta_q f.$$

Thus, by using repeatedly the triangle inequality and the dominated convergence theorem, we obtain

$$\begin{aligned} E(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_b}) &\leq \sum_{q \geq q_0+1} E \left(\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n})\Delta_q f B_q f / \sqrt{n}| \right) + 2 \sum_{q \geq q_0+1} E \left(\sup_{f \in \mathcal{F}} P_{Q_0, \mathbf{g}_n} \Delta_q f B_q f / \sqrt{n} \right). \end{aligned} \quad (45)$$

Consider the first term in the RHS expression of (45). Fix $q > q_0$. Note that $f, f' \in \mathcal{F}_{q,i}$ implies $\Delta_q f B_q f = \Delta_q f' B_q f'$. So, the supremum $\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n})\Delta_q f B_q f / \sqrt{n}|$ is actually a maximum over a set of cardinality N_q . Moreover, for each $f \in \mathcal{F}$, (i) $0 \leq \Delta_q f B_q f \leq \Delta_{q-1} f B_q f \leq \sqrt{n} a_{q-1}$, hence $\sup_{h \in \mathcal{F}_b} \|h\|_\infty \leq a_{q-1}$, and (ii) $P_{Q_0, g^r}(\Delta_q f B_q f)^2 \leq 2^{-2q}$. Apply Lemma 6 with $\mathcal{F} = \mathcal{F}_b$, C the whole probability space, $U = a_{q-1}$, $K = 4U$, $R = 4\kappa 2^{-2q}/3$ (an upper-bound on $nR_{n,4U}(\Delta_q f B_q f / \sqrt{n})$ valid uniformly in $f \in \mathcal{F}$ by Lemma 7): it holds that

$$\begin{aligned} nE \left(\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n})\Delta_q f B_q f / \sqrt{n}| \right) &\lesssim 2^{-q} \sqrt{\text{Log } N_q} + a_{q-1} \text{Log } N_q \\ &= 2^{-q} \sqrt{\text{Log } N_q} + 2^{-q+2} \frac{\text{Log } N_q}{\sqrt{\text{Log } N_q}} \\ &\lesssim 2^{-q} \sqrt{\text{Log } N_q}. \end{aligned} \quad (46)$$

Consider now the second term in (45). Fix $q > q_0$ and $f \in \mathcal{F}$. Since $B_q f = 1$ only if $\sqrt{n} a_q < \Delta_q f$, it follows that

$$\sqrt{n} a_q P_{Q_0, g_i} \Delta_q f B_q f \leq P_{Q_0, g_i} (\Delta_q f)^2 B_q f \leq 2^{-2q}$$

for every $1 \leq i \leq n$. Therefore,

$$\sup_{f \in \mathcal{F}} P_{Q_0, \mathbf{g}_n} \Delta_q f B_q f / \sqrt{n} \leq 2^{-2q} / n a_q \lesssim 2^{-q} \sqrt{\text{Log } N_q} / n. \quad (47)$$

By (45), summing up (46) and (47) over $q > q_0$ finally yields

$$nE(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_b}) \lesssim \sum_{q \geq q_0+1} 2^{-q} \sqrt{\text{Log } N_q}. \quad (48)$$

Class \mathcal{F}_c . Fix $q > q_0$. Note that $f, f' \in \mathcal{F}_{q,i}$ implies $(\pi_q f - \pi_{q-1} f) A_q f = (\pi_q f' - \pi_{q-1} f') A_q f'$. So, the supremum $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_c}$ is actually a maximum over a set of cardinality N_q . Moreover, for each $f \in \mathcal{F}$, (i) $|\pi_q f - \pi_{q-1} f| A_{q-1} f \leq \Delta_{q-1} f A_{q-1} f \leq \sqrt{n} a_{q-1}$, hence $\sup_{h \in \mathcal{F}_c} \|h\|_\infty \leq a_{q-1}$, from which we also deduce that (ii) $P_{Q_0, g^r}((\pi_q f - \pi_{q-1} f) A_{q-1} f)^2 \leq P_{Q_0, g^r}(\Delta_q f)^2 \leq 2^{-2q}$. Therefore, the same reasoning as the one which lead us to (48) applies again, and we obtain

$$nE(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_c}) \lesssim \sum_{q \geq q_0+1} 2^{-q} \sqrt{\text{Log } N_q}. \quad (49)$$

Combining (43), (44), (48), (49), and (41) completes the proof. \square

To prove Proposition 4, we must study the convergence in probability of empirical processes indexed by estimated functions. Lemma 10 below provides sufficient conditions to derive such convergences. The version of this lemma under a i.i.d. sampling scheme is given by (van der Vaart and Wellner, 2007, Theorem 2.2). Here, we provide its extension to the current targeted adaptive sampling scheme. The proof of Lemma 10 hinges on Lemma 9.

