Published in final edited form as:

Scand Stat Theory Appl. 2014 March ; 41(1): 104–140. doi:10.1111/sjos.12013.

Estimation and testing in targeted group sequential covariate-adjusted randomized clinical trials

A. Chambaz^{1,*} and M. J. van der Laan²

¹MAP5, Université Paris Descartes and CNRS

²University of California, Berkeley

Abstract

This article is devoted to the construction and asymptotic study of adaptive group sequential covariate-adjusted randomized clinical trials analyzed through the prism of the semipara-metric methodology of targeted maximum likelihood estimation (TMLE). We show how to build, as the data accrue group-sequentially, a sampling design which targets a user-supplied optimal design. We also show how to carry out a sound TMLE statistical inference based on such an adaptive sampling scheme (therefore extending some results known in the i.i.d setting only so far), and how group-sequential testing applies on top of it. The procedure is robust (*i.e.*, consistent even if the working model is misspecified). A simulation study confirms the theoretical results, and validates the conjecture that the procedure may also be efficient.

Keywords

Adaptive design; asymptotic normality; canonical distribution; clinical trial; contiguity; group-sequential testing; robustness; targeted maximum likelihood methodology

1 Introduction

This article is devoted to the construction and asymptotic study of adaptive group sequential covariate-adjusted randomized clinical trials (RCTs) analyzed through the prism of the targeted maximum likelihood estimation (TMLE) methodology. The considered observed data structure writes as $O = (W, A, Y)$ where W is a vector of baseline covariates, A is the treatment assignment and Y the primary outcome of interest. We assume that A is binary and that Y is one-dimensional. Typical parameters of scientific interest are $\Psi_+ = E\{E(Y|A = 1, W) - E(Y|A = 0, W)\}$ (additive scale, which we consider hereafter) or $\Psi_\times = \log E\{E(Y|A = 1, W)\} - \log E\{E(Y|A = 0, W)\}$ (multiplicative scale). Such parameters can be interpreted causally whenever one is willing to assume the existence of a full data structure $X = (W, Y(0), Y(1))$ containing the vector of baseline covariates and the two counterfactual outcomes under the two possible treatments, and such that $Y = Y(A)$ and A conditionally independent of X given W . If so indeed, $\Psi_+ = EY(1) - EY(0)$ and $\Psi_\times = \log EY(1) - \log EY(0)$. Let us now explain what we mean by adaptive group sequential covariate-adjusted RCT.

*This collaboration was initiated while Antoine Chambaz was a visiting scholar at UC Berkeley, supported in part by a Fulbright Research Grant and the CNRS.

Adaptive group sequential covariate-adjusted RCT.

By *adaptive covariate-adjusted design*, we mean in the setting of this article a clinical trial design that allows the investigator to dynamically modify its course through data-driven adjustment of the randomization probability based on data accrued so far, without negatively impacting the statistical integrity of the trial. Moreover, the patient's baseline covariates may be taken into account for the random treatment assignment. This definition is slightly adapted from (Golub, 2006). In particular, we assume with a view to the definition of pre-specified sampling plans given in (Emerson, 2006) that, prior to collection of the data, the trial protocol specifies the parameter of scientific interest, the inferential method and confidence level to be used when constructing a confidence interval for the latter parameter.

Furthermore, we assume that the investigator specifies beforehand in the trial protocol a criterion of special interest which yields a notion of *optimal randomization scheme* that we therefore wish to target. For instance, the criterion could translate the necessity to minimize the number of patients assigned to their corresponding inferior treatment arm, subject to level and power constraints. Or the criterion could translate the necessity that a result be available as quickly as possible, subject to level and power constraints. The sole restriction on the criterion is that it must yield an optimal randomization scheme which can be approximated from the data accrued so far. The two examples above comply with this restriction.

We focus in this article on targeted maximum likelihood estimation, as introduced by van der Laan and Rubin (2006) in the independent identically distributed (i.i.d) setting. The extension to the setting of adaptive RCTs was first considered in (van der Laan, 2008), upon which this article relies. In addition, we choose to consider specifically the second criterion cited above. Consequently, the optimal randomization scheme is the so-called *Neyman allocation*, which minimizes the asymptotic variance of the TMLE of the parameter of interest. We emphasize that there is nothing special about targeting the Neyman allocation, the whole methodology applying equally well to a large class of optimal randomization schemes derived from a variety of valid criteria.

By *adaptive group sequential design*, we refer to the possibility to adjust the randomization scheme only by blocks of c patients, where $c \geq 1$ is a pre-specified integer (the case that $c = 1$ corresponds to a fully sequential adaptive design). The expression also refers to the fact that group sequential testing methods can be equally well applied on top of adaptive designs, an extension that we do not consider here. Although all our results (and their proofs) still hold for any $c \geq 1$, we consider the case $c = 1$ in the theoretical part of the article for simplicity's sake, but the case that $c > 1$ is considered in the simulation study.

Short bibliography.

The literature on adaptive designs is vast and our review is not comprehensive. Quite misleadingly, the expression “adaptive design” has also been used in the literature for sequential testing and, in general, for designs that allow data-adaptive stopping times for the whole study (or for certain treatment arms) which achieve the wished type I and type II errors requirements when testing a null hypothesis against its alternative.

Of course, data-adaptive randomization schemes have a long history that goes back to the 1930s, and we refer to Section 1.2 in (Hu and Rosenberger, 2006), Section 17.4 in (Jennison and Turnbull, 2000) and (Rosenberger, 1996) for a comprehensive historical perspective.

Many articles are devoted to the study of “response adaptive designs”, an expression implicitly suggesting that those designs only depend on past responses of previous patients and not on the corresponding covariates. We refer to (Hu and Rosenberger, 2006; Chambaz and van der Laan, 2011) for a bibliography on that topic. On the contrary, covariate-adjusted response adaptive (CARA) randomizations tackle the so-called issue of heterogeneity (*i.e.*, the use of covariates in adaptive designs), by dynamically calculating the allocation probabilities on the basis of previous responses and current and past values of certain covariates. In this view, this article studies a new type of CARA procedure. The interest in CARA procedures is more recent, and there is a steadily growing number of articles dedicated to their study, starting with (Rosenberger et al., 2001; Bandyopadhyay and Biswas, 2001), then (Atkinson and Biswas, 2005; Zhang et al., 2007; Zha; Shao et al., 2010) among others. The latter articles are typically concerned with the convergence (almost sure and in law) of the allocation probabilities vector and of the estimator of the parameter *in a correctly specified parametric model* ((Shao et al., 2010) is devoted to the testing issue).

By contrast, the consistency and asymptotic normality results that we obtain in this article are robust to model misspecification. Thus, they notably contribute significantly to solving the question raised by the Food & Drug Administration (2006):

When is it valid to modify randomization based on results, for example, in a combined phase 2/3 cancer trial?

Finally, this article mainly relies on (Chambaz and van der Laan, 2009, 2011; van der Laan, 2008), the latter technical report paving the way to robust and more efficient estimation based on adaptive RCTs in a variety of other setting (including the case that the outcome Y is a possibly censored time-to-event).

Organization.

We first set the statistical framework in Section 2. The rationale of our adaptive covariate-adjusted designs is presented in Section 3, and we complete its formal definition in Section 4 where we also detail how the TMLE methodology operates. The asymptotic study is carried out in Section 5 (estimation) and in Section 6 (group sequential testing). The simulation study is developed in Section 7 (estimation) and in Section 8 (group sequential testing). The proofs are relegated to the Appendix.

2 Statistical framework

We tackle the asymptotic study of *adaptive group sequential designs* in the case of RCTs with covariate, binary treatment and one-dimensional primary outcome of interest. Thus, the observed data structure writes as $O = (W, A, Y)$, where $W \in \mathcal{W}$ consists of some baseline covariates, $A \in \mathcal{A} = \{0, 1\}$ denotes the assigned binary treatment, and $Y \in \mathcal{Y}$ is the primary outcome of interest. For example, Y can indicate whether the treatment has been successful or not ($\mathcal{Y} = \{0, 1\}$); or Y can count the number of times an event of interest has

occurred under the assigned treatment during a period of follow-up ($\mathcal{Y} = \mathbb{N}$) or Y can measure a quantity of interest after a given time has elapsed ($\mathcal{Y} = \mathbb{R}$). Although we will focus on the last case in this article, the methodology applies equally well to each example cited above.

Let us denote by P_0 the true distribution of the observed data structure O in the population of interest. We see P_0 as a specific element of the non-parametric set $\mathcal{M}$ of all possible observed data distributions. Note that, in order to avoid some technicalities, we assume (or rather: impose) that all elements of $\mathcal{M}$ are dominated by a common measure. The parameter of scientific interest is the marginal effect of treatment $\alpha = 1$ relative to treatment $\alpha = 0$ on the additive scale, or *risk difference*: $\psi_0 = E_{P_0}\{E_{P_0}[Y|A = 1, W] - E_{P_0}[Y|A = 0, W]\}$. Of course, other choices such as the log-relative risk (the counterpart of the risk difference on the multiplicative scale) could be considered, and dealt with along the same lines. The risk difference can be interpreted causally for instance in the counterfactual framework.

For all $P \in \mathcal{M}$, let us introduce the shorthand notation $Q_W(P)(W) = P[W]$, $g(P)(A|W) = P[A|W]$, $Q_{Y|A,W}(P)(O) = P[Y|A, W]$. We use the alternative notation $P = P_{Q,g}$ with $Q = Q(P) \equiv (Q_W(P), Q_{Y|A,W}(P))$ and $g = g(P)$. Equivalently, $P_{Q,g}$ is the data generating distribution such that $Q(P_{Q,g}) = Q$ and $g(P_{Q,g}) = g$. In particular, we denote $Q_0 = Q(P_0) = (Q_W(P_0), Q_{Y|A,W}(P_0))$. We also introduce the notation $\mathcal{Q} = \{Q(P) : P \in \mathcal{M}\}$ for the non-parametric set of all possible values of Q , and $\mathcal{G} = \{g(P) : P \in \mathcal{M}\}$ for the non-parametric set of all possible values of g . Setting $\bar{Q}(P)(A, W) = E_P[Y|A, W]$ and $\bar{Q}_0 = \bar{Q}(P_0)$ (with a slight abuse, we also write sometimes $\bar{Q}(Q)$ instead of $\bar{Q}(P_{Q,g})$), we define in greater generality

$$\Psi(P) = E_P\{\bar{Q}(P)(1, W) - \bar{Q}(P)(0, W)\} \quad (1)$$

over the whole set $\mathcal{M}$, so that ψ_0 equivalently writes as $\psi_0 = \Psi(P_0)$. This notation also emphasizes the fact that $\Psi(P)$ only depends on P through $\bar{Q}(P)$ and $Q_W(P)$, justifying the alternative notation $\Psi(P_{Q,g}) = \Psi(Q)$. The following proposition summarizes the most fundamental properties enjoyed by Ψ .

Proposition 1 (efficient influence curve). *The functional Ψ is pathwise differentiable at every $P \in \mathcal{M}$ relative to the maximal tangent space. The efficient influence curve of Ψ at $P_{Q,g} \in \mathcal{M}$ is characterized by*

$$D^\star(P_{Q,g})(O) = D_1^\star(Q)(W) + D_2^\star(P_{Q,g})(O), \quad (2)$$

where

$$\begin{aligned} D_1^\star(Q)(W) &= \bar{Q}(1, W) - \bar{Q}(0, W) - \Psi(Q), \text{ and} \\ D_2^\star(P_{Q,g})(O) &= \frac{2A-1}{g(A|W)}(Y - \bar{Q}(A, W)). \end{aligned}$$

The variance $\text{Var}_P D^\star(P)(O)$ is the lower bound of the asymptotic variance of any regular estimator of $\Psi(P)$ in the i.i.d setting. Furthermore, even if $Q \neq Q_0$,

$$E_{P_0} D^\star(P_{Q,g})(O) = 0 \text{ implies } \Psi(Q) = \Psi(Q_0) \quad (3)$$

when $g = g(P_0)$.

The implication (3) is the key to the robustness of the targeted maximum likelihood introduced and studied in this article. It is another justification of our interest in the pathwise differentiability of the functional Ψ and its efficient influence curve.

3 Data generating mechanism for adaptive design

The purpose of adaptive group sequential design as we consider it in this article is to adjust the randomization scheme as the data accrue. We first formally describe in Section 3.1 the data generating mechanism through the expression of the likelihood function and also in terms of causal graphs. Then we discuss the general issue of choosing an optimal design to target in Section 3.2, before specifically describing an optimal design of interest in Section 3.3 and showing how to target it in Section 3.4.

3.1 Data generating mechanism and related likelihood

In order to formally describe the data generating mechanism, we need to state a starting assumption: During the course of the clinical trial, it is possible to recruit independently the patients from a stationary population. In the counterfactual framework, this is equivalent to supposing that it is possible to sample as many independent copies of the full data structure as required. Let us denote by $O_i = (W_i, A_i, Y_i)$ the i th observed data structure. We also find convenient to introduce $\mathbf{O}_n = (O_1, \dots, O_n)$ and for every $i = 0, \dots, n$, $\mathbf{O}_n(i) = (O_1, \dots, O_i)$ (with convention $\mathbf{O}(0) = \emptyset$). We denote by $\mathcal{O}$ the set where the observed data structure O takes its values.

By adjusting the randomization scheme as the data accrue, we mean that the n th treatment assignment A_n is drawn from $g_n(\cdot | W_n)$, where $g_n(\cdot | W)$ is a conditional distribution (or treatment mechanism) given the covariate W which additionally depends on past observations O_{n-1} . Since the sequence of treatment mechanisms cannot reasonably grow in complexity as the sample size increases, we will only consider data-adaptive treatment mechanisms such that $g_n(\cdot | W)$ depends on $\mathbf{O}_{n-1}$ only through a finite-dimensional summary measure $Z_n = \phi_n(\mathbf{O}_{n-1})$, where the measurable function ϕ_n maps $\mathcal{O}^{n-1}$ onto $\mathbb{R}^d$ for some fixed $d \geq 0$ ($d=0$ corresponds to the case that $g_n(\cdot | W)$ actually does not adapt). For instance, $Z_{n+1} = \phi_{n+1}(\mathbf{O}_n) \equiv (n^{-1} \sum_{i=1}^n Y_i \mathbf{1}\{A_i = 0\}, n^{-1} \sum_{i=1}^n Y_i \mathbf{1}\{A_i = 1\})$ characterizes a proper summary measure of the past, which keeps track of the mean outcome in each treatment arm. Another sequence of mappings ϕ_n will be at the core of the adaptive methodology that we study in depth in this article, see (9).

Formally, the data generating mechanism is specified by the following factorization of the likelihood of $\mathbf{O}_n$:

$$\prod_{i=1}^n \left\{ Q_W^{(P_0)}(W_i) \times Q_{Y|A, W}^{(P_0)}(O_i) \right\} \times \prod_{i=1}^n g_i(A_i|W_i),$$

which suggests the introduction of $\mathbf{g}_n = (g_1, \dots, g_n)$, referred to as the *design of the study*, and the expression “ $\mathbf{O}_n$ is drawn from $(Q_0, \mathbf{g}_n)$ ”. Likewise, the likelihood of $\mathbf{O}_n$ under $(Q, \mathbf{g}_n)$ (where $Q = (Q_W, Q_{Y|A, W}) \in \mathcal{Q}$ is a candidate value for Q_0) is

$$\prod_{i=1}^n \left\{ Q_W(W_i) \times Q_{Y|A, W}(O_i) \right\} \times \prod_{i=1}^n g_i(A_i|W_i), \quad (4)$$

where we emphasize that the second factor is known. Thus we will refer with a slight abuse of terminology to $\sum_{i=1}^n \log Q_W(W_i) + \log Q_{Y|A, W}(O_i)$ as the log-likelihood of $\mathbf{O}_n$ under $(Q,$

$\mathbf{g}_n)$. Furthermore, given $\mathbf{g}_n$, we introduce the notation $P_{Q_0, \mathbf{g}_i} f \equiv E_{P_{Q_0, \mathbf{g}_i}} [f(O_i) | \mathbf{O}_n(i-1)]$

for any possibly vector-valued measurable f defined on $\mathcal{O}$.

Another equivalent characterization of the data generating mechanism involves the causal graph of Fig. 1. It is seen again firstly that W_n is drawn independently from the past $\mathbf{O}_{n-1}$; secondly that A_n is a deterministic function of W_n , the summary measure Z_n (which depends on $\mathbf{O}_{n-1}$), and a new independent source of randomness (in other words, it is drawn conditionally on (W_n, Z_n) and conditionally independently of the past $\mathbf{O}_{n-1}$); thirdly that Y_n is a deterministic function of (A_n, W_n) and a new independent source of randomness (in other words, it is drawn conditionally on (A_n, W_n) and conditionally independently of the past $\mathbf{O}_{n-1}$); then that the next summary measure Z_{n+1} is obtained as a function of $\mathbf{O}_{n-1}$ and $\mathbf{O}_n = (W_n, A_n, Y_n)$ (i.e., as a function of $\mathbf{O}_n$; here, the causal graph grants the access to a new independent source of randomness, but it is useless in our setting), and so on.