Lemma 10 (convergence of empirical processes indexed by estimated functions). *For each $n \geq 1$, let $\mathcal{F}_n = \{f_{\theta, \eta} : \theta \in \Theta, \eta \in T_n\}$ be a class of measurable, real-valued functions, with envelope function F_n . Suppose the following holds:*

(a) *The sequence $\{F_n\}_{n \geq 1}$ satisfies the Lindeberg condition: $P_{Q_0, g^r} F_n^2 = O(1)$ and, for every $\delta > 0$, $P_{Q_0, g^r} F_n^2 \mathbf{1}\{F_n > \delta \sqrt{n}\} = o(1)$.*

(b) *If $\delta_n = o(1)$, then it holds that $J(\delta_n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$.*

If $\eta_n \in T_n$ is such that $\sup_{\theta \in \Theta} P_{Q_0, g^r}(f_{\theta, \eta_n} - f_{\theta, \eta_0})^2 = o_P(1)$ for some $\eta_0 \in \cap_{p \geq 1} \cup_{n \geq p} T_n$, then $\sup_{\theta \in \Theta} |\sqrt{n}(P_n - P_{Q_0, g_n})(f_{\theta, \eta_n} - f_{\theta, \eta_0})| = o_P(1)$.

Proof of lemma 10. Define the random class $\tilde{\mathcal{F}}_n^0 \equiv \{f_{\theta, \eta_n} - f_{\theta, \eta_0} : \theta \in \Theta\}$. We wish to prove that $\sqrt{n}\|P_n - P_{Q_0, g_n}\|_{\tilde{\mathcal{F}}_n^0} = o_P(1)$. Set arbitrarily $\alpha > 0, \varepsilon > 0$, and introduce $\mathcal{F}_n^0 \equiv \{f_{\theta, \eta} - f_{\theta, \eta_0} : \theta \in \Theta, \eta \in T_n\}$, which admits $2F_n$ as an envelope function. For every $\delta > 0$, it holds that $J(\delta, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) \lesssim J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})$. Consequently, by (b) and (ii) in Lemma 12 below, there exists $\delta_0 > 0$ and $n_0 \geq 1$ such that, for all $n \geq n_0$, $J(\delta_0, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \alpha\varepsilon$. Define $T_n^0(\delta_0) \equiv \{\eta \in T_n : \sup_{\theta \in \Theta} P_{Q_0, g^r}(f_{\theta, \eta} - f_{\theta, \eta_0}) \leq \delta_0^2\}$, and $\mathcal{F}_n^0(\delta_0) \equiv \{f_{\theta, \eta} - f_{\theta, \eta_0} : \theta \in \Theta, \eta \in T_n^0(\delta_0)\} \subset \mathcal{F}_n^0$. By assumption, there exists $n_1 \geq 1$ such that $P(\eta_n \notin T_n^0(\delta_0)) \leq \varepsilon$ whenever $n \geq n_1$.

Set $n \geq \max(n_0, n_1)$. By the Markov inequality, and because $\eta_n \in T_n^0(\delta_0)$ implies $\tilde{\mathcal{F}}_n^0 \subset \mathcal{F}_n^0(\delta_0)$, it holds that

$$\begin{aligned} P\left(\sqrt{n}\|P_n - P_{Q_0, g_n}\|_{\tilde{\mathcal{F}}_n^0} \geq \alpha\right) &\leq P(\eta_n \notin T_n^0(\delta_0)) + \alpha^{-1} E\left(\sqrt{n}\|P_n - P_{Q_0, g_n}\|_{\tilde{\mathcal{F}}_n^0} \mathbf{1}\{\eta_n \in T_n^0(\delta_0)\}\right) \\ &\leq \varepsilon + \alpha^{-1} E\left(\sqrt{n}\|P_n - P_{Q_0, g_n}\|_{\mathcal{F}_n^0(\delta_0)} \mathbf{1}\{\eta_n \in T_n^0(\delta_0)\}\right) \\ &\leq \varepsilon + \alpha^{-1} E\left(\sqrt{n}\|P_n - P_{Q_0, g_n}\|_{\mathcal{F}_n^0(\delta_0)}\right). \end{aligned} \quad (50)$$

By Lemma 9, whose conditions are met,

$$\begin{aligned} E\left(\sqrt{n}\|P_n - P_{Q_0, g_n}\|_{\mathcal{F}_n^0(\delta_0)}\right) &\lesssim J(\delta_0, \mathcal{F}_n^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}}) + \sqrt{n} P_{Q_0, g^r} F_n \mathbf{1}\{F_n > \sqrt{n} a_n(\delta_0)/2\} \\ &\lesssim J(\delta_0, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) \\ &\quad + a_n(\delta_0)^{-1} P_{Q_0, g^r} F_n^2 \mathbf{1}\{F_n > \sqrt{n} a_n(\delta_0)/2\}, \end{aligned} \quad (51)$$

where $a_n(\delta_0) \equiv \delta_0 / \sqrt{\text{Log } N(\delta_0, \mathcal{F}_n^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}})}$. By (i) in Lemma 12 below, $m \mapsto J(\delta_0, \mathcal{F}_m^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}})$ is a bounded function. We also know that, for all $m \geq 1$,

$$J(\delta_0, \mathcal{F}_m^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}}) \geq \delta_0 \sqrt{\text{Log } N(\delta_0, \mathcal{F}_m^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}})} = \delta_0^2 / a_m(\delta_0).$$