Finally, it is interesting in practice to adapt the design group sequentially. This can be simply formalized. For a given pre-specified integer $c \geq 1$ ($c = 1$ corresponds to a fully sequential adaptive design), proceeding c -group sequentially simply amounts to imposing

$\phi_{(r-1)c+1}(\mathbf{O}_{(r-1)c}) = \dots = \phi_{rc}(\mathbf{O}_{rc-1})$ for all $r \geq 1$. Then the c treatment assignments $A_{(r-1)c+1}, \dots, A_{rc}$ in the r th c -group are all drawn from the same conditional distribution $g_{(r-1)c}(\cdot | W)$.

Yet, although all our results (and their proofs) still hold for any $c \geq 1$, we prefer to consider in the rest of this section and in Sections 4 and 5 the case that $c = 1$ for simplicity's sake. On the contrary, the simulation study carried out in Section 7 involves some $c > 1$.

3.2 On the user-supplied design to target

One of the most important features of the adaptive group sequential design methodology is that it *targets* a *user-supplied* specific design of special interest. This specific design is

generally an optimal design with respect to a criterion which translates what the investigator cares for the most. Specifically, one could care the most for the well-being of the target population, wishing that a result be available as quickly as possible and aspiring therefore to the highest efficiency (*i.e.*, the ability to reach a conclusion as quickly as possible subject to level and power constraints). Or one could care the most for the well-being of the subjects participating to the clinical trial (therefore trying to minimize the number of patients assigned to their corresponding inferior treatment arms, subject to level and power constraints). Obviously, these are only two important examples from a large class of potentially interesting criteria. The sole purpose of the criterion is to generate a random element in $\mathcal{G}$ of the form $g_n = gZ_n$ where $Z_n = \phi_n(\mathbf{O}_{n-1})$ is a finite-dimensional summary measure of $\mathbf{O}_{n-1}$.

We decide to focus in this article on the first example, but it must be clear that the methodology applies to a variety of other criteria (see (van der Laan, 2008) for other examples).

3.3 From the Neyman allocation to the optimal design

So, the objective is now clearly identified: we wish to adapt the design as the data accrue in order to learn from the data then mimic a specific design which guarantees the highest efficiency, our (user-supplied) *optimal design*.

By Proposition 1, the asymptotic variance of any regular estimator of the risk difference $\Psi(Q_0)$ has lower bound $\text{Var}_{P_{Q_0, g}} D^\star(P_{Q_0, g})(O)$ if the estimator relies on data sampled independently from $P_{Q_0, g}$. Now,

$$\text{Var}_{P_{Q_0, g}} D^\star(P_{Q_0, g})(O) = E_{Q_0} (\bar{Q}_0(1, W) - \bar{Q}_0(0, W) - \Psi(Q_0))^2 + E_{Q_0} \left(\frac{\sigma^2(Q_0)(1, W)}{g(1|W)} + \frac{\sigma^2(Q_0)(0, W)}{g(0|W)} \right),$$

where $\sigma^2(Q_0)(A, W)$ denotes the conditional variance of Y given (A, W) under Q_0 . We use the notation E_{Q_0} above (for the expectation with respect to the marginal distribution of W under P_0) in order to emphasize the fact that the treatment mechanism g only appears in the second term of the right-hand side sum. Furthermore, it holds P_0 -almost surely that

$$\frac{\sigma^2(Q_0)(1, W)}{g(1|W)} + \frac{\sigma^2(Q_0)(0, W)}{g(0|W)} \geq (\sigma(Q_0)(1, W) + \sigma(Q_0)(0, W))^2,$$

with equality if and only if

$$g(1|W) = \frac{\sigma(Q_0)(1, W)}{\sigma(Q_0)(1, W) + \sigma(Q_0)(0, W)} \quad (5)$$

P_0 -almost surely. Therefore, the following lower bound holds for all $g \in \mathcal{G}$:

$$\text{Var}_{P_{Q_0, g}} D^{\star}(P_{Q_0, g})(O) \geq E_{Q_0} (\bar{Q}_0(1, W) - \bar{Q}_0(0, W) \Psi(Q_0))^2 + E_{Q_0} (\sigma(Q_0)(1, W) + \sigma(Q_0)(0, W))^2,$$

with equality if and only if $g \in \mathcal{G}$ is characterized by (5). This optimal design is known in the literature as the *Neyman allocation* (see (Hu and Rosenberger, 2006) page 13). This result notably makes clear that the most efficient treatment mechanism assigns with higher probability a patient with covariate vector W to the treatment arm such that the variance of the outcome Y in this arm is the largest, regardless of the mean of the outcome (*i.e.*, whether the arm is inferior or superior).

Due to logistical reasons, it might be preferable to consider only treatment mechanisms that assign treatment in response to a subvector V of the baseline covariate vector W . In addition, if W is complex, targeting the optimal Neyman allocation might be too ambitious.

Therefore, we will consider the important case where V is a discrete covariate with finitely many values in the set $\mathcal{V} = \{1, \dots, v\}$. The covariate V indicates subgroup membership for a collection of v subgroups of interest. We decide to restrict the search of an optimal design to the set $\mathcal{G}_1 \subset \mathcal{G}$ of those treatment mechanisms which only depend on W through V . The same calculations as above yield straightforwardly that, for all $g \in \mathcal{G}_1$,

$$\text{Var}_{P_{Q_0, g}} D^{\star}(P_{Q_0, g})(O) \geq E_{Q_0} (\bar{Q}_0(1, W) - \bar{Q}_0(0, W) - \Psi(Q_0))^2 + E_{Q_0} (\bar{\sigma}(Q_0)(1, V) + \bar{\sigma}(Q_0)(0, V))^2,$$

where $\bar{\sigma}^2(Q_0)(a, V) = E_{Q_0} [\sigma^2(Q_0)(a, W) | V]$ for $a \in \mathcal{A}$ with equality if and only if g coincides with $g^{\star}(Q_0)$, characterized by

$$g^{\star}(Q_0)(1 | V) = \frac{\bar{\sigma}(Q_0)(1, V)}{\bar{\sigma}(Q_0)(1, V) + \bar{\sigma}(Q_0)(0, V)} \quad (6)$$

P_0 -almost surely. Hereafter, we refer to $g^{\star}(Q_0)$ as *the optimal design*.

3.4 Targeting the optimal design

Because $g^{\star}(Q_0)$ is characterized as the minimizer over $g \in \mathcal{G}_1$ of the variance under $P_{Q_0, g}$ of the efficient influence curve at $P_{Q_0, g}$ we propose to construct $g_{n+1} \in \mathcal{G}_1$ as the minimizer over $g \in \mathcal{G}_1$ of an estimator of the latter variance based on past observations $\mathbf{O}_n$.

We proceed by recursion. We first set $g_1 = g^b$, the so-called balanced treatment mechanism such that $g^b(1 | W) = \frac{1}{2}$ for all $W \in \mathcal{W}$, and assume that $\mathbf{O}_n$ has already been sampled from (Q_0, g_n) as described in Section 3.1, the sample size being large enough to guarantee

$\sum_{i=1}^n \mathbf{1}\{V_i = v\} > 0$ for all $v \in \mathcal{V}$ (if n_0 is the smallest sample size such that the previous condition is met, then we set $g_1 = \dots = g_{n_0} = g^b$).

The issue is now to construct g_{n+1} . Let us assume for the time being that we already know how to construct an estimator Q_n of Q_0 based on $\mathbf{O}_n$ (hence the estimators $\bar{Q}_n = \bar{Q}(Q_n)$ of $\bar{Q}_0$ and $\Psi(Q_n)$ of $\Psi(Q_0) = \Psi_0$).¹ Then, for all $g \in \mathcal{G}_1$,

$$\begin{aligned} S_n(g) &= \frac{1}{n} \sum_{i=1}^n \left\{ D_1^\star(Q_n)(W_i)^2 + 2D_1^\star(Q_n)(W_i)D_2^\star(P_{Q_n, g})(O_i) \frac{g(A_i|V_i)}{g_i(A_i|V_i)} \right. \\ &\quad \left. + D_2^\star(P_{Q_n, g})(O_i)^2 \frac{g(A_i|V_i)}{g_i(A_i|V_i)} \right\} - \left(\frac{1}{n} \sum_{i=1}^n D_1^\star(Q_n)(W_i) + D_2^\star(P_{Q_n, g})(O_i) \frac{g(A_i|V_i)}{g_i(A_i|V_i)} \right)^2 \\ &= \frac{1}{n} \sum_{i=1}^n \frac{(Y_i - \bar{Q}_n(A_i, W_i))^2}{g(A_i|V_i)g_i(A_i|V_i)} \\ &\quad + \left\{ \frac{1}{n} \sum_{i=1}^n D_1^\star(Q_n)(W_i)^2 + 2D_1^\star(Q_n)(W_i)D_2^\star(P_{Q_n, g})(O_i) \frac{g(A_i|V_i)}{g_i(A_i|V_i)} \right. \\ &\quad \left. - \left(\frac{1}{n} \sum_{i=1}^n D_1^\star(Q_n)(W_i) + D_2^\star(P_{Q_n, g})(O_i) \frac{g(A_i|V_i)}{g_i(A_i|V_i)} \right)^2 \right\} \end{aligned}$$

estimates $\text{Var}_{P_{Q_0, g}} D^\star(O; P_{Q_0, g})$ (the weighting provides the adequate tilt of the empirical distribution; it is not necessary to weight the terms corresponding to $D_1^\star$ because they do not depend on the treatment mechanism). Now, only the first term in the rightmost expression still depends on g . The same calculations as above straightforwardly yields that $S_n(g)$ is minimized at $g_{n+1}^\star \in \mathcal{G}_1$ characterized by

$$g_{n+1}(1|v) = \frac{s_{v,n}(1)}{s_{v,n}(1) + s_{v,n}(0)}, \quad (7)$$

for all $v \in \mathcal{V}$, where for each $(v, a) \in \mathcal{V} \times \mathcal{A}$,

$$s_{v,n}^2(a) = \frac{\frac{1}{n} \sum_{i=1}^n \frac{(Y_i - \bar{Q}_n(A_i, W_i))^2}{g_i(A_i|V_i)} \mathbf{1}\{(V_i, A_i) = (v, a)\}}{\frac{1}{n} \sum_{i=1}^n \mathbf{1}\{V_i = v\}}.$$

Yet, instead of considering the above characterization, we find it more convenient to define

¹ The reasoning is not circular by virtue of the chronological ordering as it is summarized in Fig. 1 for instance.

$$g_{n+1}^{\star}(1|v) = \frac{\sigma_{v,n}(1)}{\sigma_{v,n}(1) + \sigma_{v,n}(0)}, \quad (8)$$

for all $v \in \mathcal{V}$, where for each $(v, a) \in \mathcal{V} \times \mathcal{A}$,

$$\sigma_{v,n}^2(a) = \frac{\frac{1}{n} \sum_{i=1}^n \frac{(Y_i - \bar{Q}_n(A_i, W_i))^2}{g_i(A_i|V_i)} \mathbf{1}\{(V_i, A_i) = (v, a)\}}{\frac{1}{n} \sum_{i=1}^n \frac{\mathbf{1}\{(V_i, A_i) = (v, a)\}}{g_i(A_i|V_i)}}.$$

Note that $s_{v,n}^2(a)$ and $\sigma_{v,n}^2(a)$ share the same numerator, and that the different denominators converge to the same limit. Substituting $\sigma_{v,n}^2(a)$ for $s_{v,n}^2(a)$ is convenient because one naturally interprets the former as an estimator of the conditional variance of Y given $(A, V) = (a, v)$ based on $\mathbf{O}_n$, a fact that we use in Section 4.2. Finally, we emphasize that $g_{n+1}^{\star} = g_{n+1}^{z_{n+1}}$ for the summary measure of the past $\mathbf{O}_n$

$$Z_{n+1} = \phi_{n+1}(\mathbf{O}_n) \equiv ((\sigma_{v,n}^2(0), \sigma_{v,n}^2(1)) : v \in \mathcal{V}). \quad (9)$$

The rigorous definition of the design $\mathbf{g}_n^{\star} = (g_1^{\star}, \dots, g_n^{\star})$ follows by recursion, but it is still subject to knowledge about how to construct an estimator Q_n of Q_0 based on $\mathbf{O}_n$. Because this last missing piece of the formal definition of the adaptive group sequential design data generating mechanism is also the core of the TMLE procedure, we address it in Section 4.

4 TMLE in adaptive covariate-adjusted RCTs

We assume hereafter that $\mathbf{O}_n$ has already been sampled from the $(Q_0, g_n^{\star})$ -adaptive sampling scheme. In this section, we construct an estimator Q_n (actually denoted by Q_n^*) of Q_0 , therefore yielding the characterization of $g_{n+1}^{\star}$, and completing the formal definition of the adaptive design $\mathbf{g}_n^{\star}$ that we initiated in Section 3. In particular, the next observed data structure $\mathbf{O}_{n+1}$ can be drawn from $(Q_0, g_{n+1}^{\star})$, and it makes sense to undertake the asymptotic study of the properties of the TMLE methodology based on adaptive group sequential sampling.

As in the i.i.d framework, the TMLE procedure maps an initial substitution estimator $\Psi(Q_n^0)$ of ψ_0 into an update $\psi_n^* = \Psi(Q_n^*)$ by fluctuating the initial estimate Q_n^0 of Q_0 . The construction of Q_n^0 is presented and studied in Section 4.1. In Section 4.2, a straightforward

application of the main result of Section 4.1 shows that the adaptive design converges. How to fluctuate Q_n^0 and stretch it optimally to Q_n^* is presented and studied in Section 4.3.

4.1 Initial maximum likelihood based substitution estimator Working model.

In order to construct the initial estimate Q_n^0 of Q_0 , we consider a *working model* Q_n^w . With a slight abuse of notation, the elements of Q_n^w are denoted by $(Q_W(P_n), Q_{Y|A, W}(\theta))$ for some parameter $\theta \in \Theta$, where $Q_W(P_n)$ is the empirical marginal distribution of W . Specifically, the working model Q_n^w is chosen in such a way that

$$Q_{Y|A, W}(\theta)(O) = \frac{1}{\sqrt{2\pi\sigma_V^2(A)}} \exp\left[-\frac{(Y - m(A, W; \beta_V))^2}{2\sigma_V^2(A)}\right].$$

This notably implies that for any $P_\theta \in \mathcal{M}$ such that $Q_{Y|A, W}(P_\theta) = Q_{Y|A, W}(\theta)$, the conditional mean $\bar{Q}(P_\theta)(A, W)$, which we also denote by $\bar{Q}(\theta)$, satisfies $\bar{Q}(\theta)(A, W) = m(A, W; \beta_V)$, the right-hand side expression being a linear combination of variables extracted from (A, W) and indexed by the regression vector β_V (of dimension b). Defining

$$\theta(v) = (\beta_v, \sigma_v^2(0), \sigma_v^2(1))^T \in \Theta_v \subset \mathbb{R}^b \times \mathbb{R}_+^* \times \mathbb{R}_+^* \quad (10)$$

for each $v \in \mathcal{V}$, the complete parameter is given by $\theta = (\theta(1)^T, \dots, \theta(v)^T)^T \in \Theta$, where $\theta = \prod_{v=1}^V \theta_v$. We impose the following condition on the parameterization:

PARAM. The parameter set Θ is compact. Furthermore, the linear parameterization is identifiable: for all $v \in \mathcal{V}$, if $m(a, w; \beta_v) = m(a, w; \beta'_v)$ for all $a \in \mathcal{A}$ and $w \in \mathcal{W}$ (compatible with v), then necessarily $\beta_v = \beta'_v$.

Characterizing Q_n^0 .—Let us set a reference fixed design $g^r \in \mathcal{G}_1$. We now characterize Q_n^0 by letting

$$Q_n^0 = (Q_W(P_n), Q_{Y|A, W}(\theta_n)), \quad (11)$$

Where

$$\theta_n = \arg \max_{\theta \in \Theta} \sum_{i=1}^n \log Q_{Y|A, W}(\theta)(O_i) \frac{g^r(A_i|V_i)}{g_i^\star(A_i|V_i)} \quad (12)$$

is a *weighted* maximum likelihood estimator with respect to the working model. Thus, the v th component $\theta_n(v)$ of θ_n satisfies

$$\theta_n(v) = \arg \min_{\theta(v) \in \Theta_{v,i}} \sum_{i=1}^n \left(\log \sigma_v^2(A_i) + \frac{(Y_i - m(A_i, W_i; \beta_v))^2}{\sigma_v^2(A_i)} \right) \frac{g^r(A_i|V_i)}{g_i^\star(A_i|V_i)} \mathbf{1}\{V_i = v\}$$

for every $v \in \mathcal{V}$. Note that this initial estimate Q_n^0 of Q_0 yields the initial maximum likelihood based substitution estimator $\Psi(Q_n^0)$ of ψ_0 :

$$\Psi(Q_n^0) = \frac{1}{n} \sum_{i=1}^n \bar{Q}(\theta_n^0)(1, W_i) - \bar{Q}(\theta_n^0)(0, W_i).$$

Studying Q_n^0 through θ_n .—For simplicity, let us introduce, for all $\theta \in \Theta$, the additional notations

$$\ell_{\theta,0} = \log Q_{Y|A,W}(\theta), \quad \dot{\ell}_{\theta,0} = \frac{\partial}{\partial \theta} \ell_{\theta,0}, \quad \ddot{\ell}_{\theta,0} = \frac{\partial^2}{\partial \theta^2} \ell_{\theta,0}.$$

The first asymptotic property of θ_n that we derive concerns its consistency (see Theorem 5 in (van der Laan, 2008)).

Proposition 2 (consistency of θ_n). *Assume that:*

A1. *There exists a unique interior point $\theta_0 \in \Theta$ such that $\theta_0 = \arg \max_{\theta \in \Theta} P_{Q_0, g^r} \ell_{\theta,0}$.*

A2. *The matrix $-P_{Q_0, g^r} \ddot{\ell}_{\theta_0,0}$ is positive definite.*

Provided that $\mathcal{O}$ is a bounded set, θ_n consistently estimates θ_0 .

The proof of Proposition 2 is given in Section A.1.

The limit in probability of θ_n has a nice interpretation in terms of projection of $Q_{Y|A,W}(P_0)$ onto $\{Q_{Y|A,W}(\theta): \theta \in \Theta\}$. Preferring to discuss this issue in terms of data generating distribution rather than conditional distribution, let us set $Q_{\theta_0} = (Q_W(P_0), Q_{Y|A,W}(\theta_0))$ and

assume that $P_{Q_0, g^r} \log Q_{Y|Q,W}(P_0)$ is well defined (this weak assumption concerns Q_0 , not

g^r , and holds for instance when $|\log Q_{Y|A,W}(P_0)|$ is bounded). Then **A1** is equivalent to $P_{Q_{\theta_0}, g^r}$ being the unique Kullback-Leibler projection of P_{Q_0, g^r} onto the set

$$\left\{ P \in \mathcal{M}: \exists \theta \in \Theta \text{ s.t. } Q_{Y|A,W}(P) = Q_{Y|A,W}(\theta), \text{ and } Q_W(P) = Q_W(P_0), g(P) = g^r \right\}.$$

In addition to being consistent, θ_n actually satisfies a central limit theorem if supplementary mild conditions are met. The latter central limit theorem is embedded in a more general result that we state in Section 4.3, see Proposition 5.

Furthermore, maximizing a weighted version of the log-likelihood is a technical twist that makes the theoretical study of the properties of θ_n easier. Indeed, the unweighted maximum likelihood estimator

$$t_n = \arg \max_{\theta \in \Theta} \sum_{i=1}^n \log Q_{Y|A, W}(\theta)(O_i)$$

targets the parameter

$$T_{\bar{g}_n}(Q_0) = \arg \max_{\theta \in \Theta} \sum_{i=1}^n P_{Q_0, \bar{g}_n} \star \log Q_{Y|A, W}(\theta) = \arg \max_{\theta \in \Theta} P_{Q_0, \bar{g}_n} \ell_{\theta, 0},$$

where $\bar{g}_n = n^{-1} \sum_{i=1}^n g_i \star$ and $P_{Q_0, \bar{g}_n} f \equiv E_{P_{Q_0, \bar{g}_n}} [f(O_n) | \mathbf{O}_n(n-1)]$ for any measurable f defined on $\mathcal{O}$. Therefore, t_n asymptotically targets the limit, if it exists, of $T_{\bar{g}_n}(Q_0)$. Assuming that $\bar{g}_n$ converges itself to a fixed design $g_\infty \in \mathcal{G}$, then t_n asymptotically targets parameter $T_{g_\infty}(Q_0)$. The latter parameter is very difficult to interpret and to analyze, as it depends directly and indirectly (through g_∞) on Q_0 .

4.2 Convergence of the adaptive design

Consider the mapping $G^\star$ from Θ to $\mathcal{G}_1$ (respectively equipped with the Euclidean distance and, for instance, the distance $d(g, g') = \sum_{v \in \mathcal{V}} |g(1|v) - g'(1|v)|$) such that, for any $\theta \in \Theta$, for any $(a, v) \in \mathcal{A} \times \mathcal{V}$,

$$G^\star(\theta)(a|v) = \frac{\sigma_v(a)}{\sigma_v(1) + \sigma_v(0)}. \quad (13)$$

Equation (13) characterizes $G^\star$ which is obviously continuous. Since $\mathbf{g}_n^\star$ is adapted in such a way that $g_n^\star = G^\star(\theta_n)$ (see (8)), Proposition 2 and the continuous mapping theorem (see Theorem 1.3.6 in (van der Vaart and Wellner, 1996)) straightforwardly imply the following result (the convergence in probability yields the convergence in L^1 because $\mathbf{g}_n^\star$ is uniformly bounded).

Proposition 3 (convergence of $g_n^\star$). *Under the assumptions of Proposition 2, the adaptive design $g_n^\star$ converges in probability and in L^1 to the limit design $G^\star(\theta_0)$.*

The convergence of the adaptive design $\mathbf{g}_n^\star$ is a crucial result. It is noteworthy that the limit design $G^\star(\theta_0)$ equals the optimal design $g^\star(Q_0)$ if the working model is correctly specified (which never happens in real-life applications), but not necessarily otherwise. Furthermore, the relationship $g_n^\star = G^\star(\theta_n)$ also entails the possibility to derive the convergence in distribution of $\sqrt{n}(g_n^\star - G^\star(\theta_n))$ to a centered Gaussian distribution with known variance by application of the delta-method ($G^\star$ is differentiable) from a central limit theorem on θ_n (see Proposition 5 and Theorem 3.1 in (van der Vaart, 1998)).

4.3 TMLE

Fluctuating Q_n^0 .—The second step of the TMLE procedure stretches the initial estimate $\Psi(Q_n^0)$ in the direction of the parameter of interest, through a maximum likelihood step over a well-chosen fluctuation of Q_n^0 . The latter fluctuation of Q_n^0 is just a one-dimensional parametric model $\{Q_n^0(\varepsilon): \varepsilon \in \mathcal{E}\} \subset \mathcal{Q}$ indexed by the parameter $\varepsilon \in \mathcal{E}$, $\mathcal{E} \subset \mathbb{R}$ being a bounded interval which contains a neighborhood of the origin. Specifically, we set for all $\varepsilon \in \mathcal{E}$:

$$Q_n^0(\varepsilon) = (Q_W(P_n), Q_{Y|A, W}(\theta_n, \varepsilon)), \quad (14)$$

where for any $\theta \in \Theta$,

$$Q_{Y|A, W}(\theta, \varepsilon)(O) = \frac{1}{\sqrt{2\pi\sigma_V^2(A)}} \exp\left\{-\frac{(Y - \bar{Q}(\theta)(A, W) - \varepsilon H^\star(\theta)(A, W))^2}{2\sigma_V^2(A)}\right\}, \quad (15)$$

with

$$H^\star(\theta)(A, W) = \frac{2A - 1}{G^\star(\theta)(A|V)} \sigma_V^2(A).$$

In particular, the fluctuation goes through Q_n^0 at $\varepsilon = 0$ (i.e., $Q_n^0(0) = Q_n^0$). Let $P_n^0(\varepsilon) \in \mathcal{M}$ be a data generating distribution such that $Q_{Y|A, W}(P_n^0(\varepsilon)) = Q_{Y|A, W}(\theta_n, \varepsilon)$. The conditional mean $\bar{Q}(P_n^0(\varepsilon))$, which we also denote by $\bar{Q}(\theta_n, \varepsilon)$, is

$$\bar{Q}(\theta_n, \varepsilon)(A, W) = \bar{Q}(\theta_n)(A, W) + \varepsilon H^\star(\theta_n)(A, W).$$

Furthermore, the score at $\varepsilon = 0$ of $P_n^0(\varepsilon)$ equals

$$\frac{\partial}{\partial \varepsilon} \log P_n^0(\varepsilon)[O] \Big|_{\varepsilon=0} = \frac{2A - 1}{G^\star(\theta_n)(A|V)} (Y - \bar{Q}(\theta_n)(A, W)) = D_2^\star(P_{Q_n^0, G^\star(\theta_n)})(O),$$

the second component of the efficient influence curve of Ψ at $P_{Q_n^0, G^\star(\theta_n)} = P_{Q_n^0, g_n^\star}$, see (2)

in Proposition 1 (recall that $g_n^\star = G^\star(\theta_n)$).

Characterizing Q_n^* yields the TMLE.—We characterize the update Q_n^* of Q_n^0 in the fluctuation $\{Q_n^0(\varepsilon): \varepsilon \in \mathcal{E}\}$ by

$$Q_n^* = Q_n^0(\varepsilon_n),$$

where

$$\varepsilon_n = \arg \max_{\varepsilon \in \mathcal{E}} \sum_{i=1}^n \log Q_{Y|A, W}(\theta_n, \varepsilon)(O_i) \frac{g_n^\star(A_i|V_i)}{g_i^\star(A_i|V_i)} \quad (16)$$

is a *weighted* maximum likelihood estimator with respect to the fluctuation. It is worth noting that ε_n is known in closed form (we assume, without serious loss of generality, that $\mathcal{E}$ is large enough for the maximum to be achieved in its interior). Denoting the v th component $\theta_n(v)$ of θ_n by $(\beta_{v,n}, \sigma_{v,n}^2(0), \sigma_{v,n}^2(1))$, it holds that

$$\varepsilon_n = \frac{\sum_{i=1}^n (Y_i - \bar{Q}(\theta_n)(A_i, W_i)) \frac{2A_i - 1}{g_i^\star(A_i|V_i)}}{\sum_{i=1}^n \frac{\sigma_{V_i, n}^2(A_i)}{g_n^\star(A_i|V_i)g_i(A_i|V_i)}}.$$

The notation Q_n^* for this first update of Q_n^0 is a reference to the fact that the TMLE procedure, which is in greater generality an iterative procedure, converges here in one single step. Indeed, say that one fluctuates Q_n^* as we fluctuate Q_n^0 *i.e.*, by introducing

$$Q_n^1(\varepsilon) = (Q_{W(P_n)}, Q_{Y|A, W}(\theta_n, \varepsilon_n, \varepsilon))$$

with $Q_{Y|A, W}(\theta, \varepsilon', \varepsilon)$ equal to the right-hand side of (15) where one substitutes $\bar{Q}(\theta, \varepsilon')$ for $\bar{Q}(\theta)$. Say that one then defines the weighted maximum likelihood ε'_n as the right-hand side of (16) where one substitutes $Q_{Y|A, W}(\theta_n, \varepsilon_n, \varepsilon)$ for $Q_{Y|A, W}(\theta_n, \varepsilon)$. Then it follows that $\varepsilon'_n = 0$ so that the “updated” $Q_n^*(\varepsilon'_n) = Q_n^*$. The updated estimator Q_n^* of Q_0 maps into the TMLE $\psi_n^* = \Psi(Q_n^*)$ of the risk difference $\psi_0 = \Psi(Q_0)$:

$$\psi_n^* = \frac{1}{n} \sum_{i=1}^n \bar{Q}(\theta_n, \varepsilon_n)(1, W_i) - \bar{Q}(\theta_n, \varepsilon_n)(0, W_i). \quad (17)$$

The asymptotic study of ψ_n^* relies on a central limit theorem for $(\theta_n, \varepsilon_n)$, which we discuss in Section 5.3.

5 Asymptotics

5.1 Studying Q_n^* through $(\theta_n, \varepsilon_n)$: consistency

We now state and comment on a consistency result for the stacked estimator $(\theta_n, \varepsilon_n)$ which complements Proposition 2 (see Theorem 8 in (van der Laan, 2008)). For simplicity, let us generalize the notation ℓ_{Q_0} introduced in Section 4.1 by setting, for all $(\theta, \varepsilon) \in \Theta \times \mathcal{E}$,

$$\ell_{\theta, \varepsilon} = \log Q_{Y|A, W}(\theta, \varepsilon).$$

Moreover, let us set, for all $(\theta, \varepsilon) \in \Theta \times \mathcal{E}$,

$$Q_{\theta, \varepsilon} = (Q_W(P_0), Q_{Y|A, W}(\theta, \varepsilon)). \quad (18)$$

Proposition 4 (consistency of $(\theta_n, \varepsilon_n)$). *Suppose that **A1** and **A2** from Proposition 2 hold. In addition, assume that:*

A3. *There exists a unique interior point $\varepsilon_0 \in \mathcal{E}$ such that $\varepsilon_0 = \arg \max_{\varepsilon \in \mathcal{E}} P_{Q_0, G^\star(\theta_0)} \ell_{\theta_0, \varepsilon}$.*

(i) *It holds that $\Psi(Q_{\theta_0, \varepsilon_0}) = \Psi(Q_0)$.*

(ii) *Provided that $\mathcal{O}$ is a bounded set, $(\theta_n, \varepsilon_n)$ consistently estimates $(\theta_0, \varepsilon_0)$.*

The proof of Proposition 4 is given in Section A.1.

We already discussed the interpretation of the limit in probability of θ_n in terms of Kullback-Leibler projection. Likewise, the limit in probability ε_0 of ε_n enjoys such an interpretation.

Let us assume that $P_{Q_0, G^\star(\theta_0)} \log Q_{Y|A, W}(P_0)$ is well defined (this weak assumption

concerns Q_0 , not $G^\star(\theta_0)$, and holds for instance when $|\log Q_{Y|A, W}(P_0)|$ is bounded). Then

A3 is equivalent to $P_{Q_{\theta_0, \varepsilon_0}, G^\star(\theta_0)}$ being the unique the Kullback-Leibler projection of

$P_{Q_0, G^\star(\theta_0)}$ onto the set

$$\left\{ P \in M: \exists \varepsilon \in \mathcal{E} \text{ s.t. } Q(P) = Q_{\theta_0, \varepsilon} \text{ and } g(P) = G^\star(\theta_0) \right\}.$$

Of course, the most striking property that Q_0 enjoys is (i): even if $Q_0 \notin \{\bar{Q}(\theta, \varepsilon) : (\theta, \varepsilon) \in \Theta \times \mathcal{E}\}$, it holds that $\Psi(Q_{\theta_0, \varepsilon_0}) = \Psi(Q_0)$. This remarkable equality and the convergence of $(\theta_n, \varepsilon_n)$ to $(\theta_0, \varepsilon_0)$ are evidently the keys to the consistency of $\psi_n^* = \Psi(Q_n^*)$. We investigate in Section 5.3 how the consistency result stated in Proposition 4 translates into the consistency of the TMLE.

5.2 Studying Q_n^* through $(\theta_n, \varepsilon_n)$: central limit theorem

We now state and comment on a central limit theorem for the stacked estimator $(\theta_n, \varepsilon_n)$ (see also Theorem 9 in (van der Laan, 2008)).

Let us introduce, for all $(\theta, \varepsilon) \in \Theta \times \mathcal{E}$,

$$D(\theta, \varepsilon)(O, Z) = \frac{1}{g_Z(A|V)} \left(\dot{\ell}_{\theta, 0}^T(O) g^r(A|V), \frac{\partial}{\partial \varepsilon} \ell_{\theta, \varepsilon}(O) G^\star(\theta)(A|V) \right)^\top, \quad (19)$$

and $\tilde{D}(\theta, \varepsilon)(O) = g_Z(A|V) D(\theta, \varepsilon)(O, Z)$.

Proposition 5 (central limit theorem for $(\theta_n, \varepsilon_n)$). *Suppose that **A1**, **A2** and **A3** from Propositions 2 and 4 hold. In addition, assume that **A4**. If a deterministic function F is such that $F(O) = 0$ $P_{Q_0, G^\star(\theta_0)}$ -almost surely, then $F = 0$. Then the following asymptotic linear expansion holds:*

$$\sqrt{n}((\theta_n, \varepsilon_n) - (\theta_0, \varepsilon_0)) = S_0^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n D(\theta_0, \varepsilon_0)(O_i, Z_i) + o_p(1), \quad (20)$$

where

$$S_0 = E_{Q_0, G^\star(\theta_0)}$$

$$\begin{pmatrix} \ddot{\ell}_{\theta_0, 0}(O) \frac{g^r(A|V)}{G^\star(\theta_0)(A|V)} & 0 \\ \left[\frac{\partial^2}{\partial \theta \partial \varepsilon} \ell_{\theta, \varepsilon}(O) G^\star(\theta)(A|V) \right]^\top \Big|_{(\theta, \varepsilon) = (\theta_0, \varepsilon_0)} & \frac{1}{G^\star(\theta_0)(A|V)} \frac{\partial^2}{\partial \varepsilon^2} \ell_{\theta_0, \varepsilon}(O) \Big|_{\varepsilon = \varepsilon_0} \end{pmatrix}$$

(21)

is an invertible matrix. Furthermore, (20) entails that $\sqrt{n}((\theta_n, \varepsilon_n) - (\theta_0, \varepsilon_0))$ converges in distribution to the centered Gaussian distribution with covariance matrix $S_0^{-1} \Sigma_0 (S_0^{-1})^\top$, where

$$\Sigma_0 = E_{Q_0, G^\star(\theta_0)} \left(\frac{\tilde{D}(\theta_0, \varepsilon_0) \tilde{D}(\theta_0, \varepsilon_0)^\top (O)}{G^\star(\theta_0)(A|V)^2} \right) \quad (22)$$

is a positive definite symmetric matrix. Moreover, S_0 is consistently estimated by

$$S_n = \frac{1}{n} \sum_{i=1}^n \begin{pmatrix} \ddot{\ell}_{\theta_n, 0}(O_i) \frac{g^r(A_i|V_i)}{gz_i(A_i|V_i)} & 0 \\ \left[\frac{\partial^2}{\partial \theta \partial \varepsilon} \ell_{\theta, \varepsilon}(O_i) G^\star(\theta)(A_i|V_i) \right]^\top \Big|_{(\theta, \varepsilon) = (\theta_n, \varepsilon_n)} \frac{1}{gz_i(A_i|V_i)} & \frac{\partial^2}{\partial \varepsilon^2} \ell_{\theta_n, \varepsilon}(O_i) \Big|_{\varepsilon = \varepsilon_n} \frac{gz_i(A_i|V_i)}{gz_i(A_i|V_i)} \end{pmatrix}$$

and Σ_0 is consistently estimated by

$$\Sigma_n = \frac{1}{n} \sum_{i=1}^n D(\theta_n, \varepsilon_n) D(\theta_n, \varepsilon_n)^\top (O_i, Z_i). \quad (23)$$

The proof of Proposition 5 is given in Section A.2. We investigate in Section 5.3 how the above central limit theorem translates into a central limit theorem for the TMLE.