In particular, $m \mapsto a_m(\delta_0)$ must be bounded away from 0. Let $c > 0$ be such that $a_m(\delta_0) \geq c$ for all $m \geq 1$. With this in mind, (51) implies

$$E\left(\sqrt{n}\|P_n - P_{Q_0, g_n}\|_{\mathcal{F}_n^0(\delta_0)}\right) \leq J(\delta_0, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) + c^{-1} P_{Q_0, g^r} F_n^2 \mathbf{1}\{F_n > \sqrt{nc}/2\},$$

where $J(\delta_0, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \alpha\varepsilon$ by construction. Assumption (b) guarantees that there exists $n_2 \geq 1$ such that $m \geq n_2$ implies $P_{Q_0, g^r} F_m^2 \mathbf{1}\{F_m > \sqrt{mc}/2\} \leq \alpha\varepsilon$. In summary, provided that $n \geq \max(n_0, n_1, n_2)$, (50) and (51) yield $P(\sqrt{n}\|P_n - P_{Q_0, g_n}\|_{\tilde{\mathcal{F}}_n^0} \geq \alpha) \leq 3\varepsilon$. In other words, $\sqrt{n}\|P_n - P_{Q_0, g_n}\|_{\tilde{\mathcal{F}}_n^0} = o_P(1)$. This completes the proof. \square

The next two lemmas proved useful in our demonstrations.

Lemma 11. *Let \mathcal{F} be a uniformly bounded class of measurable, real-valued functions. Let h, h' be two measurable, bounded, real-valued functions. We do not assume that $h, h' \in \mathcal{F}$. Set $\delta > 0$.*

- Define \mathcal{F}' equal either to $\{f - h : f \in \mathcal{F}\}$, or $\{f|h| : f \in \mathcal{F}\}$, or $\{f|h| + f'|h'| : f, f' \in \mathcal{F}\}$, or $\{|f| : f \in \mathcal{F}\}$, or $\{f^2 : f \in \mathcal{F}\}$, or $\{\phi(f) : f \in \mathcal{F}\}$ where ϕ is non-decreasing and Lipschitz, or $\{h \log(f) + (1 - h) \log(1 - f) : f \in \mathcal{F}\}$ if the functions in \mathcal{F} and h take their values in $[0, 1]$ and are uniformly bounded away from 0 and 1. It holds that $J(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \lesssim J(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$.
- Define $\mathcal{F}' = \{\sqrt{f} : f \in \mathcal{F}\}$ if the functions in \mathcal{F} are non-negative. It holds that $J(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \lesssim J(\sqrt{\delta}, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$.

Proof of Lemma 11. Fix $\delta > 0$ and $M > 0$ such that $\sup_{f \in \mathcal{F}} \|f\|_\infty \leq M < \infty$. Let $N \equiv N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, and consider a collection of δ -brackets $\{[l_i, u_i] : i \leq N\}$ that covers \mathcal{F} .

- Case $\mathcal{F}' = \{f - h : f \in \mathcal{F}\}$. The collection of δ -brackets obtained by substituting $[l_i - h, u_i - h]$ for $[l_i, u_i]$, all $i \leq N$, covers \mathcal{F}' . This proves the first claim.
- Case $\mathcal{F}' = \{f|h| : f \in \mathcal{F}\}$. The collection of $\delta\|h\|_\infty$ -brackets obtained by substituting $[l_i|h|, u_i|h|]$ for $[l_i, u_i]$, all $i \leq N$, covers \mathcal{F}' . This proves the second claim.
- Case $\mathcal{F}' = \{f|h| + f'|h'| : f, f' \in \mathcal{F}\}$. The collection of brackets consisting of $[l_i|h| + l_j|h'|, u_i|h| + u_j|h'|]$, all $i, j \leq N$, covers \mathcal{F}' . Consider $i, j \leq N$, and set $\gamma_{ij} \equiv u_i|h| + u_j|h'|$, $\lambda_{ij} \equiv l_i|h| + l_j|h'|$, $c \equiv 2\sqrt{\|h\|_\infty^2 + \|h'\|_\infty^2}$: it holds that $P_{Q_0, g^r}(\gamma_{ij} - \lambda_{ij})^2 \leq c^2 \delta^2$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta/c, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})^2$, from which the third claim follows.

- Case $\mathcal{F}' = \{|f| : f \in \mathcal{F}\}$. Set $f \in \mathcal{F}$, and assume without loss of generality that $f \in [l_1, u_1]$. Define $F_+ \equiv \mathbf{1}\{f > 0\}$, $F_- \equiv \mathbf{1}\{f < 0\}$, $G_+ \equiv \mathbf{1}\{l_1 > 0\}$, $G_- \equiv \mathbf{1}\{u_1 < 0\}$, and $G_0 \equiv \mathbf{1}\{l_1 \leq 0 \leq u_1\}$. Then

$$F_+(l_1)_+ + F_-(u_1)_- \leq |f| \leq F_+u_1 - F_-l_1$$

with

$$F_+(l_1)_+ + F_-(u_1)_- = G_+l_1 - G_-u_1 \equiv \lambda_1,$$

and

$$\begin{aligned} F_+u_1 - F_-l_1 &= G_+u_1 - G_-l_1 + G_0(F_+u_1 - F_-l_1) \\ &\leq G_+u_1 - G_-l_1 + G_0(u_1 - l_1) \equiv \gamma_1. \end{aligned}$$

Thus $\lambda_1 \leq |f| \leq \gamma_1$, where $\gamma_1 - \lambda_1 = u_1 - l_1$, hence $P_{Q_0, g^r}(\gamma_1 - \lambda_1)^2 \leq \delta^2$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, from which the fourth claim follows.