5.3 Consistency and asymptotic normality of the TMLE

In this section, we finally state and comment on the asymptotic properties of the TMLE ψ_n^* .

TMLE is consistent and asymptotically Gaussian.—In the first place, the TMLE ψ_n^* is robust: it is a consistent estimator even when the working model is misspecified (which is always the case in real-life applications).

Proposition 6 (consistency of ψ_n^*). *Suppose that A1, A2, A3 and A4 from Propositions 2, 4 and 5 hold. Then the TMLE ψ_n^* consistently estimates the risk difference ψ_0 .*

If the design of the clinical trial was fixed (and consequently, the n first observations were i.i.d), then the TMLE would be a robust estimator of ψ_0 : Even if the working model is misspecified, then the TMLE still consistently estimates ψ_0 because the treatment mechanism is known (or can be consistently estimated, if one wants to gain in efficiency).

Thus, the robustness of the TMLE stated in Proposition 6 is the *expected* counterpart of the TMLE's robustness in the latter i.i.d setting: *Expected* because the TMLE solves a martingale estimating function that is unbiased for ψ_0 at misspecified Q and correctly specified g_i $i = 1, \dots, n$.

In the second place, the TMLE ψ_n^* is asymptotically linear, and therefore satisfies a central limit theorem. To see this, let us introduce the real-valued function ϕ on $\Theta \times \mathcal{E}$ such that $\phi(\theta, \epsilon) = \Psi(Q_{\theta, \epsilon})$ (see (18) for the definition of $Q_{\theta, \epsilon}$). The function ϕ is differentiable on the interior of $\Theta \times \mathcal{E}$, and we denote $\phi'_{\theta, \epsilon}$ its gradient at (θ, ϵ) . The latter gradient satisfies

$$\phi'_{\theta, \epsilon} = E_{Q_{\theta, \epsilon}, G^{\star}(\theta)} \left\{ D^{\star}(P_{Q_{\theta, \epsilon}}, G^{\star}(\theta))(O) \left(\frac{\partial}{\partial \theta} \ell_{\theta, \epsilon}^{\top}(O), \frac{\partial}{\partial \epsilon} \ell_{\theta, \epsilon}(O) \right)^{\top} \right\}. \quad (24)$$

Note that the right-hand side expression cannot be computed explicitly because the marginal distribution $Q_W(P_0)$ is unknown. By the law of large numbers (independent case), we can build an estimator ϕ'_n of $\phi'_{\theta_0, \epsilon_0}$ as follows. For B a large number (say $B = 10^4$), simulate B independent copies $\tilde{O}_b$ of O from the data generating distribution $P_{Q_n^*, G^{\star}(\theta_n)}$, then compute

$$\phi'_n = \frac{1}{B} \sum_{b=1}^B D^{\star} \left(P_{Q_n^*, G^{\star}(\theta_n)} \right) (\tilde{O}_b) \left(\frac{\partial}{\partial \theta} \ell_{\theta, \epsilon_n}^{\top}(\tilde{O}_b) \Big|_{\theta=\theta_n}, \frac{\partial}{\partial \epsilon} \ell_{\theta_n, \epsilon}(\tilde{O}_b) \Big|_{\epsilon=\epsilon_n} \right)^{\top}. \quad (25)$$

Proposition 7 (central limit theorem for ψ_n^*). Suppose that **A1**, **A2**, **A3** and **A4** from Propositions 2, 4 and 5 hold. Then the following asymptotic linear expansion holds:

$$\sqrt{n}(\psi_n^* - \psi_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \text{IC}(O_i, Z_i) + o_P(1), \quad (26)$$

where

$$\text{IC}(O, Z) = D_1^{\star}(Q_{\theta_0, \epsilon_0})(W) + \phi_{\theta_0, \epsilon_0}'^{\top} S_0^{-1} D(\theta_0, \epsilon_0)(O, Z). \quad (27)$$

Furthermore, (26) entails that $\sqrt{n}(\psi_n^* - \psi_0)$ converges in distribution to the centered Gaussian distribution with a variance consistently estimated by

$$s_n^2 = \frac{1}{n} \sum_{i=1}^n D_1^\star(Q_n^*)(W_i)^2 + \frac{2}{n} \sum_{i=1}^n D_1^\star(Q_n^*)(W_i) \phi_n' S_n^{-1} D(\theta_n, \varepsilon_n)(O_i, Z_i) + (\phi_n' S_n^{-1} \sum_n (\phi_n' S_n^{-1})^\top).$$

Proposition 7 is the backbone of the statistical analysis of adaptive group sequential RCTs as constructed in Section 4. In particular, denoting the $(1 - \alpha)$ -quantile of the standard normal distribution by $\xi_{1-\alpha}$, the proposition guarantees that the asymptotic level of the confidence interval

$$\left[\psi_n^* \pm \frac{s_n}{\sqrt{n}} \xi_{1-\alpha/2} \right] \quad (28)$$

for the risk difference ψ_0 is $(1 - \varepsilon)$. The proofs of Propositions 6 and 7 are given in Section A.3.

Extensions.—We *conjecture* that the influence function IC computed at (O, Z) , see (27), is equal to

$$D_1^\star(Q_{\theta_0, \varepsilon_0})(W) + D_2^\star(P_{Q_{\theta_0, \varepsilon_0}, G^\star(\theta_0)})(O) \frac{G^\star(\theta_0)(A|V)}{gZ(A|V)}.$$

This conjecture is backed by the simulations that we carry out and present in Section 7. We will tackle the proof of the conjecture in future work. Let us assume for the moment that the conjecture is true. Then the asymptotic linear expansion (26) now implies that the asymptotic variance of $\sqrt{n}(\psi_n^* - \psi_0)$ can be consistently estimated by

$$s_n^{\star 2} = \frac{1}{n} \sum_{i=1}^n \left(D_1^\star(Q_n^*)(W_i) + D_2^\star(P_{Q_n^*, G^\star(\theta_n)})(O_i) \frac{G^\star(\theta_n)(A_i|V_i)}{gZ_i(A_i|V_i)} \right)^2,$$

another independent argument also showing that $s_n^{\star 2}$ converges toward

$$\text{Var}_{Q_0, G^\star(\theta_0)} D_2^\star(P_{Q_{\theta_0, \varepsilon_0}, G^\star(\theta_0)})(O)$$

i.e., the variance under the fixed design $P_{Q_0, G^\star(\theta_0)}$ of the efficient influence curve at

$$P_{Q_{\theta_0, \varepsilon_0}, G^\star(\theta_0)}.$$

Furthermore, the most essential characteristic of the joint methodologies of design adaptation and targeted maximum likelihood estimation is certainly the utmost importance of the role played by the *likelihood*. In this view, the targeted maximized log-likelihood of the data

$$\sum_{i=1}^n \left\{ \log Q_{W^n}(P_n)(W_i) + \log g_i^\star(A_i|W_i) + \log Q_{Y|A,W}(\theta_n, \varepsilon_n)(O_i) \right\}$$

provides us with a quantitative measure of the quality of the fit (targeted toward the parameter of interest). It is therefore possible, for example, to use that quantity for the sake of selection of different working models for Q_0 . As with TMLE for i.i.d data, we can use likelihood based crossvalidation to select among more general initial estimators indexed by fine-tuning parameters. The validity of such TMLEs for group sequential adaptive designs as studied here is outside the scope of this article.

6 Application to group sequential testing

We derive in this section a *group sequential testing procedure*, that is a testing procedure which repeatedly tries to make a decision at intervals rather than once all data are collected, or than after every new observation is obtained (such a testing procedure would be said fully sequential). We refer to (Jennison and Turnbull, 2000; Proschan et al., 2006) for a general presentation of group sequential testing procedures. The TMLE group sequential testing procedure is formally described in Section 6.1, and some arguments justifying why it should work well (as validated by the simulation study carried out in Section 8) are exposed in Section 6.2.

6.1 Description of the TMLE group sequential testing procedure

The problem at stake is to test the null “ $\Psi(Q) = \psi_0$ ” against “ $\Psi(Q) > \psi_0$ ” with asymptotic type I error α and asymptotic type II error β at some $\psi_1 > \psi_0$. We want to proceed group sequentially with $K - 2$ steps, based on the multi-dimensional t -statistic

$$(T_1^*, \dots, T_K^*) = \left(\frac{\sqrt{N_k}(\psi_{N_k}^* - \psi_0)}{s_{N_k}} \right)_{k \leq K}. \quad (29)$$

Here, $N_1, \dots, N_K$ are random sample sizes whose realizations on a specific trajectory depend on how fast the *information* accrues as the data are collected. In order to quantify this notion of information, we decide to consider the inverse $\frac{n}{s_n^2}$ of the estimated variance of the TMLE

ψ_n^* based on the n first observations $\mathbf{O}_n$ (as a proxy to its true, finite sample, inverse variance). Given a reference maximum committed information $I_{\max}$ and K increasingly order proportions $0 < p_1 < \dots < p_K = 1$, we set for every $k - K$

$$N_k = \inf \left\{ n \geq 1 : \frac{n}{s_n^2} \geq p_k I_{\max} \right\}.$$

The characterization of $I_{\max}$ depends on how we wish to “spend” the type I and type II errors at each step of the group sequential procedure and on how demanding is the power requirement (*i.e.*, how close ψ_1 is to ψ_0). Say that our spending strategies are summarized by the K -tuples of positive numbers $(\alpha_1, \dots, \alpha_K)$ and $(\beta_1, \dots, \beta_K)$ such that $\sum_{k=1}^K \alpha_k = \alpha$ and $\sum_{k=1}^K \beta_k = \beta$.

Now, let $(Z_1, \dots, Z_K)$ be distributed from the centered Gaussian distribution with covariance matrix $\mathcal{C} = (\sqrt{p_k \wedge l' p_k \vee l'})_{k, l \leq K}$. We assume that there exists a unique value $I > 0$, our $I_{\max}$, such that there exist a rejection boundary $(a_1, \dots, a_K)$ and a futility boundary $(b_1, \dots, b_K)$ satisfying $a_K = b_K$, $P(Z_1 \leq a_1) = \alpha_1$, $P(Z_1 + (\psi_1 - \psi_0)\sqrt{p_1 I} \leq b_1) = \beta_1$, and for every $1 \leq k < K$,

$$P(\forall j \leq k, b_j < Z_j < a_j \text{ and } Z_{k+1} \geq a_{k+1}) = \alpha_{k+1},$$

$$P(\forall j \leq k, b_j < Z_j + (\psi_1 - \psi_0)\sqrt{p_j I} < a_j \text{ and } Z_{k+1} + (\psi_1 - \psi_0)\sqrt{p_{k+1} I} \leq b_{k+1}) = \beta_{k+1}.$$

Note that the closer ψ_1 is to ψ_0 , the larger is $I_{\max}$ (actually, $\psi_1 \mapsto \psi_1 \sqrt{I_{\max}}$ is both upper bounded and bounded away from zero). Heuristically, the closer ψ_1 is to ψ_0 , the more difficult it is to decide between the null and its alternative while preserving the required type II error at ψ_1 , the more information is needed to proceed.

The targeted maximum likelihood group sequential testing procedure finally goes as follows: starting from $k = 1$,

- if $T_k^* \geq a_k$ then reject the null and stop accruing data,
- if $T_k^* \leq b_k$ then fail rejecting the null and stop accruing data,
- if $b_k < T_k^* < a_k$ then set $k \leftarrow k + 1$ and repeat.

If $(T_1^*, \dots, T_K^*)$ had the same distribution as $(Z_1, \dots, Z_K)$, then the latter rule would yield a testing procedure with the required type I error and type II error at the specified alternative parameter.

Clearly, our decision to target the optimal design $G^\star(\theta_0)$ which reduces as much as possible (over $\mathcal{G}_1$ and for our choice of working model) the asymptotic variance of the TMLE ψ_n^* guarantees, at least informally, that each N_k is stochastically smaller under $(Q_0, g_n^\star)$ -adaptive sampling than it would have been had another fixed design been used (or targeted). Thus,

resorting to the $(Q_0, g_n^\star)$ -adaptive sampling is likely to result in an earlier conclusion than what would have yielded another fixed (or targeted) sampling scheme.

6.2 Rationale of the TMLE group sequential testing procedure

The next proposition partially justifies the characterization of the TMLE group sequential testing procedure of Section 6.1. First, let us define $n_k = \lceil np_k \rceil$ (the smallest integer not smaller than np_k) for every $k \leq K$, then

$$(T_1, \dots, T_K) = \left(\frac{\sqrt{n_k}(\psi_{n_k}^* - \psi_0)}{s_{n_k}} \right)_{k \leq K}.$$

This other multi-dimensional t -statistic is a substitute to $(T_1^*, \dots, T_K^*)$ whose asymptotic study is easier to carry out. In particular, it is possible to derive the limit distribution of $(T_1, \dots, T_K)$ under the null and under a sequence of contiguous alternatives. The limit distribution is called the *canonical distribution* (see Theorems 11 in 12 in (van der Laan, 2008), Theorem 3 in (Chambaz and van der Laan, 2011) and Theorem 2.1 in (Zhu and Hu, 2010), where a similar result is obtained through a different approach based on the study of the limit distribution of a stochastic process defined over $(0, 1]$).

Proposition 8. Suppose that **A1**, **A2**, **A3** and **A4** from Propositions 2, 4 and 5 hold. Consider $f \in L^2(P_{Q_0, G^\star(\theta_0)})$, $f \neq 0$, with $P_{Q_0, G^\star(\theta_0)} f = 0$ and a fluctuation $\{Q_0(h): h \in \mathcal{H}\} \subset \mathcal{Q}$ of Q_0 with score f . Set $Q_{f/\sqrt{n}} = Q_0(1/\sqrt{n})$ for all $n \geq n_0$ (n_0 is such that $1/\sqrt{n_0} \in \mathcal{H}$). Assume also that **A5**. The score f is bounded, it is not proportional to $D^\star(P_{Q_0, G^\star(\theta_0)})$,

$$E_P \left[f(O) \middle| A, W \right] = 0, P_{Q_0, G^\star(\theta_0)} D^\star(P_{Q_0, G^\star(\theta_0)}) f > 0, \text{ and } P_{Q_0, G^\star(\theta_0)} \text{IC} f > 0.$$

The sequence $(Q_{f/\sqrt{n}})_{n \geq n_0}$ defines a sequence $(\psi_n)_{n \geq n_0}$ of contiguous parameters (“from direction f ”, which only fluctuates the conditional distribution of Y given (A, W)), with

$$\psi_n = \psi(Q_{f/\sqrt{n}}) > \psi(Q_0) \text{ for } n \text{ large enough. Introduce } \mu(f) = (\sqrt{p_1}, \dots, \sqrt{p_K}) \frac{P_{Q_0, G^\star(\theta_0)} \text{IC} f}{\sqrt{P_{Q_0, G^\star(\theta_0)} \text{IC}^2}}.$$

- (i) Under (Q_0, g_n) -adaptive sampling, $(T_1, \dots, T_K)$ converges in distribution to the centered Gaussian distribution with covariance matrix $\mathcal{C}$.
- (ii) Under $(Q_{f/\sqrt{n}}, g_n)$ -adaptive sampling, $(T_1, \dots, T_K)$ converges in distribution to the Gaussian distribution with mean $\mu(f)$ and covariance matrix $\mathcal{C}$.

The proof of Proposition 8 is given in Section A.4.

The rationale follows. Say that one concludes from Proposition 8 that $(i) (T_1^\star, \dots, T_K^\star)$ is also approximately distributed as $(Z_1, \dots, Z_K)$ under $(Q_0, g_n^\star)$ -adaptive sampling. Likewise, say that $(\sqrt{N_k}(\psi_{N_k}^* - \psi_1)/s_{N_k})_{k \leq K}$ is also approximately distributed as $(Z_1, \dots, Z_K)$ under $(Q_1, g_n^\star)$ -adaptive sampling, with $\Psi(Q_1) = \psi_1 > \psi_0$. Say that one is willing to substitute $p_k I_{\max}$ to $N_k/s_{N_k}^2$ for each $k \leq K$. Then $(ii) (T_1^*, \dots, T_K^*)$ is approximately distributed as $\sqrt{I_{\max}}(\psi_1 - \psi_0)(\sqrt{p_1}, \dots, \sqrt{p_K}) + (Z_1, \dots, Z_K)$ under $(Q_1, g_n^\star)$ -adaptive sampling. It appears that (i) and (ii) suffice to guarantee the wished asymptotic control of type I and type II errors. The rationale is validated by the simulation study undertaken in Section 8.

7 Simulation study of the performances of TMLE in adaptive covariate-adjusted RCTs

In this section, we carry out a simulation study of the performances of TMLE in adaptive group sequential RCTs as exposed in the previous sections. We present the simulation scheme in Section 7.1. The working model upon which the TMLE methodology relies is described in Section 7.2. How the TMLE-based confidence intervals behave is considered in Section 7.3 (empirical coverage) and in Section 7.4 (empirical width). An illustrating example is finally presented in Section 7.5.

7.1 Simulation scheme

We characterize the component $Q_0 = Q(P_0)$ of the true distribution P_0 of the observed data structure $O = (W, A, Y)$ as follows:

- the baseline covariate $W = (U, V)$ where U is uniformly distributed over the unit interval $[0; 1]$ and the subgroup membership covariate $V \in \mathcal{V} = \{1, 2, 3\}$ (hence $\mathcal{V} = 3$) satisfies

$$P_0(V = 1) = \frac{1}{2}, P_0(V = 2) = \frac{1}{3} \text{ and } P_0(V = 3) = \frac{1}{6};$$

- the conditional distribution of Y given (A, W) is the Gamma distribution characterized by the conditional mean

$$\bar{Q}_0(A, W) = 2U^2 / (2U + 1 + \rho \left(AV + \frac{1-A}{1+V} \right)) \quad (30)$$

with $\rho=1$ (we will set ρ to another value in Section 8) and the conditional variance

$$\text{Var}_{P_0}(Y | A, W) = \left(U + A(1+V) + \frac{1-A}{1+V} \right)^2.$$

In particular, the risk difference $\psi_0 = \psi(Q_0)$, our parameter of interest, is known in closed form:

$$\psi_0 = \frac{91}{72} \simeq 1.264,$$

and so is the variance

$$v^b(Q_0) = \text{Var}_{Q_{0,g^b}} D^{\star}(O; P_{Q_{0,g^b}})$$

of the efficient influence curve under balanced sampling. The numerical value of $v^b(Q_0)$ is reported in Table 1.

We target the design which (a) depends on the baseline covariate $W = (U, V)$ only through V (i.e., belongs to $\mathcal{G}_1$) and (b) minimizes the variance of the efficient influence curve of the parameter of interest ψ . The latter treatment mechanism $g^{\star}(Q_0)$ and optimal efficient asymptotic variance

$$v^{\star}(Q_0) = \text{Var}_{Q_{0,g^{\star}(Q_0)}} D^{\star}(O; P_{Q_{0,g^{\star}(Q_0)}})$$

are also known in closed form, and numerical values are reported in Table 1.

Let $n = (100, 250, 500, 750, 1000, 2500, 5000)$ be a sequence of sample sizes. We estimate $M = 1000$ times the risk difference $\psi_0 = \psi(Q_0)$ based on $O_{n\tau}^m(n_i), m = 1, \dots, M, i = 1, \dots, 7$, under

- i.i.d (Q_0, g^b) -balanced sampling,
- i.i.d $(Q_0, g^{\star}(Q_0))$ -optimal sampling,
- $(Q_0, g_{n\tau}^{\star})$ -adaptive sampling.

Finally, we emphasize that the observed data structure $O = (W, A, Y)$ is not bounded, whereas O is assumed bounded in Propositions 2, 3, 4, 5, 6 and 7.

7.2 Specifying the working model, reference design and a few twists On the working model.

For each $v \in \mathcal{V}$, let us denote $\theta(v) = (\beta_v, \sigma_v^2(0), \sigma_v^2(1)) \in \Theta_v, \Theta_v \subset \mathbb{R}^3 \times \mathbb{R}_+^* \times \mathbb{R}_+^*$ being compact, where the regression vector $\beta_v = (\beta_{v1}, \beta_{v2}, \beta_{v3})$ ($b = 3$ in (10)), then $\theta = (\theta_1, \theta_2, \theta_3) \in \Theta = \Theta_1 \times \Theta_2 \times \Theta_3$.

Following the description in Section 4.1, the working model $\mathcal{Q}_n^w$ that the TMLE methodology relies on is characterized by the conditional likelihood of Y given (A, W) :

$$Q_{Y|A,W}(\theta)(O) = \frac{1}{\sqrt{2\pi\sigma_V^2(A)}} \exp \left\{ -\frac{(Y - m(A, W; \beta_V))^2}{2\sigma_V^2(A)} \right\},$$

with the specific choice of conditional mean $\bar{Q}(\theta)(A, W)$ of Y given (A, W) :

$$\bar{Q}(\theta)(a, w) = m(a, w; \beta_V) = \beta_{v,1} + \beta_{v,2}u + \beta_{v,3}a$$

for all $a \in \mathcal{A}$ and $w = (u, v) \in \mathcal{W} = \mathbb{R} \times \mathcal{V}$. As required, condition **PARAM** is met.

Obviously, the working model is heavily misspecified:

- a Gaussian conditional likelihood is used instead of a Gamma conditional likelihood,
- the parametric forms of the conditional expectation and variance are wrong too.

On the reference design.—Regarding the choice of a reference fixed design $g^r \in \mathcal{G}_1$ (see Section 4.1), we select $g^r = g^b$ (the balanced design). The parameter θ_0 only depends on Q_0 and the working model, but its estimator θ_n depends on g^r , which may affect negatively its performances. Therefore, we propose to dilute the impact of the choice of g^r as an initial reference design as follows.

For a given sample size n , we first compute a first estimate θ_n^1 of θ_0 as in (12) but with $\lceil n/4 \rceil$ (the smallest integer not smaller than $n/4$) substituted for n in the sum. Then θ_n is computed as in (12) but this time with $G^*(\theta_n^1)(A_i | V_i)$ substituted for $g^r(A_i | V_i)$.

The proofs can be adapted in order to incorporate this modification of the procedure. We refer the interested reader to Section 8.5 in (van der Laan, 2008).

On additional details.—We decide arbitrarily to update the design each time $c = 25$ new observations are sampled. In addition, the first update only occurs when there are at least five completed observations in each treatment arm and for all V -strata. Thus, the minimal sample size at first update is 30. It can be shown that, under initialization using the balanced design, the expected sample size at the first sample size at which there are at least 5 observations in each arm equals 75. Finally, as a precautionary measure, we systematically apply a thresholding to the updated treatment mechanism: Using the notation of Section 4, $\max\{\delta, \min\{1 - \delta, g_i^*(A_i | V_i)\}\}$ is substituted for $g_i^*(A_i | V_i)$ in all computations. We arbitrarily choose $\delta = 0.01$.

7.3 Empirical coverage of the TMLE confidence intervals

We now invoke the central limit theorem stated in Proposition 7 to construct confidence intervals for the risk difference. Let us introduce, for all types of sampling and each sample size n_i , the confidence intervals

$$\mathcal{J}_{n_i, m} = \left[\psi_{n_i}^*(\mathbf{O}_{n_7}^m(n_i)) \pm \sqrt{\frac{v_{n_i}(\mathbf{O}_{n_7}^m(n_i))}{n_i} \xi_1 - \alpha/2} \right], \quad m = 1, \dots, M$$

where the definition of the variance estimator $v_n(O_{n\tau}^m(n))$ based on the n first observations $O_{n\tau}^m(n)$ depends on the sampling scheme:

- under i.i.d (Q_0, g^b) -balanced sampling, $v_n(O_{n\tau}^m(n))$ is the estimator of the asymptotic variance of the TMLE $\psi(Q_{n, iid}^*)$:

$$v_n(\mathbf{O}_{n_7}^m(n)) = \frac{1}{n} \sum_{i=1}^n D^{\star}(P_{Q_{n, iid}^*}^b)(O_i^m)^2, \quad (31)$$

- under i.i.d $(Q_0, g^{\star}(Q_0))$ -optimal sampling, $v_n(O_{n\tau}^m(n))$ is defined as in (31), replacing g^b with $g^{\star}(Q_0)$,
- under $(Q_0, g_n^{\star})$ -adaptive sampling, $v_n(O_{n\tau}^m(n)) = s_n^{*2}(O_{n\tau}^m(n))$, the estimator of the conjectured asymptotic variance of $\sqrt{n}(\psi_n^*(O_{n\tau}^m(n)) - \psi_0)$ computed on the n first observations $O_{n\tau}^m(n)$.

We are interested in the empirical coverage (reported in Table 2, top rows)

$$c_{n_i} = \frac{1}{M} \sum_{m=1}^M \mathbf{1}\{\psi_0 \in \mathcal{J}_{n_i, m}\}$$

guaranteed for each sampling scheme and every $i = 1, \dots, 7$ by $\{\mathcal{J}_{n_i, m} : m = 1, \dots, M\}$. The rescaled empirical coverage proportions M_{cn_i} should have a Binomial distribution with

parameter $(M, 1 - \alpha)$ with $\alpha = \alpha$ for every $i = 1, \dots, 7$. This property can be tested in terms of a standard binomial test, the alternative stating that $\alpha > \alpha$. This results in a collection of 7 p-values for each sampling scheme, as reported in Table 2 (bottom rows).

Considering each sampling scheme (*i.e.*, each row of Table 2) separately, we conclude that the $(1 - \alpha)$ -coverage cannot be declared defective under

- i.i.d (Q_0, g^b) -balanced sampling for any sample size $n_i \quad n_3 = 500$,
- i.i.d $(Q_0, g^{\star}(Q_0))$ -optimal sampling for any sample size $n_i \quad n_2 = 250$,
- $(Q_0, g_n^{\star})$ -adaptive sampling for any sample size $n_i \quad n_1 = 100$,

adjusting for multiple testing in terms of the Benjamini and Yekutieli procedure for controlling the False Discovery Rate at level 5%.

This is a remarkable result that not only validates the theory but also provides us with insight into the finite sample properties of the TMLE procedure based on adaptive sampling. The fact that the TMLE procedure behaves better under adaptive sampling scheme than under balanced i.i.d sampling scheme at sample size $n_1 = 100$ may not be due to mere chance only. Although the TMLE procedure based on an adaptive sampling scheme is initiated under the balanced sampling scheme (so that each stratum consists at the beginning of comparable numbers of patients assigned to each treatment arm, allowing to estimate, at least roughly, the required parameters), it starts deviating from it (as soon as every (A, V) -stratum counts 5 patients) each time 25 new observations are accrued. The poor performance of the TMLE procedure based on optimal i.i.d sampling scheme at sample size n_1 is certainly due to the fact that, by starting directly from the optimal sampling scheme (a choice we would not recommend in practice), too few patients from stratum $V = 3$ are assigned to treatment arm $A = 0$ among the n_1 first subjects. At larger sample sizes, the TMLE procedure performs equally well under adaptive sampling scheme and under both i.i.d schemes in terms of coverage.

7.4 Empirical width of the TMLE confidence intervals

Now that we know that the TMLE-based confidence intervals based on $(Q_0, g_n^\star)$ -adaptive sampling are valid confidence regions, it is of interest to compare the widths of the latter $(Q_0, g_n^\star)$ -adaptive sampling confidence intervals with their counterparts obtained under i.i.d (Q_0, g^b) -balanced or $(Q_0, g^*(Q_0))$ -optimal sampling schemes.

For this purpose, we compare on one hand, for each sample size n_i , the empirical distribution of $\{\sqrt{v_n(O_{n\tau}^m(n_i))} : m = 1, \dots, M\}$ as in (31) (i.e., the empirical distribution of width of the TMLE-based confidence intervals at sample size n_i obtained under i.i.d (Q_0, g^b) -balanced sampling, up to the factor $2\xi_1 - \alpha/2/\sqrt{n_i}$) to the empirical distribution of $\{s_n^\star(O_{n\tau}^m(n_i)) : m = 1, \dots, M\}$ (i.e., the empirical distribution of the width of the TMLE-based confidence intervals at sample size n_i obtained under $(Q_0, g_n^\star)$ adaptive sampling, up to the factor $2\xi_1 - \alpha/2/\sqrt{n_i}$ in terms of the two-sample Kolmogorov-Smirnov test, where the alternative states that the confidence intervals obtained under adaptive sampling are stochastically smaller than their counterparts under i.i.d balanced sampling. This results in 7 p -values, all equal to zero, which we nonetheless report in Table 3 (bottom row). In order to get a sense of how much narrower the confidence intervals obtained under adaptive sampling are, we also compute and report in Table 3 (top row) the ratios of empirical average widths

$$\frac{\frac{1}{M} \sum_{m=1}^M s_n^\star(O_{n\tau}^m(n_i))}{\frac{1}{M} \sum_{m=1}^M \sqrt{v_n(O_{n\tau}^m(n_i))}} \quad (32)$$

for each sample size n_i . Informally, this shows a 12% gain in width.

On the other hand, we also compare, for each sample size n_i , the empirical distribution of $\{\sqrt{v_n(O_{n\tau}^m(n_i))} : m = 1, \dots, M\}$ as in (31) but replacing g^b by $g^*(Q_0)$ (i.e., the empirical distribution of width of the TMLE-based confidence intervals at sample size n_i obtained under i.i.d $(Q_0, g^*(Q_0))$ -optimal sampling, up to the factor $2\xi_1 - \alpha/2/\sqrt{n_i}$ to the empirical distribution of $\{s_n^*(O_{n\tau}^m(n_i)) : m = 1, \dots, M\}$ (i.e., the empirical distribution of the width of the TMLE-based confidence intervals at sample size n_i obtained under $(Q_0, g_n^\star)$ -adaptive sampling, up to the factor $2\xi_1 - \alpha/2/\sqrt{n_i}$ in terms of the two-sample Kolmogorov-Smirnov test, where the alternative states that the confidence intervals obtained under adaptive sampling are stochastically larger than their counterparts under i.i.d optimal sampling. This results in 7 p -values that we report in Table 3 (bottom row). In order to get a sense of how similar the confidence intervals obtained under adaptive and i.i.d optimal sampling schemes are, we also compute and report for each sample size n_i in Table 3 (top row) the ratios of empirical average widths as in (32) replacing again g^b by $g^*(Q_0)$ in the definition (31) of $v_n(O_{n\tau}^m(n))$. Informally, this shows that the confidence intervals obtained under adaptive sampling are even slightly narrower in average than their counterparts obtained under i.i.d optimal sampling.

7.5 Illustrating example

So far we have been concerned with distributional results, answering to the questions *Does the confidence interval $[\psi_n^* \pm s_n^* \xi_1 - \alpha/2/\sqrt{n}]$ provide the wished coverage?* (yes, even for moderate sample size: see Section 7.3), *How does its width compare with the width of the confidence interval obtained under either i.i.d sampling scheme?* (well: see Section 7.4). In this section, we focus on a particular simulated trajectory (we select arbitrarily the first one, associated to $O_{n_7}^1$) for the sake of illustration.

Some interesting features of the selected simulated trajectory are apparent in Fig. 7.5 and Table 4.

For instance, we can follow the convergence of the TMLE ψ_n^* toward the true risk difference ψ_0 in the top plot of Fig. 7.5 and in the fifth column of Table 4. Similarly, the middle plot of Fig. 7.5 and the second to fourth columns of Table 4 illustrate the convergence of $g_n^\star$ toward $G^*(\theta_0)$, as stated in Proposition 3. What these plots and columns also teach us is that, in spite of the misspecified working model, the learned design $G^*(\theta_0)$ seems very close to the optimal treatment mechanism $g^*(Q_0)$ for the chosen simulation scheme and working model used in our simulation study. Moreover, the last column of Table 4 illustrates how the confidence intervals $[\psi_n^* \pm s_n^\star \xi_1 - 0.05/2/\sqrt{n}]$ shrink around the true risk difference ψ_0 as the sample size increases.

Yet, the bottom plot of Fig. 7.5 may be the most interesting of the three. It obviously illustrates the convergence of $s_n^{\star 2}$ toward $\text{Var}_{Q_0, G^{\star}(\theta_0)} D^{\star}(P_{Q_{\theta_0, \varepsilon_0}, G^{\star}(\theta_0)})(O)$ i.e., toward the variance under the fixed design $P_{Q_0, G^{\star}(\theta_0)}$ of the efficient influence curve at $P_{Q_{\theta_0, \varepsilon_0}, G^{\star}(\theta_0)}$. Hence, it also teaches us that the latter limit seems very close to the optimal asymptotic variance $v^*(Q_0)$ for the chosen simulation scheme and working model used in our simulation study. *More importantly, $s_n^{\star 2}$ strikingly converges to $v^*(Q_0)$ from below.* This finite sample characteristic may reflect the fact that the true finite sample variance of $\sqrt{n}(\psi_n^* - \psi_0)$ might be lower than $v^*(Q_0)$. Studying further this issue in depth is certainly very delicate, and goes beyond the scope of this article.

8 Simulation study of the performances of the TMLE group sequential testing procedure

In this section, we resume the simulation study undertaken in Section 7 in order to investigate the performances of the TMLE group sequential testing procedure presented in Section 6. We describe the simulation scheme in Section 8.1, then evaluate how the testing procedure behaves in terms of empirical type I and type II errors in Section 8.2 and how it behaves in terms of empirical distribution of sample size at decision in Section 8.3.