- Case $\mathcal{F}' = \{f^2 : f \in \mathcal{F}\}$. Set $f \in \mathcal{F}$, and assume without loss of generality that $f \in [l_1, u_1]$. Let $[\lambda_1, \gamma_1]$ be the bracket that we just built. The inequalities $\lambda_1 \geq 0$ and $f^2 \leq M^2$ imply that $\lambda_1^2 \leq f^2 \leq \min(\gamma_1^2, M^2)$. Set $\lambda_2 \equiv \lambda_1$, $\gamma_2 \equiv \sqrt{\min(\gamma_1^2, M^2)}$ so that $\lambda_2^2 \leq f^2 \leq \gamma_2^2$. Obviously, $\gamma_2^2 - \lambda_2^2 \leq 2\gamma_2(\gamma_2 - \lambda_2) \leq 2M(\gamma_1 - \lambda_1)$, hence $P_{Q_0, g^r}(\gamma_2^2 - \lambda_2^2)^2 \leq 4M^2\delta^2$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta/2M, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, from which the fifth claim follows.

- Case $\mathcal{F}' = \{\phi(f) : f \in \mathcal{F}\}$. Say that ϕ is c -Lipschitz. The collection of $c\delta$ -brackets obtained by substituting $[\phi(l_i), \phi(u_i)]$ for $[l_i, u_i]$, all $i \leq N$, covers \mathcal{F}' . This proves the sixth claim.
- Case $\mathcal{F}' = \{h \log(f) + (1 - h) \log(1 - f) : f \in \mathcal{F}\}$. Set $f \in \mathcal{F}$, and assume without loss of generality that $f \in [l_1, u_1]$ and $0 < \inf_{f \in \mathcal{F}} f \leq l_1 \leq u_1 < \sup_{f \in \mathcal{F}} f < 1$. Define $\lambda_3 \equiv h \log(l) + (1 - h) \log(1 - l)$ and $\gamma_3 \equiv h \log(u) + (1 - h) \log(1 - u)$. It holds that $\lambda_3 \leq h \log(f) + (1 - h) \log(1 - f) \leq \gamma_3$. Moreover, $0 \leq \gamma_3 - \lambda_3 \lesssim (u_1 - l_1)$ because \log is Lipschitz on any compact subset of $(0, 1)$. Consequently, there exists $c \geq 1$ such that $P_{Q_0, g^r}(\gamma_3 - \lambda_3)^2 \leq c^2 \delta^2$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta/c, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, from which the seventh claim follows.

- *Case $\mathcal{F}' = \{\sqrt{f} : f \in \mathcal{F}\}$.* Set $f \in \mathcal{F}$, and assume without loss of generality that $f \in [l_1, u_1]$ and $l_1 \geq 0$. Then $\sqrt{l_1} \leq \sqrt{f} \leq \sqrt{u_1}$. Moreover, $(\sqrt{u_1} - \sqrt{l_1})^2 \leq (\sqrt{u_1} - \sqrt{l_1})(\sqrt{u_1} + \sqrt{l_1}) = u_1 - l_1$. The Cauchy-Schwarz inequality yields $P_{Q_0, g^r}(u_1 - l_1) \leq \sqrt{P_{Q_0, g^r}(u_1 - l_1)^2} \leq \sqrt{\delta}$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\sqrt{\delta}, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, from which the eighth claim follows.

This completes the proof. \square

Lemma 12. *For each $n \geq 1$, let \mathcal{F}_n be a class of measurable, real-valued functions such that $\delta_n = o(1)$ implies $J(\delta_n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$. Then (i) $J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = O(1)$ for every $\delta > 0$, and (ii) for every $\varepsilon > 0$, there exist $\delta > 0$ and $n_1 \geq 1$ such that $J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \varepsilon$ for all $n \geq n_1$.*

Proof of Lemma 12. We prove (i) and (ii) by contradiction.

Suppose there exists $\delta > 0$ such that $\limsup_{n \rightarrow \infty} J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = \infty$. Without loss of generality, we can assume that $J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \geq 2^{2n}$ for each $n \geq 1$. Now,

$$J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = J(\delta/2, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) + \int_{\delta/2}^{\delta} \sqrt{\log N(\varepsilon, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon,$$

with

$$2J(\delta/2, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \geq \delta \sqrt{\log N(\delta/2, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})} \geq 2 \int_{\delta/2}^{\delta} \sqrt{\log N(\varepsilon, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon.$$

Therefore, $J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \geq 2^{2n}$ implies $J(\delta/2, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \geq 2^{2n}/2$ hence, by recursion,

$$J(\delta/2^n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \geq 2^{2n}/2^n = 2^n.$$

The sequence $\{\delta_n\}_{n \geq 1}$ given by $\delta_n = \delta/2^n$ satisfies $\delta_n = o(1)$ and $\lim_{n \rightarrow \infty} J(\delta_n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = \infty$, in contradiction with the assumption of the lemma. This completes the proof of (i).

Now, assume that there exists $\varepsilon > 0$ such that, for all $\delta > 0$, there exists $n_1 \geq 1$ for which $J(\delta, \mathcal{F}_{n_1}, \|\cdot\|_{2, P_{Q_0, g^r}}) > \varepsilon$. In particular, we can construct by recursion an increasing sequence $\{\varphi(n)\}_{n \geq 1}$ such that, for all $n \geq 1$, $J(1/n, \mathcal{F}_{\varphi(n)}, \|\cdot\|_{2, P_{Q_0, g^r}}) > \varepsilon$. This induces the existence of a sequence $\{\delta_n\}_{n \geq 1}$ such that $\delta_n = o(1)$ and $\limsup_{n \rightarrow \infty} J(\delta_n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) > \varepsilon$, in contradiction with the assumption of the lemma. This completes the proof of (ii). \square

References

- A. C. Atkinson and A. Biswas. Adaptive biased-coin designs for skewing the allocation proportion in clinical trials with normal responses. *Stat. Med.*, 24(16):2477–2492, 2005.
- U. Bandyopadhyay and A. Biswas. Adaptive designs for normal responses with prognostic factors. *Biometrika*, 88(2):409–419, 2001.
- P. J. Bickel, C. A. J. Klaassen, Y. Ritov, and J. A. Wellner. *Efficient and adaptive estimation for semiparametric models*. Springer-Verlag, New York, 1998. Reprint of the 1993 original.
- A. Biswas, R. Bhattacharya, and E. Park. On a class of optimal covariate-adjusted response-adaptive designs for survival outcomes. *Statistical methods in medical research*, 2014.
- A. Chambaz and M. J. van der Laan. Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate: Theoretical study. *Int. J. Biostat.*, 7(1), 2011a. Article 10.
- A. Chambaz and M. J. van der Laan. Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate: Simulation study. *Int. J. Biostat.*, 7(1), 2011b. Article 11.

- A. Chambaz and M. J. van der Laan. Estimation and testing in targeted group sequential covariate-adjusted randomized clinical trials. Technical report 299, Division of Biostatistics, University of California, Berkeley, 2011c.
- A. Chambaz and M. J. van der Laan. Inference in targeted group sequential covariate-adjusted randomized clinical trials. *Scandinavian Journal of Statistics*, 41(1):104–140, 2013.
- Y. I. Chang and E. Park. Sequential estimation for covariate-adjusted response-adaptive designs. *J. Korean Statistical Society*, 42(1):105–116, 2013.
- J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2010. URL <http://www.jstatsoft.org/v33/i01/>.
- F. Hu and W. F. Rosenberger. Optimality, variability, power: evaluating response-adaptive randomization procedures for treatment comparisons. *Journal of the American Statistical Association*, 98(463):671–678, 2003.
- F. Hu and W. F. Rosenberger. *The theory of response-adaptive randomization in clinical trials*, volume 525. John Wiley & Sons, 2006.
- T. Huang, Z. Liu, and F. Hu. Longitudinal covariate-adjusted response-adaptive randomized designs. *J. Statistical Planning and Inference*, 143(10):1816–1827, 2013.
- A. Ivanova. A play-the-winner type urn model with reduced variability. *Metrika*, 58:1–13, 2003.
- C. Jennison and B. W. Turnbull. *Group Sequential Methods with Applications to Clinical Trials*. Chapman & Hall/CRC, Boca Raton, FL, 2000.
- K. Moore and M. J. van der Laan. Covariate adjustment in randomized trials with binary outcomes: targeted maximum likelihood estimation. *Stat. Med.*, 28(1):39–64, 2009.
- J. Pearl. *Causality*. Cambridge University Press, Cambridge, 2000. Models, reasoning, and inference.
- D. Pollard. *Convergence of stochastic processes*. Springer Series in Statistics. Springer-Verlag, New York, 1984. ISBN 0-387-90990-7.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014. URL <http://www.R-project.org/>.
- W. F. Rosenberger. New directions in adaptive designs. *Statist. Sci.*, 11:137–149, 1996.
- W. F. Rosenberger and F. Hu. Maximizing power and minimizing treatment failures in clinical trials. *Clinical Trials*, 1(141-147), 2004.
- W. F. Rosenberger, A. N. Vidyashankar, and D. K. Agarwal. Covariate-adjusted response-adaptive designs for binary response. *Journal of biopharmaceutical statistics*, 11(4):227–236, 2001.
- W. F. Rosenberger, O. Sverdlov, and F. Hu. Adaptive randomization for clinical trials. *J Biopharm Stat*, 22(4):719–36, 2012.
- M. Rosenblum. *Robust Analysis of RCTs Using Generalized Linear Models*, chapter 11. Springer-Verlag, New York, 2011.
- P. K. Sen and J. M. Singer. *Large sample methods in statistics*. Chapman & Hall, New York, 1993. ISBN 0-412-04221-5. An introduction with applications.
- J. Shao and X. Yu. Validity of tests under covariate-adaptive biased coin randomization and generalized linear models. *Biometrics*, 69(4):960–969, 2013.