8.1 The simulation scheme, continued

We wish to test the null “ $\Psi(P) = \psi_0$ ” against its alternative “ $\Psi(P) > \psi_0$ ” ($\psi_0 = \frac{91}{72}$) with prescribed type I error $\alpha = 5\%$ and type II error $\beta = 10\%$ at the alternative $\psi_1 = \psi_0 + 0.4$. Depending on whether we want to investigate the empirical behaviors of the TMLE group sequential testing procedure with respect to (i) type I or (ii) type II errors, we test $M = 1000$ times the above hypotheses based on $3 \times M$ independent datasets obtained under

- (i) empirical type I error study:
 - i.i.d (Q_0, g^b) -balanced sampling,
 - i.i.d $(Q_0, g^*(Q_0))$ -optimal sampling,
 - $(Q_0, g_n^{\star})$ -adaptive sampling;
- (ii) empirical type II error study:
 - i.i.d (Q_1, g^b) -balanced sampling,
 - i.i.d $(Q_1, g^*(Q_0))$ -optimal sampling,
 - $(Q_1, g_n^{\star})$ -adaptive sampling,

where Q_1 is defined as Q_0 in Section 7.1 with $\rho = \psi_1/\psi_0$ (so that $\Psi(Q_1) = \psi_1$).

We decide to proceed in $K = 4$ steps, at proportions $(p_1, p_2, p_3, p_4) = (0.25, 0.50, 0.75, 1)$. We choose the α - and β -spending strategies characterized by the equalities $\sum_{l=1}^k \alpha_l = p_k^2 \alpha$ and $\sum_{l=1}^k \beta_l = p_k^2 \beta$, $k = 1, \dots, K$. This set of conditions characterizes the whole group sequential testing procedure, see Table 5 for the resulting numerical values.

8.2 Empirical type I and type II errors of the TMLE group sequential testing procedure

We report in Table 6 the empirical type I and type II errors obtained during the course of the simulation study. The values are strikingly close to each other in both cases. Even better, the empirical type I and type II errors obtained under $(Q, g_n^\star)$ -adaptive sampling are the lowest in both cases.

Since the number of times the null was falsely rejected for its alternative is Binomial with parameter (M, α) , it is possible to test rigorously if $\alpha = \alpha$ (as it should) against $\alpha > \alpha$. This yields three p -values that we report in Table 6. Because there are less than 50 wrong decisions for all sampling schemes, we naturally get large p -values, confirming (if necessary) that the type I error is under control.

Similarly, the number of times the null was falsely not rejected for its alternative is also Binomial with parameter (M, b) . Thus, it is possible to test rigorously if $b = \beta$ (as it should) against $b > \beta$. This yields three small p -values ($p < 0.001$ for the i.i.d balanced sampling scheme, 0.002 for the i.i.d optimal sampling scheme and 0.003 for the adaptive sampling scheme). This confirms the impression that the number of wrong decisions are all significantly larger than what random deviations from the reference distribution are likely to allow. Put bluntly, the type II error is not under control. However, there is no real need to worry. Indeed, if one rather tests $b = 12\%$ against $b > 12\%$, then the three p -values (reported in Table 6) are large.

In summary, for the considered scenario, the TMLE group sequential testing procedure described in Section 6.1 performs at least as well under the $(Q, g_n^\star)$ -adaptive sampling scheme as under both i.i.d (Q, g^b) -balanced and $(Q, g^*(Q))$ -optimal sampling schemes. The type I error control meets the requirements. The type II error control is defective, but only slightly defective.

8.3 Empirical distribution of sample size at decision of the TMLE group sequential testing procedure

Now that we know that the TMLE group sequential testing procedure performs at least as well under adaptive sampling scheme as under both i.i.d sampling schemes, let us compare the performances *in terms of sample size at decision*.

For this purpose, we simply report the average sample sizes at decision for each sampling scheme when evaluating the type I and type II errors control, see Table 7. The gain in average sample size at decision obtained by resorting to the $(Q, g_n^\star)$ -adaptive sampling scheme instead of the i.i.d (Q, g^b) -balanced sampling scheme is dramatic: In average, one

needs approximately 16% less observations in order to reach a conclusion under the $(Q, g_n^\star)$ -adaptive sampling scheme relative to the i.i.d (Q, g^b) -balanced sampling scheme. Furthermore, it appears that it is even more efficient to resort to the i.i.d $(Q, g^*(Q))$ -optimal sampling scheme: In average, one needs approximately 6% more observations in order to reach a conclusion under the i.i.d $(Q, g_n^\star)$ -adaptive sampling scheme relative to the i.i.d (Q, g^b) -balanced sampling scheme.

In summary, for the considered scenario, the TMLE group sequential testing procedure described in Section 6.1 reaches a decision more quickly under the $(Q, g_n^\star)$ -adaptive sampling scheme than under the i.i.d (Q, g^b) -balanced sampling scheme, with a gain in average sample size at decision of approximately 16%. Furthermore, the TMLE group sequential testing procedure reaches a decision more slowly under the $(Q, g_n^\star)$ -adaptive sampling scheme than under the i.i.d $(Q, g^*(Q))$ -optimal sampling scheme, but the loss in average sample size at decision reduces to approximately 6%.

A Proofs

Let us introduce for convenience the notations $\theta_\varepsilon = (\theta, \varepsilon)$, $\theta_0^\varepsilon = (\theta_0, \varepsilon_0)$ and $\theta_n^\varepsilon = (\theta_n, \varepsilon_n)$.

Moreover, let us recall that if μ is a probability measure and f a measurable function, then μf denotes the integral $\int f d\mu$.

A.1 Proof of Propositions 2 and 4, and a useful remark on Proposition 3 Proof of Proposition 2.

The cornerstone of the proof is to interpret θ_n as the solution in θ of the martingale estimating equation $\sum_{i=1}^n D_1(\theta)(O_i, Z_i) = 0$, where Z_i (defined in (9)) is the finite dimensional summary measure of past observation $\mathbf{O}_n(i-1)$ such that $g_i^\star$ depends on $\mathbf{O}_n(i-1)$ only through Z_i (hence the notation $g_i^\star = g_{Z_i}^\star$) and $D_1(\theta)(O, Z) = \dot{\ell}_{\theta,0}(O)g^T(A|V)/g_Z(A|V)$.

Note first that, for all $i \leq n$,

$$P_{Q_0, g_i^\star} D_1(\theta_0) = E_{P_{Q_0, g^r}} [\dot{\ell}_{\theta_0, 0}(O_i) \mid O_n(i-1)] = 0$$

by **A1** (changing the order of differentiation and integration is permitted by the dominated convergence theorem, because $\mathcal{O}$ is bounded and Θ is compact—see **PARAM**). Observe then that, by definition of θ_n , we have:

$$\left\| \frac{1}{n} \sum_{i=1}^n P_{Q_0, g_i^*} D_1(\theta_n) \right\| = \left\| \frac{1}{n} \sum_{i=1}^n D_1(\theta_n)(O_i, Z_i) - P_{Q_0, g_i^*} D_1(\theta_0) \right\| \leq \sup_{\theta \in \Theta} \left\| \frac{1}{n} \sum_{i=1}^n D_1(\theta_n)(O_i, Z_i) - P_{Q_0, g_i^*} D_1(\theta) \right\| \equiv M_n.$$

Now, since (a) $\sup_{\theta \in \Theta} \|D_1(\theta)\|_\infty < \infty$ (because $\mathcal{O}$ is bounded and Θ is compact) and (b) the standard entropy of $\mathcal{F} = \{D_1(\theta) : \theta \in \Theta\}$ for the supremum norm satisfies

$$H(\mathcal{F}, \|\cdot\|_\infty, \varepsilon) = \mathcal{O}(\log 1/\varepsilon) \text{ (see (van der Vaart, 1998), example 19.7), we can apply}$$

(componentwise) the Kolmogorov strong law of large numbers for martingales as in Theorem 8 of (Chambaz and van der Laan, 2009). This yields the almost sure convergence of M_n hence of $n^{-1} \sum_{i=1}^n P_{Q_0, g_i^*} D_1(\theta_n)$, to 0.

By a Taylor expansion of $\theta \mapsto P_{Q_0, g_i^*} \ell_{\theta, 0}$ at θ_0 (changing the order of differentiation and integration is permitted here too for the same reasons as above), it holds that, for all $i \leq n$ (recall that $P_{Q_0, g_i^*} D_1(\theta_0) = 0$),

$$P_{Q_0, g_i^*} D_1(\theta) = (P_{Q_0, g_i^*} \ell_{\theta_0, 0}''(\theta_n - \theta_0)) + o_P(\|\theta_n - \theta_0\|)$$

From this we deduce that $(P_{Q_0, g_i^*} \ell_{\theta_0, 0}'')(\theta_n - \theta_0)$ converges in probability to 0. Because $P_{Q_0, g_i^*} \ell_{\theta_0, 0}''$ is positive definite by **A2**, this implies the result.

Proof of Proposition 4.—The proof of (i) is very simple and typical of robust statistical studies. Indeed,

$$0 = P_{Q_0, G^*}(\theta_0) \frac{\partial}{\partial \varepsilon} \ell_{\theta_0, \varepsilon} \Big|_{\varepsilon = \varepsilon_0}$$

the latter expression also writing as

$$E_{Q_0, G^*(\theta_0)} \left\{ \frac{2A-1}{G^*(\theta_0)(A|V)} (Y - \bar{Q}(\theta_0^\varepsilon)(A, W)) \right\} = E_{Q_0, G^*(\theta_0)}$$

$$\{(\bar{Q}_0(1, W) - \bar{Q}(\theta_0^\varepsilon)(1, W)) - (\bar{Q}_0(0, W) - \bar{Q}(\theta_0^\varepsilon)(0, W))\} = \psi(Q_0) - \psi(Q_{\theta_0^\varepsilon}),$$

hence the result.

The proof of (ii) fundamentally relies on the fact that θ_n^ε solves the martingale estimating equation $\sum_{i=1}^n D(\theta^\varepsilon)(O_i, Z_i) = 0$, where $D(\theta^\varepsilon)$ is the extension of $D_1(\theta)$ defined in (19). Here too, $P_{Q_0, g_i^\star} D(\theta_0^\varepsilon) = 0$ for all $i \leq n$, and applying the Kolmogorov strong law of large numbers

for martingales (see Theorem 8 in (Chambaz and van der Laan, 2009)) yields that $n^{-1} \sum_{i=1}^n P_{Q_0, g_i^\star} D(\theta_n^\varepsilon)$ converges to zero almost surely. This entails the convergence in probability of θ_n^ε to θ_0^ε by a Taylor expansion of $\theta_\varepsilon \mapsto (P_{Q_0, g_i^\star} \dot{\ell}_{\theta, 0}^T, P_{Q_0, G^\star(\theta)} \frac{\partial}{\partial \varepsilon} \ell_{\theta, \varepsilon})$ at θ_0^ε .

Note that **A3** is a clear counterpart to **A1** from Proposition 2 but that there is no counterpart to **A2** from Proposition 2 in Proposition 4. Indeed, it automatically holds in the framework

of the proposition that $-P_{Q_0, G^\star(\theta_0)} \frac{\partial^2}{\partial \varepsilon^2} \ell_{\theta_0^\varepsilon} \Big|_{\varepsilon = \varepsilon_0} > 0$, the proof requiring that the latter

quantity be different from zero.

A useful remark on Proposition 3.—We deduce in Proposition 3 the convergence in probability and in L^1 of the adaptive design $g_n^\star$ to the limit design $G^*(\theta_0)$ from the convergence in probability of θ_n to θ_0 as obtained in Proposition 2. It is crucial for us that $n^{-1} \sum_{i=1}^n g_i^\star$ and $n^{-1} \sum_{i=1}^n (g_i^\star)^{-1}$ also converge to $G^*(\theta_0)$ and $1/G^*(\theta_0)$ (with an obvious notational shortcut) in probability and in L^1 . Fortunately, this is true because (a) the positive random variables $g_i^\star$ are uniformly bounded away from 0 and (b) if a sequence X_n converges in L^1 to 0 then $n^{-1} \sum_{i=1}^n X_i$ also converges in L^1 to 0 (by convexity of the L^1 -norm). Let us put this in a lemma:

Lemma 1. *For all $(a, v) \in \mathcal{A} \times \mathcal{V}$, $n^{-1} \sum_{i=1}^n g_i^*(a|v)$ and $n^{-1} \sum_{i=1}^n (g_i^*(a|v))^{-1}$ converge in probability and in L^1 to $G^*(\theta_0)(a|v)$ and $G^*(\theta_0)(a|v)^{-1}$.*

A.2 Proof of Proposition 5

The proof of Proposition 5 still relies on the facts that (a) θ_n^ε solves the martingale estimating equation $\sum_{i=1}^n D(\theta^\varepsilon)(O_i, Z_i) = 0$ and (b) $P_{Q_0, g_i^\star} D(\theta_0^\varepsilon) = 0$ for all $i \leq n$.

Consider the following equality:

$$\text{LHT} \equiv \frac{1}{n} \sum_{i=1}^n [D(\theta_n^\varepsilon)(O_i, Z_i) - D(\theta_0^\varepsilon)(O_i, Z_i)] = -\frac{1}{n} \sum_{i=1}^n [D(\theta_n^\varepsilon)(O_i, Z_i) - P_{Q_0, g_i^*} D(\theta_0^\varepsilon)]$$

(33)

A Taylor expansion of LHT at θ_0^ε first yields that

$$\text{LHT} = \frac{1}{n} \sum_{i=1}^n \left. \frac{\partial}{\partial \theta^\varepsilon} D(\theta^\varepsilon)(O_i, Z_i) \right|_{\theta^\varepsilon = \theta_0^\varepsilon} (\theta_n^\varepsilon - \theta_0^\varepsilon) + o_P(\|\theta_n^\varepsilon - \theta_0^\varepsilon\|) = A_n(\theta_n^\varepsilon - \theta_0^\varepsilon) + B_n(\theta_n^\varepsilon - \theta_0^\varepsilon) + o_P(\|\theta_n^\varepsilon - \theta_0^\varepsilon\|)$$

where $A_n = E[n^{-1} \sum_{i=1}^n \left. \frac{\partial}{\partial \theta^\varepsilon} D(\theta^\varepsilon)(O_i, Z_i) \right|_{\theta^\varepsilon = \theta_0^\varepsilon}]$ and

$$B_n = n^{-1} \sum_{i=1}^n \left. \frac{\partial}{\partial \theta^\varepsilon} D(\theta^\varepsilon)(O_i, Z_i) \right|_{\theta^\varepsilon = \theta_0^\varepsilon} - A_n.$$

The matrix A_n satisfies $A_n = n^{-1} E[\sum_{i=1}^n P_{Q_0, g_i^*} \left. \frac{\partial}{\partial \theta^\varepsilon} D(\theta^\varepsilon) \right|_{\theta^\varepsilon = \theta_0^\varepsilon}]$. Now, for each $i \leq n$,

$$P_{Q_0, g_i^*} \left. \frac{\partial}{\partial \theta^\varepsilon} D(\theta^\varepsilon) \right|_{\theta^\varepsilon = \theta_0^\varepsilon} = S_0 \quad (34)$$

where the latter matrix, defined in (21) and independent of i , is deterministic (this is due to the weighting). Thus, $A_n = S_0$. Moreover, $A_n = S_0$ is invertible because its determinant

equals the product of $P_{Q_0, G^\star(\theta_0)} \frac{\partial^2}{\partial \varepsilon^2} \ell_{\theta_0, \varepsilon | \varepsilon = \varepsilon_0} < 0$ and

$$\det E_{Q_0, G^\star(\theta_0)} \left[\ddot{\ell}_{\theta_0, 0^{(O)}} \frac{g^r(A|V)}{G^*(\theta_0)(A|V)} \right] = \det P_{Q_0, g^r \ddot{\ell}_{\theta_0, 0^{(O)}}}$$

which is negative by **A2**. Furthermore, (34) and $A_n = S_0$ (deterministic) also imply that B_n can be rewritten as

$$B_{n,n} \frac{1}{n} \sum_{i=1}^n \left[\frac{\partial}{\partial \theta^\varepsilon} D(\theta^\varepsilon)(O_i, Z_i) \Big|_{\theta^\varepsilon = \theta_0^\varepsilon} - P_{Q_0, g_i^*} \frac{\partial}{\partial \theta^\varepsilon} D(\theta^\varepsilon) \Big|_{\theta^\varepsilon = \theta_0^\varepsilon} \right]$$

and applying (componentwise) a standard law of large numbers for martingales (see for instance Theorem 2.4.2 in (Sen and Singer, 1993)) yields that B_n converges to 0 almost surely (note that $\sup_{\theta^\varepsilon \in \Theta \times \varepsilon} \left\| \frac{\partial}{\partial \theta^\varepsilon} D(\theta^\varepsilon) \right\|_\infty < \infty$). We emphasize that this proves the

convergence in probability of S_n to S_0 as stated in the proposition. Consequently, $\text{LHT} = S_0(\theta_n^\varepsilon - \theta_0^\varepsilon) + o_p(\|\theta_n^\varepsilon - \theta_0^\varepsilon\|)$ and (33) entails

$$\sqrt{n}(\theta_n^\varepsilon - \theta_0^\varepsilon) = -S_0^{-1}M_n + o_p(\sqrt{n}\|\theta_n^\varepsilon - \theta_0^\varepsilon\|) \quad (35)$$

with $M_n = n^{-1/2} \sum_{i=1}^n \left[D(\theta_0^\varepsilon)(O_i, Z_i) - P_{Q_0, g_i^*} D(\theta_0^\varepsilon) \right]$. It mainly remains to show that M_n satisfies a central limit theorem.