- J. Shao, X. Yu, and B. Zhong. A theory for testing hypotheses under covariate-adaptive randomization. *Biometrika*, 97(2):347–360, 2010.
- O. Sverdlov, W. F. Rosenberger, and Y. Ryznik. Utility of covariate-adjusted response-adaptive randomization in survival trials. *Statistics in Biopharmaceutical Research*, 5(1):38–53, 2013.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- M. J. van der Laan and J. M. Robins. *Unified methods for censored longitudinal data and causality*. Springer Series in Statistics. Springer-Verlag, New York, 2003.
- M. J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *Int. J. Biostat.*, 2(1), 2006.
- A. W. van der Vaart. *Asymptotic statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998.
- A. W. van der Vaart and J. A. Wellner. *Weak Convergence*. Springer, 1996.
- A. W. van der Vaart and J. A. Wellner. Empirical processes indexed by estimated functions. *Lecture Notes-Monograph Series*, pages 234–252, 2007.
- R. van Handel. On the minimal penalty for markov order estimation. *Probability Theory and Related Fields*, 150:709–738, 2011.
- L. J. Wei and S. Durham. The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association*, 73(840-843), 1978.
- L-X. Zhang and F-F. Hu. A new family of covariate-adjusted response-adaptive designs and their properties. *Appl. Math. J. Chinese Univ. Ser. B*, 24(1):1–13, 2009.
- L-X. Zhang and W. F. Rosenberger. Adaptive randomization in clinical trials. In K. Hinkelmann, editor, *Design and Analysis of Experiments: Special Designs and Applications*, volume 3, chapter 7. John Wiley & Sons, 2012.
- L-X. Zhang, F. Hu, S. H. Cheung, and W. S. Chan. Asymptotic properties of covariate-adjusted response-adaptive designs. *Ann. Statist.*, 35(3):1166–1182, 2007.

Figure 1: **Performance TMLE across 500 simulations.** Each row corresponds to a performance measure (top: bias, middle: sample variance, bottom: MSE). Each column corresponds to a working model for the optimal randomization scheme (left: mis-specified working model \mathcal{G}_1^m , right: correctly specified working model \mathcal{G}_1^c). The red and green dots correspond to our CARA RCT with different working models for the conditional response (red: LASSO working model $\mathcal{Q}_{1,n}^\ell$, green: parametric working model $\mathcal{Q}_{1,n}^p$). The blue dots correspond to a RCT with a fixed design set to the balanced randomization scheme g^b and $\mathcal{Q}_{1,n}^p$ as (fixed) parametric working model for the conditional response.

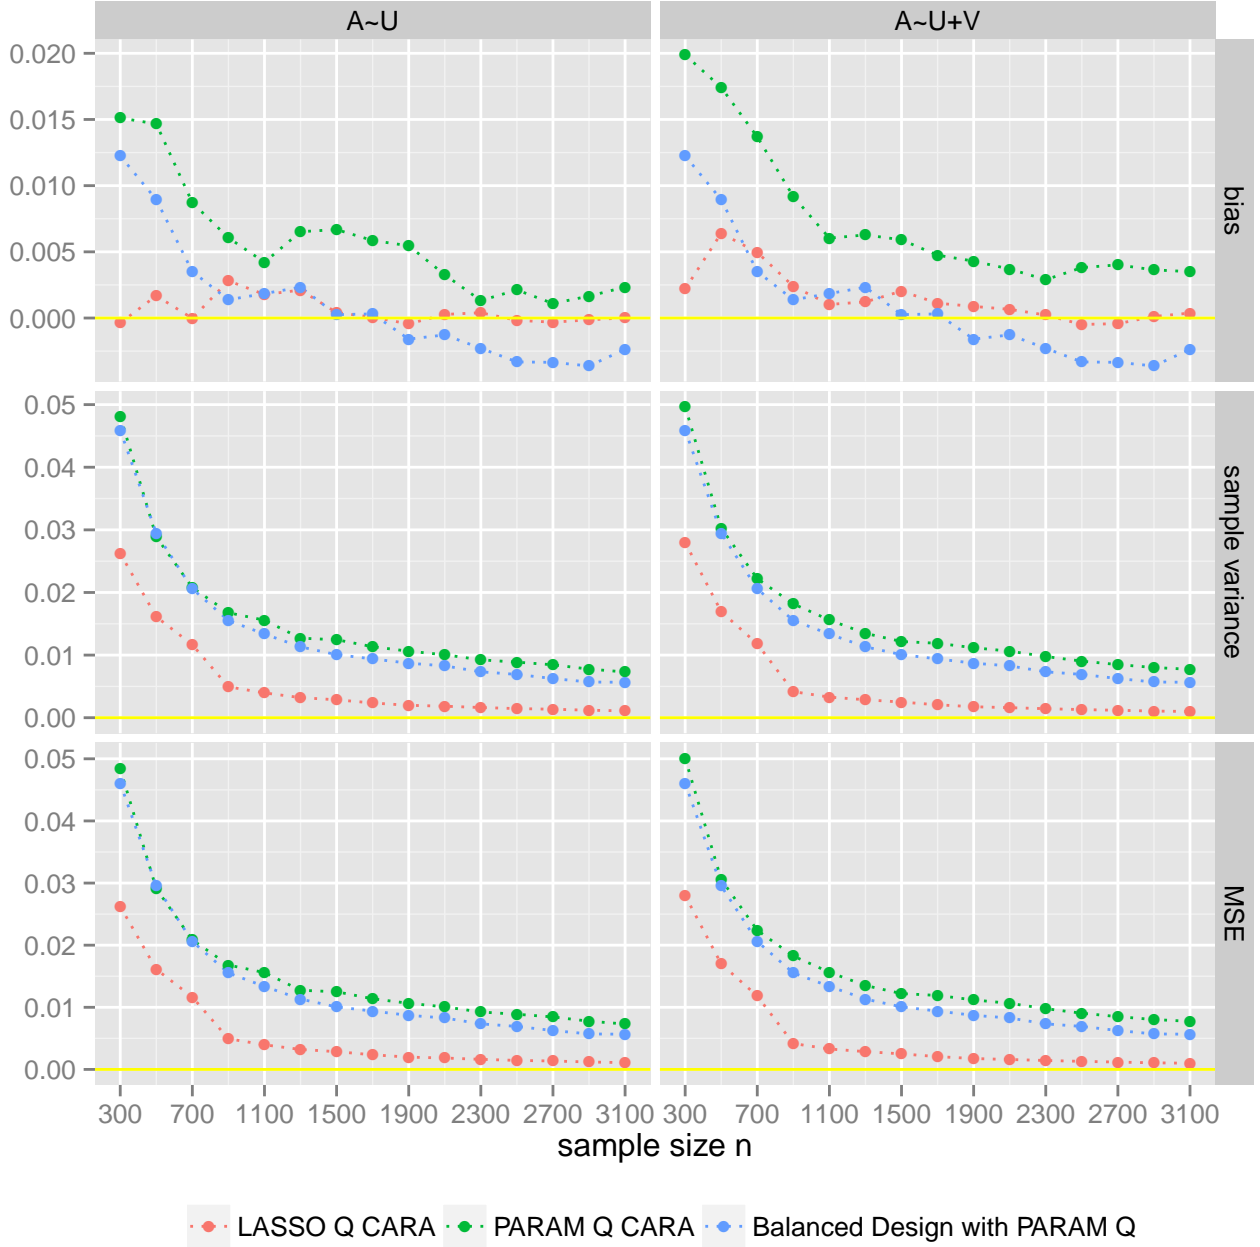


Figure 2: **Empirical coverage of CLT-based 95% CIs across 500 simulations.** Each column corresponds to a working model for the optimal randomization scheme (left: mis-specified working model \mathcal{G}_1^m , right: correctly specified working model \mathcal{G}_1^c). The red and green dots correspond to our CARA RCT with different working models for the conditional response (red: LASSO working model $\mathcal{Q}_{1,n}^\ell$, green: parametric working model $\mathcal{Q}_{1,n}^p$). The blue dots correspond to a RCT with a fixed design set to the balanced randomization scheme g^b and $\mathcal{Q}_{1,n}^p$ as (fixed) parametric working model for the conditional response. The yellow lines indicate the confidence levels 95% and 94%.