For this purpose, we apply Theorem 10 in (Chambaz and van der Laan, 2009). Define

$C_n = \sum_{i=1}^n P_{Q_0, g_i^*} D(\theta_0^\varepsilon) D(\theta_0^\varepsilon)^T$. It holds that

$$\begin{aligned} \frac{1}{n} C_n &= \frac{1}{n} \sum_{i=1}^n E_{Q_0, G^*(\theta_0)} \left[\frac{\tilde{D}(\theta_0^\varepsilon) \tilde{D}(\theta_0^\varepsilon)^T(O)}{g_i^*(A|V) G^*(\theta_0)(A|V)} \mathbf{1}_{O_n(i-1)} \right] \\ &= \sum_{(a,v) \in A \times V} E_{Q_0, G^*(\theta_0)} \left[\frac{\tilde{D}(\theta_0^\varepsilon) \tilde{D}(\theta_0^\varepsilon)^T(O)}{G^*(\theta_0)(A|V)} \mathbf{1}_{(A,V)=(a,v)} \right] \frac{1}{n} \sum_{i=1}^n \frac{1}{g_i^*(a|v)}. \end{aligned} \quad (36)$$

Now, by Lemma 1, (36) implies that $n^{-1} E C_n = \Sigma_0(1 + o(1))$, where Σ_0 (defined in (22)) is positive definite when **A4** is met. Indeed, assume on the contrary that there exists a vector $u \neq 0$ such that $u^T \Sigma_0 u = 0$. Then necessarily $\tilde{D}(\theta_0^\varepsilon)^T u = 0 P_{Q_0, G^*(\theta_0)}$ -almost surely, which contradicts **A4**. Using (36) again we also see that

$$\begin{aligned} \frac{1}{n} (C_n - E C_n) &= \sum_{(a,v) \in A \times V} E_{Q_0, G^*(\theta_0)} \left[\frac{\tilde{D}(\theta_0^\varepsilon) \tilde{D}(\theta_0^\varepsilon)^T(O)}{G^*(\theta_0)(A|V)} \mathbf{1}_{(A,V)=(a,v)} \right] \\ &\quad \times \left(\frac{1}{n} \sum_{i=1}^n \frac{1}{g_i^*(a|v)} - E \left[\frac{1}{n} \sum_{i=1}^n \frac{1}{g_i^*(a|v)} \right] \right) \end{aligned}$$

from which we deduce that $n^{-1}(C_n - EC_n)$ converges to 0 in probability by Lemma 1. Consequently, M_n converges in distribution to the centered Gaussian law with covariance matrix Σ_0 . In addition, $n^{-1} \sum_{i=1}^n D(\theta_0^\varepsilon) D(\theta_0^\varepsilon)^\top (O_i, Z_i)$ converges in probability to Σ_0 . Since (a) $\theta^\varepsilon \mapsto D(\theta^\varepsilon) D(\theta^\varepsilon)^\top$ is differentiable at θ_0^ε , (b) its derivative is bounded, and (c) we already know that $\|\theta_0^\varepsilon - \theta_0\| = o_P(1)$, yet another Taylor expansion allows to derive that Σ_n (defined in (23)) equals $\Sigma_0 (1 + o_P(1))$.

Let us go back to (35), knowing now that M_n satisfies a central limit theorem. We derive from it that $\sqrt{n} \|\theta_n^\varepsilon - \theta_0^\varepsilon\| = O_P(1) + o_P(\sqrt{n} \|\theta_n^\varepsilon - \theta_0^\varepsilon\|)$ hence $\|\theta_n^\varepsilon - \theta_0^\varepsilon\| = O_P(1)$. Thus, (35) does entail (20) since $P_{Q_{0g_i^*}} D(\theta_0^\varepsilon) = 0$ for all $i \leq n$. The stated central limit theorem on $\sqrt{n}(\theta_n^\varepsilon - \theta_0^\varepsilon)$ readily follows.

A.3 Proof of Propositions 6 and 7

Proof of Proposition 6.—The proof of Proposition 6 is twofold and relies on the decomposition

$$\psi_n^* - \psi_0 = (\psi_n^* - \psi(Q_n^\sim)) + (\psi(Q_n^\sim) - \psi_0), \quad (37)$$

where

$$Q_n^\sim = (QW(P_0), Q_{Y|A, W}(\theta_n^\varepsilon)). \quad (38)$$

First, a continuity argument and the convergence in probability of the stacked estimator θ_n^ε to θ_0^ε entail the convergence in probability of $\psi(Q_n^\sim)$ to $\psi(Q_{\theta_0^\varepsilon}) = \psi_0$ (the equality holds by (i) in Proposition 4). The conclusion follows because $\psi_n^* - \psi(Q_n^\sim)$ converges almost surely to zero. Indeed,

$$\psi_n^* - \psi(Q_n^\sim) = (P_{W,n} - P_{W,0}) P_{\theta_n^\varepsilon} \quad (39)$$

where, for all $\theta^\varepsilon \in \Theta \times \varepsilon$, $q\theta^\varepsilon = \bar{Q}(\theta^\varepsilon)(1, \cdot) - \bar{Q}(\theta^\varepsilon)(0, \cdot)$ and, with a slight abuse of notation, $P_{W,n}$ (respectively, $P_{W,0}$) denotes the empirical (respectively, true) marginal distribution of W . Thus,

$$|\psi_n^* - \psi(Q_n^\sim)| \leq \sup_{\theta^\varepsilon \in \Theta \times \varepsilon} |(P_{W,n} - P_{W,0}) q\theta^\varepsilon|, \quad (40)$$

where (a) $\sup_{\theta^\varepsilon \in \Theta \times \varepsilon} \|q\theta^\varepsilon\|_\infty < \infty$ and (b) the standard entropy of $\mathcal{F} = \{q\theta^\varepsilon : \theta^\varepsilon \in \Theta \times \varepsilon\}$ for

the supremum norm satisfies $H(\mathcal{F}, \|\cdot\|_\infty, \varepsilon) = \mathcal{O}(\log 1/\varepsilon)$ (see (van der Vaart, 1998), example

19.7). Therefore, the class $\mathcal{F}$ is $P_{W,0}$ -Glivenko-Cantelli and the right-hand side term of (40) converges to 0 $P_{W,0}$ -almost surely by the Glivenko-Cantelli theorem (i.i.d framework; see for instance Theorem 19.4 in (van der Vaart, 1998)).

Proof of Proposition 7.—The proof of (26) relies again on (37). It is easy to derive the asymptotic linear expansion of the first term. Indeed, we can rewrite (39) as

$$\begin{aligned} \sqrt{n}(\psi_n^* - \psi(Q_n)) &= \sqrt{n}(P_{W,n} - P_{W,0})q\theta_0^\varepsilon + \sqrt{n}(P_{W,n} - P_{W,0})(q\theta_n^\varepsilon - q\theta_0^\varepsilon) \\ &= \sqrt{n}P_{W,n}(q\theta_0^\varepsilon - P_{W,0}q\theta_0^\varepsilon) + o_p(1) \\ &= \sqrt{n}P_{W,n}D_1(Q_{\theta_0^\varepsilon}) + o_p(1), \end{aligned} \quad (41)$$

which holds subject to checking that $\sqrt{n}(P_{W,n} - P_{W,0})(q\theta_n^\varepsilon - q\theta_0^\varepsilon) = o_p(1)$. Since the class $\mathcal{F}$ introduced previously for the proof of Proposition 6 is also $P_{W,0}$ -Donsker, this is a consequence of Lemma 19.24 in (van der Vaart, 1998), provided that $X_n = P_{W,0}(q\theta_n^\varepsilon - q\theta_0^\varepsilon)^2$ converges in probability to 0. Now, θ_n^ε converges to θ_0^ε in probability, and the function $(w, \theta^\varepsilon) \mapsto q\theta^\varepsilon(w)$ continuously maps $\mathcal{W} \times \Theta \times \varepsilon$ onto $\mathbb{R}$ (thus it is uniformly bounded). Consequently, X_n (which is obviously non-negative) is upper bounded, so that it is equivalent to show that X_n converges to 0 in L^1 . By the Fubini theorem,

$$EX_n = \int E[(q\theta_n^\varepsilon(w) - q\theta_0^\varepsilon(w))^2] dP_{W,0}(w).$$

By the continuous mapping theorem (see Theorem 1.3.6 in (van der Vaart and Wellner, 1996)), $(q\theta_n^\varepsilon(w) - q\theta_0^\varepsilon(w))^2$ converges to 0 in probability (hence in L^1 , the variables being bounded) for all $w \in \mathcal{W}$. Therefore, the integrand of the right-hand side integral in the previous display converges pointwise to 0. Since it is also bounded, we conclude by the dominated convergence theorem that EX_n converges to 0. This completes the study of the first term of the decomposition (37).

Regarding the second term of the latter decomposition, we derive its asymptotic linear expansion by the delta-method (see Theorem 3.1 in (van der Vaart, 1998)) from that of $\sqrt{n}(\theta_n^\varepsilon - \theta_0^\varepsilon)$, see (20). Specifically,

$$\sqrt{n}(\psi(Q_n) - \psi_0) = \sqrt{n}(\phi(\theta_n^\varepsilon) - \phi(\theta_0^\varepsilon)) = \phi_{\theta_0^\varepsilon}^T S_0^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n D(\theta_0^\varepsilon)(O_i, Z_i) + o_p(1). \quad (42)$$

Combining (41) and (42) yields the wished asymptotic linear expansion (26) and the closed-form expression of the related influence function IC as given in (27).

A central limit theorem for real-valued martingales (see Theorem 9 in (Chambaz and van der Laan, 2009)) applied to (26) yields the stated convergence and validates the use of s_n^2 as an estimator of the asymptotic variance. To see this, note that $P_{Q_0, g_i^*} \text{IC} = 0$ for all $i \leq n$ and

define $c_n = \sum_{i=1}^n P_{Q_0, g_i^*} \text{IC}^2$. Now we emphasize that, for every $i \leq n$, firstly

$$P_{Q_0, g_i^*} D_1^*(Q_{\theta_0^\varepsilon})^2 = P_{Q_0, G^*(\theta_0)} D_1^*(Q_{\theta_0^\varepsilon})^2,$$

secondly

$$\begin{aligned} 2P_{Q_0, g_i^*} D_1^*(Q_{\theta_0^\varepsilon}) \phi_{\theta_0^\varepsilon}^T S_0^{-1} D(\theta_0^\varepsilon) &= 2\phi_{\theta_0^\varepsilon}^T S_0^{-1} P_{Q_0, g_i^*} D_1^*(Q_{\theta_0^\varepsilon}) D(\theta_0^\varepsilon) \\ &= 2\phi_{\theta_0^\varepsilon}^T S_0^{-1} P_{Q_0, G^*(\theta_0)} \left[\frac{D_1^*(Q_{\theta_0^\varepsilon}) \tilde{D}(\theta_0^\varepsilon)}{G^*(\theta_0)} \right] \end{aligned}$$

and thirdly

$$P_{Q_0, g_i^*} (\phi_{\theta_0^\varepsilon}^T S_0^{-1} D(\theta_0^\varepsilon))^2 = \phi_{\theta_0^\varepsilon}^T S_0^{-1} (P_{Q_0, g_i^*} D(\theta_0^\varepsilon) D(\theta_0^\varepsilon)^T) (\phi_{\theta_0^\varepsilon}^T S_0^{-1})^T.$$

By combining the last three equalities, we thus obtain that

$$\begin{aligned} \frac{1}{n} c_n &= P_{Q_0, G^*(\theta_0)} D_1^*(Q_{\theta_0^\varepsilon})^2 + 2\phi_{\theta_0^\varepsilon}^T S_0^{-1} P_{Q_0, G^*(\theta_0)} \left[\frac{D_1^*(Q_{\theta_0^\varepsilon}) \tilde{D}(\theta_0^\varepsilon)}{G^*(\theta_0)} \right] + \phi_{\theta_0^\varepsilon}^T S_0^{-1} \frac{1}{n} c_n (\phi_{\theta_0^\varepsilon}^T S_0^{-1})^T \\ & \quad (43) \end{aligned}$$

where C_n was introduced in the proof of Proposition 5. Since we already showed in the latter proof that $n^{-1}EC_n = \Sigma_0(1 + o_p(1))$, (43) notably yields the convergence of $n^{-1}EC_n$ toward

$$\begin{aligned} & P_{Q_0, G^*(\theta_0)} D_1^*(Q_{\theta_0^\varepsilon})^2 + 2\phi_{\theta_0^\varepsilon}^T S_0^{-1} P_{Q_0, G^*(\theta_0)} \left| \frac{D_1^*(Q_{\theta_0^\varepsilon}) \tilde{D}(\theta_0^\varepsilon)}{G^*(\theta_0)} \right| + \phi_{\theta_0^\varepsilon}^T S_0^{-1} \Sigma_0 (\phi_{\theta_0^\varepsilon}^T S_0^{-1})^T \\ &= P_{Q_0, G^*(\theta_0)} \left(D_1^*(Q_{\theta_0^\varepsilon}) + \phi_{\theta_0^\varepsilon}^T S_0^{-1} \frac{\tilde{D}(\theta_0^\varepsilon)}{G^*(\theta_0)} \right)^2 \equiv s^2 \end{aligned}$$

which is positive when **A4** is met. In addition, (43) also entails that $n^{-1}(c_n - EC_n)$ converges to 0 in probability, therefore implying the convergence in distribution of $\sqrt{n}(\psi_n^* - \psi_0)$ to the centered Gaussian distribution with variance s^2 as well as the convergence in probability of $n^{-1} \sum_{i=1}^n \text{IC}(O_i, Z_i)^2$ to s^2 . Since (a) $\theta^\varepsilon \mapsto D_1^*(Q_{\theta^\varepsilon})$ and $\theta^\varepsilon \mapsto D(\theta^\varepsilon)$ are differentiable at

θ_0^ε , (b) both with uniformly bounded derivatives, (c) we already know that

$\|\theta_n^\varepsilon - \theta_0^\varepsilon\| = o_p(1)$, and (d) $\phi'_{\theta_n^\varepsilon}$ and S_0^{-1} are consistently estimated by ϕ'_n and S_n^{-1} , yet

another Taylor expansion (and the continuous mapping theorem, see for instance Theorem 1.3.6 in (van der Vaart and Wellner, 1996)) finally yields the stated convergence in probability of s_n^2 to s^2 , therefore completing the proof.

A.4 Proof of Proposition 8

The fact that $\Psi(Q_{f/\sqrt{n}}) > \Psi(Q_0)$ for n large enough is a consequence of the expansion

$\Psi(Q_{f/\sqrt{n}}) = \Psi(Q_0) + P_{Q_0, G^*(\theta_0)} D^*(P_{Q_0, G^*(\theta_0)}) f/\sqrt{n} + o(1/\sqrt{n})$, which holds because Ψ is

pathwise differentiable at $P_{Q_0, g}$ (for any g) relative to the maximal tangent space, with efficient influence curve $D^*(Q_0, g)$. The rest of Proposition 8 is a corollary of the following lemma.

Lemma 2. Let Λ_n be the log-likelihood ratio of the $(Q_{f/\sqrt{n}}, g_n^*)$ experiment relative to the

(Q_0, g_n^*) experiment. Under (Q_0, g_n^*) , the vector

$(\sqrt{n}(\psi_{n_1}^* - \Psi(Q_0)), \dots, (\sqrt{n}(\psi_{n_K}^* - \Psi(Q_0)), \Lambda_n)$ converges in law to the Gaussian distribution

with mean $(0, \dots, 0, -\frac{1}{2} P_{Q_0, G^*(\theta_0)} f^2)$ and covariance matrix

$$\begin{pmatrix} P_{Q_0, G^*(\theta_0)} \text{IC}^2 \times C & P_{Q_0, G^*(\theta_0)} \text{IC} f \times (\sqrt{pK})^T \\ P_{Q_0, G^*(\theta_0)} \text{IC} f \times (\sqrt{p1}, \dots, \sqrt{pK}) & P_{Q_0, G^*(\theta_0)} f^2 \end{pmatrix}.$$

In particular, the $(Q_{f/\sqrt{n}}, g_n^\star)$ and $(Q_0, g_n^\star)$ experiments are mutually contiguous.

The limiting distribution of $(T_1, \dots, T_K)$ under $(Q_0, g_n^\star)$ is easily derived from Lemma 2. Le Cam's third lemma (see Example 6.7 in (van der Vaart, 1998)) solves the problem of obtaining the limiting distribution of $(T_1, \dots, T_K)$ under $(Q_{f/\sqrt{n}}, g_n^\star)$ from that under $(Q_0, g_n^\star)$. Only the asymptotic means are different. We do not repeat the proof here, as it is exactly the same (up to some minor variations in notation) as the proof of Lemma 5 in (Chambaz and van der Laan, 2011).

Proof of Lemma 2.—Since f is bounded and $E_{P_{Q_0, G^\star(\theta_0)}} [f(O) | A, W] = 0$ (here, one could equivalently use the notation E_{Q_0}), we can assume without loss of generality that the fluctuation $\{Q_0(h) : h \in \mathcal{H}\}$ is characterized by

$$\frac{dQ_0(h)}{dQ_0} = (1 + hf(O))$$

for all $h \in \mathcal{H} = (-1/\|f\|_\infty, 1/\|f\|_\infty)$. Let us denote by q_n and q_0 the conditional densities of Y given (A, W) under $Q_{f/\sqrt{n}}$ and Q_0 .