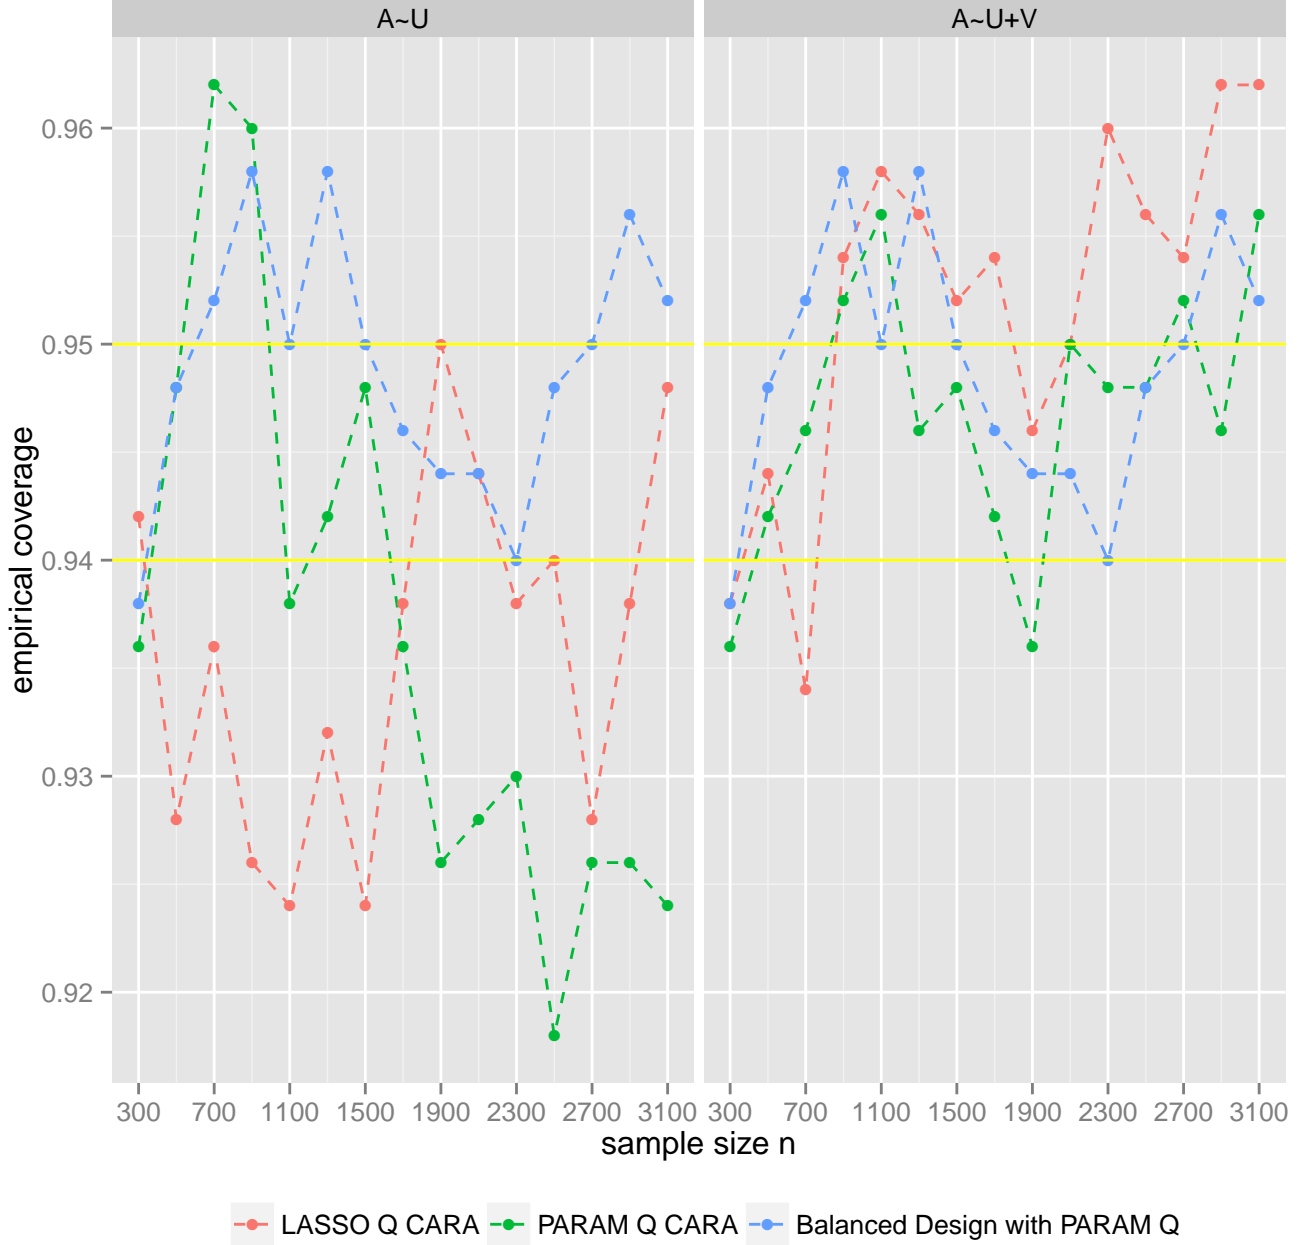


Figure 3: **Hypotheses testing to assess, across 500 simulations, the quality of coverage guaranteed by the CLT-based 95%-CI.** Each row corresponds to a null hypothesis $H_0^{1-\alpha} : \pi \geq 1 - \alpha$ (top: $\alpha = 5\%$, bottom: $\alpha = 6\%$), where π is the actual coverage guaranteed by each CI, which should satisfy by construction $\pi \geq 95\%$. Each column corresponds to a working model for the optimal randomization scheme (left: mis-specified working model \mathcal{G}_1^m , right: correctly specified working model \mathcal{G}_1^c). The red and green colors correspond to our CARA RCT with different working models for the conditional response (red: LASSO working model $\mathcal{Q}_{1,n}^l$, green: parametric working model $\mathcal{Q}_{1,n}^p$). The blue color correspond to a RCT with a fixed design set to the balanced randomization scheme g^b and $\mathcal{Q}_{1,n}^p$ as (fixed) parametric working model for the conditional response. The yellow line indicates the threshold 0.05.