The log-likelihood ratio $\Lambda_n = \sum_{i=1}^n \log q_n/q_0(O_i)$, by conditional independence of $O_1, \dots, O_n$ given $((A_1, W_1), \dots, (A_n, W_n))$ and because the marginal distribution of W is the same under $Q_{f/\sqrt{n}}$ as under Q_0 , (see (4) and note that the $g_i^\star(A_i | W_i)$ factors cancel out). The fluctuation is chosen in such a way that $q_n/q_0 = (1 + f/\sqrt{n})$, hence

$\Lambda_n = \sqrt{n}^{-1} \sum_{i=1}^n f(O_i) - \frac{1}{2} n^{-1} \sum_{i=1}^n f^2(O_i) + n^{-1} \sum_{i=1}^n f^2(O_i) R(f(O_i)/\sqrt{n})$, where the function R is characterized by $\log(1 + x) = x - \frac{1}{2}x^2 + x^2 R(x)$ (all $x > -1$). In particular, R is increasing and $R(x) \rightarrow R(0) = 0$ when $x \rightarrow 0$. Since f is bounded, the expansion is valid for n large enough. Moreover, the last term satisfies

$$\left| \frac{1}{n} \sum_{i=1}^n f^2(O_i) R(f(O_i)/\sqrt{n}) \right| \leq \frac{1}{n} \sum_{i=1}^n f^2(O_i) \times R(\|f\|_\infty/\sqrt{n}),$$

so that if $n^{-1} \sum_{i=1}^n f^2(O_i) = O_P(1)$ then it is $o_P(1)$. Furthermore, the Kolmogorov law of large numbers for martingales implies that

$$n^{-1} \sum_{i=1}^n f^2(O_i) = n^{-1} \sum_{i=1}^n f^2 P_{Q_0, g_i^\star} + o_P(1) = P_{Q_0, \bar{g}_n^\star} f^2 + o_P(1), \text{ where}$$

$$\bar{g}_n^\star = n^{-1} \sum_{i=1}^n g_i^\star \text{ Denote } \text{Var}(f)(A, W) = EQ_0[f^2(O) | A, W]. \text{ We have that}$$

$$\begin{aligned} P_{Q_0, \bar{g}_n^*} f^2 &\equiv E_{Q_0, \bar{g}_n^*} [f^2(O_n) | \mathbf{O}_n(n-1)] = E_{Q_0, \bar{g}_n^*} [\text{Var}(f)(A, W)(O_n) | \mathbf{O}_n(n-1)] \\ &= E_{Q_0} [\bar{g}_n^*(1) V \text{Var}(f)(1, W) + \bar{g}_n^*(0) W \text{Var}(f)(1, V) | \mathbf{O}_n(n-1)] \end{aligned}$$

for V the relevant subvector of W upon which each g_i^* depends. By Lemma 1, $\bar{g}_i^*$ converges in probability to $G^*(\theta_0)$. Consequently, the continuous mapping theorem (see Theorem 1.3.6 in (van der Vaart and Wellner, 1996)) yields that $P_{Q_0, \bar{g}_n^*} f^2 = P_{Q_0, G^*(\theta_0)} f^2 + o_P(1) = O_P(1)$.

In summary, we obtain the following asymptotic linear expansion:

$$\Lambda_n = \frac{1}{\sqrt{n}} \sum_{i=1}^n f(O_i) - \frac{1}{2} P_{Q_0, G^*(\theta_0)} f^2 + o_P(1). \quad (44)$$

Introduce now $Z'_i = (1\{i \leq n_1\}, \dots, 1\{i \leq n_K\})$ and the bounded (and measurable) function F such that $F(O_i, Z_i, Z'_i) = (Z'_i \text{IC}(O_i, Z_i), f(O_i))$. We show that

$$\begin{aligned} M_n &= n^{-1} \sum_{i=1}^n [F(O_i, Z_i, Z'_i) - P_{Q_0, g_i^*} F] = n^{-1} \sum_{i=1}^n F(O_i, Z_i, Z'_i) \quad (\text{since } P_{Q_0, g_i^*} F = P_{Q_0, g_i^*} f \\ &= 0 \text{ for all } i \leq n) \end{aligned}$$

satisfies a central limit theorem. The proof is very similar to the corresponding part of proof of Lemma 5 in (Chambaz and van der Laan, 2011), hence we only give an outline, focusing on the differences. First, it holds for every $1 \leq k, l \leq K$ that

$$\begin{aligned} n^{-1} \sum_{i=1}^{n_k} P_{Q_0, g_i^*} \text{IC} f &= p_k P_{Q_0, G^*(\theta_0)} \text{IC} f + o_P(1), \quad n^{-1} \sum_{i=1}^{n_k \wedge n_l} P_{Q_0, g_i^*} \text{IC}^2 = p_{k \wedge l} P_{Q_0, G^*(\theta_0)} \text{IC}^2 \\ &+ o_P(1), \quad \text{and } n^{-1} \sum_{i=1}^n P_{Q_0, g_i^*} f^2 = P_{Q_0, G^*(\theta_0)} f^2 + o_P(1) \end{aligned}$$

This entails that the matrix $E n^{-1} \sum_{i=1}^n P_{Q_0, g_i^*} F^\top F$ converges to

$$\begin{pmatrix} P_{Q_0, G^*(\theta_0)} \text{IC}^2 \times (p_k \wedge p_l)_{k, l \leq K} & P_{Q_0, G^*(\theta_0)} \text{IC} f \times (p_1, \dots, p_K)^\top \\ P_{Q_0, G^*(\theta_0)} \text{IC} f \times (p_1, \dots, p_K) & P_{Q_0, G^*(\theta_0)} f^2 \end{pmatrix},$$

which is positive definite if and only if its determinant is positive. The latter determinant equals a positive constant times $\left[P_{Q_0, G^*(\theta_0)} f^2 - (P_{Q_0, G^*(\theta_0)} \text{IC} f)^2 / P_{Q_0, G^*(\theta_0)} \text{IC}^2 \right]$, which is positive too by the Cauchy-Schwarz inequality (f is not proportional to $D^\star \left(P_{Q_0, G^*(\theta_0)} \right)$).

Therefore, Theorem 8 in (Chambaz and van der Laan, 2011) applies and $\sqrt{n} M_n$ converges in

law to the centered Gaussian distribution with the covariance matrix given in the above display. The conclusion readily follows.

References

- Atkinson AC and Biswas A . Adaptive biased-coin designs for skewing the allocation proportion in clinical trials with normal responses. *Stat. Med*, 24(16):2477–2492, 2005.16007571
- Bandyopadhyay U and Biswas A . Adaptive designs for normal responses with prognostic factors. *Biometrika*, 88(2):409–419, 2001.
- Chambaz A and van der Laan MJ . Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate. Technical report, Division of Biostatistics, University of California, Berkeley, 2009.
- Chambaz A and van der Laan MJ . Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate: Theoretical study. *Int. J. Biostat*, 7(1), 2011.
- Emerson SS . Issues in the use of adaptive clinical trial designs. *Stat. Med*, 25:3270–3296, 2006.16906553
- The Food & Drug Administration. Critical path opportunities list Technical report, U.S. Department of Health and Human Services, Food and Drug Administration, 2006.
- Golub HL . The need for mor efficient trial designs. *Stat. Med*, 25:3231–3235, 2006.16819752
- Hu F and Rosenberger WF . The theory of response adaptive randomization in clinical trials. New York Wiley, 2006.
- Jennison C and Turnbull BW . Group Sequential Methods with Applications to Clinical Trials. Chapman & Hall/CRC, Boca Raton, FL, 2000.
- Proschan MA , Lan GKK , and Wittes JT . Statistical Monitoring of Clinical Trials: A Unified Approach. Statistics for biology and health. Springer, New-York, 2006.
- Rosenberger WF , Vidyashankar AN , and Agarwal DK . Covariate-adjusted response-adaptive designs for binary response. *J. Biopharm. Statist*, 11(227–236), 2001.
- Rosenberger WF . New directions in adaptive designs. *Statist. Sci*, 11:137–149, 1996.
- Sen PK and Singer JM . Large sample methods in statistics. Chapman & Hall, New York, 1993 An introduction with applications.
- Shao J , Yu X , and Bob Zhong B . A theory for testing hypotheses under covariate-adaptive randomization. *Biometrika*, 2010.
- van der Laan MJ . The construction and analysis of adaptive group sequential designs. Technical report 232, Division of Biostatistics, University of California, Berkeley, 3 2008.
- van der Laan MJ and Rubin D . Targeted maximum likelihood learning. *Int. J. Biostat*, 2(1), 2006.
- van der Vaart AW . Asymptotic Statistics. Cambridge University Press, 1998.
- van der Vaart AW and Wellner JA . Weak Convergence and Emprical Processes. Springer-Verlag New York, 1996.
- Zhang Li-Xin , Hu Feifang , Cheung Siu Hung , and Chan Wai Sum . Asymptotic properties of covariate-adjusted response-adaptive designs. *Ann. Statist*, 35(3):1166–1182, 2007.
- Zhu H and Hu F . Sequential monitoring of response-adaptive randomized clinical trials. *Ann. Statist*, 38(4):2218–2241, 2010.

Figure 1: A causal graph describing the data generating mechanism for adaptive group sequential design.

An from nodes $\Lambda_1, \dots, \Lambda_r$ to node Υ means that there exists a deterministic function f and an independent source of randomness U such that $\Upsilon = f(\Lambda_1, \dots, \Lambda_r, U)$.

Figure 2: Illustrating the TMLE procedure under $(Q_0, g_n^\star)$ -adaptive sampling scheme.

These three plots illustrate how the TMLE procedure behaves (on a simulated trajectory) as the sample size (on x -axis, logarithmic scale; the vertical grey lines indicate the sample sizes n_i , $i = 1, \dots, 7$) increases. Top plot: we represent the sequence $\psi_n^*(O_{n_7}^1(n))$ at intermediate sample sizes n . The horizontal grey line indicates the true value of the risk difference ψ_0 . Middle plot: we represent the three sequences $g_n^\star(1|1) = G^\star(\theta_n(O_{n_7}^1(n)))(1|1)$ (bottom curve), $g_n^\star(1|2) = G^\star(\theta_n(O_{n_7}^1(n)))(1|2)$ (middle curve) and $g_n^\star(1|3) = G^\star(\theta_n(O_{n_7}^1(n)))(1|3)$ (top curve) at intermediate sample sizes n . The three horizontal grey lines indicate the optimal allocation probabilities $g^*(Q_0)(1|1)$ (bottom line), $g^*(Q_0)(1|2)$ (middle line) and $g^*(Q_0)(1|3)$ (top line). Bottom plot: we represent the sequence $s_n^{\star 2}$ of estimated asymptotic variance of $\sqrt{n}(\psi_n^* - \psi_0)$

at intermediate sample sizes n . The horizontal grey line indicates the value of the optimal variance $v^*(Q_0)$. See also Table 4.

Table 1:

Numerical values of the allocation probabilities and variance of the efficient influence curve under either the balanced g^b or the targeted optimal $g^\star(Q_0)$ sampling schemes. The ratio of the variances of the efficient influence curve under targeted optimal and balanced sampling schemes satisfies $R(Q_0) = v^\star(Q_0)/v^b(Q_0) \simeq 0.762$.

Sampling scheme (Q_0, g)	Allocation probabilities			Variance
	$g(1 v=1)$	$g(1 v=1)$	$g(1 v=1)$	
(Q_0, g^b) -balanced	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	23.864
$(Q_0, g^\star(Q_0))$ -optimal	0.707	0.799	0.849	18.181

Table 2:

Checking the adequateness of the coverage guaranteed by our simulated confidence intervals. We test if the rescaled empirical coverage Binomial random variables M_{cn_i} have parameter $(M, 1 - \alpha)$, the alternative stating that they have parameter $(M, 1 - a)$ with $a > \alpha$. We report the values c_{n_i} (top row) and corresponding p -values (bottom row, between parentheses) of the Binomial test for each sample size and each sampling scheme.

Sampling scheme	Sample size						
	n_1	n_2	n_3	n_4	n_5	n_6	n_7
i.i.d (Q_0, g^b) -balanced	0.913 (p < 0.001)	0.925 (p < 0.001)	0.939 (0.067)	0.934 (0.015)	0.945 (0.253)	0.940 (0.087)	0.946 (0.300)
i.i.d $(Q_0, g^*(Q_0))$ -optimal	0.894 (p < 0.001)	0.941 (0.111)	0.940 (0.087)	0.953 (0.688)	0.954 (0.739)	0.947 (0.351)	0.947 (0.351)
adaptive $(Q_0, g_{n_7}^*)$	0.934 (0.015)	0.939 (0.067)	0.956 (0.827)	0.945 (0.253)	0.943 (0.172)	0.933 (0.011)	0.952 (0.634)

Table 3:

Comparing the width of our confidence intervals. On one hand, we test for each sample size n_i if the TMLE-based confidence intervals obtained under $(Q_0, g_n^\star)$ -adaptive sampling are narrower stochastically than the TMLE-based confidence intervals obtained under i.i.d (Q_0, g^b) -balanced sampling in terms of the two-sample Kolmogorov-Smirnov test. On the other hand, we test for each sample size n_i if the TMLE-based confidence intervals obtained under $(Q_0, g_n^\star)$ -adaptive sampling are wider stochastically than the TMLE-based confidence intervals obtained under i.i.d $(Q_0, g^*(Q_0))$ -optimal sampling in terms of the two-sample Kolmogorov-Smirnov test. We report the p -values (bottom rows, between parentheses). In addition, we report for each sample size n_i the ratios of average widths as defined in (32).

Comparison	Sample size						
	n_1	n_2	n_3	n_4	n_5	n_6	n_7
$(Q_0, g_{n_7}^\star) \text{ vs } (Q_0, g^b)$	0.856 (0)	0.871 (0)	0.879 (0)	0.880 (0)	0.878 (0)	0.877 (0)	0.876 (0)
$(Q_0, g_{n_7}^\star) \text{ vs } (Q_0, g^\star(Q_0))$	0.962 (0.144)	0.977 (0.236)	0.992 (0.100)	0.995 (0.060)	0.997 (0.407)	1.000 (0.236)	1.000 (0.144)

Table 4:

Illustrating the TMLE procedure under $(Q_0, g_n^\star)$ -adaptive sampling scheme. This table illustrates how the TMLE procedure behaves (on a simulated trajectory) as the sample size increases. We report, at each sample size n_i , the updated adapted design $g_n^\star = G^\star(\theta_n(O_{n_7}^1(n)))$ (columns two to four), current TMLE $\psi_n^*(O_{n_7}^1(n))$ (column five) and 95% confidence interval $I_{n,1}$ as in (28) with $s_n^*(O_{n_7}^1(n))$ substituted to s_n . The true risk difference $\psi_0 = \simeq 1.264$ belongs to all confidence intervals. See also Fig. 7.5.

Sample size	Allocation probabilities			TMLE	Confidence interval
n	$g_n^\star(1 1)$	$g_n^\star(1 2)$	$g_n^\star(1 3)$	ψ_n^*	$[\psi_n^* \pm s_n^* \xi_1 - 0.05/2/\sqrt{n}]$
n_1	0.589	0.764	0.766	1.252	[0.722;1.783]
n_2	0.624	0.775	0.707	1.388	[0.974;1.802]
n_3	0.679	0.767	0.795	1.361	[1.037;1.685]
n_4	0.677	0.757	0.813	1.341	[1.068;1.615]
n_5	0.670	0.760	0.806	1.250	[1.012;1.488]
n_6	0.677	0.788	0.835	1.288	[1.126;1.451]
n_7	0.694	0.793	0.834	1.273	[1.157;1.389]

Table 5:

Specifics of the TMLE group sequential testing procedure

Ψ_0	Ψ_1	$I_{\max}$
1.264	1.664	57.927

Step k	1	2	3	4
Rejection boundary	2.734	2.305	2.005	1.715
α -spending	0.3125%	0.9375%	1.5625%	2.1875%

Step k	1	2	3	4
Futility boundary	-0.976	0.132	0.961	1.715
β -spending	0.625%	1.875%	3.125%	4.375%

Table 6:

Checking the adequateness of the type I and type II error controls on our simulated TMLE group sequential testing procedures. We test if the Binomial random numbers of times the null is falsely rejected have parameter (M, α) , the alternative stating that they have parameter (M, a) with $a > \alpha$. We also test if the Binomial random numbers of times the null is falsely not rejected have parameter $(M, 12\%)$, the alternative stating that they have parameter (M, b) with $b > 12\%$. We report the empirical type I and type II errors and corresponding p-values

Sampling scheme	Type I error ($Q = Q_0$) Empirical error (p-value)	Type II error ($Q = Q_1$) Empirical error (p-value)
i.i.d (Q, g^b)-balanced	0.040 (0.919)	0.132 (0.113)
i.i.d ($Q, g^\star(Q)$)-optimal	0.043 (0.827)	0.128 (0.203)
adaptive ($Q, g_n^\star$)	0.040 (0.919)	0.126 (0.261)

Table 7:

Comparing the average sample sizes at decision for each sampling scheme when evaluating the type I and type II errors control.

Sampling scheme	Type I error ($Q = Q_0$) Average sample size	Type II error ($Q = Q_1$) Average sample size
i.i.d (Q, g^b)-balanced	786.63	895.64
i.i.d ($Q, g^\star(Q)$)-optimal	620.71	713.47
adaptive ($Q, g_n^\star$)	661.14	746.72