

UC Irvine

UC Irvine Previously Published Works

Title

Improving Precision through Adjustment for Prognostic Variables in Group Sequential Trial Designs: Impact of Baseline Variables, Short-Term Outcomes, and Treatment Effect Heterogeneity

Permalink

<https://escholarship.org/uc/item/1p82z4dw>

Authors

Qian, Tianchen
Rosenblum, Michael
Qiu, Huitong

Publication Date

2019-10-13

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Improving Precision through Adjustment for Prognostic Variables in Group Sequential Trial Designs: Impact of Baseline Variables, Short-Term Outcomes, and Treatment Effect Heterogeneity

Tianchen Qian*, Michael Rosenblum[†] and Huitong Qiu[‡]

October 28, 2019

Abstract

In randomized trials, appropriately adjusting for baseline variables and short-term outcomes can lead to increased precision and reduced sample size. We examine the impact of such adjustment in group sequential designs, i.e., designs with preplanned interim analyses where enrollment may be stopped early for efficacy or futility. We address the following questions: how much precision gain can be obtained by appropriately adjusting for baseline variables and a short-term outcome? How is this precision gain impacted by factors such as the proportion of pipeline participants (those who enrolled but haven't yet had their primary outcomes measured) and treatment effect heterogeneity? What is the resulting impact on power and average sample size in a group sequential design? We derive an asymptotic formula that decomposes the overall precision gain from adjusting for baseline variables and a short-term outcome into contributions from factors mentioned above, for efficient estimators in the model that only

*Department of Statistics, Harvard University. qiantianchen@fas.harvard.edu

[†]Department of Biostatistics, Johns Hopkins University.

[‡]Vatic Labs, New York.

assumes randomization and independent censoring. We use our formula to approximate the precision gain from a targeted minimum loss-based estimator applied to data from a completed trial of a new surgical intervention for stroke. Our formula implies that (for an efficient estimator) adjusting for a prognostic baseline variable leads to at least as much asymptotic precision gain as adjusting for an equally prognostic short-term outcome. In many cases, such as our stroke trial application, the former leads to substantially greater precision gains than the latter. In our simulation study, we show how precision gains from adjustment can be converted into sample size reductions (even when there is no treatment effect).

Keywords: Short-term outcome; Semiparametric efficiency; Targeted minimum loss-based estimator

1 Introduction

Group sequential designs for randomized clinical trials involve interim analyses that may result in early stopping for efficacy or futility. We consider trial designs where the primary outcome is measured at a fixed time (called the delay) after enrollment. Prior to observing the primary outcome, participants may have baseline variables and a short-term outcome measured. For example, in the MISTIE-II trial ([Hanley et al., 2016](#)) for evaluating a surgical treatment for intracerebral hemorrhage the primary outcome is the modified Rankin Scale (mRS), which measures degree of disability, 180 days after enrollment. A short-term outcome is mRS 30 days after enrollment. The baseline (pre-randomization) variables include age and measures of stroke severity.

Estimators that adjust for baseline variables are used in randomized trials because of the potential for increased precision and reduced sample size. [Pocock et al. \(2002\)](#), who surveyed 50 clinical trial reports from major medical journals, found that 36 adjusted for baseline variables. Adjusting for prognostic baseline variables is recommended by regulators, e.g., [FDA and EMA \(1998\)](#); [EMA \(2015\)](#); [FDA \(2019\)](#). Also, in order to address missing outcomes, it may be useful to appropriately adjust for baseline and post-randomization variables (e.g., short-term outcomes) ([National Research Council, 2010](#); [EMA, 2010](#)).

Various methods for adjustment in randomized trials have been developed; see, e.g., [Leon et al. \(2003\)](#); [Davidian et al. \(2005\)](#); [Tsiatis et al. \(2008\)](#); [Rubin and van der Laan \(2008\)](#); [Moore and van der Laan \(2009\)](#); [Moore et al. \(2011\)](#). There has also been discussion on the benefits and limitations of adjustment, compared to the standard, unadjusted estimator ([Yang and Tsiatis, 2001](#); [Freedman, 2008](#); [Lin et al., 2013](#)). [Moore et al. \(2011\)](#) present a formula for the precision gain from adjusting for baseline variables in single stage trials.

We focus on estimating the average treatment effect for continuous or binary outcomes. We derive a formula for the asymptotic precision gain (measured by the relative efficiency compared to the unadjusted estimator) due to adjustment for baseline variables and a short-term outcome using efficient estimators, that is, estimators that extract the maximum prognostic information from these variables. The formula can be used in trial planning to approximate the precision gain from adjustment, which can translate to reduced sample size.

We show, using data of the MISTIE-II trial, how a modified version of these formulas can be used to approximate the precision gain from using the targeted minimum loss-based estimator (TMLE) of [van der Laan and Gruber \(2012\)](#) while accounting for model misspecification.

We highlight a few implications of our formula. Holding other factors fixed, larger treatment effect heterogeneity decreases the value added from adjusting for baseline variables of all participants; however, for the pipeline participants, adjusting for their baseline variables leads to increased precision only when there is treatment effect heterogeneity. Adjusting for prognostic baseline variables typically leads to precision gains even if all participants have their primary outcomes observed. In contrast, adjusting for a prognostic short-term outcome can only improve precision when there are participants with the short-term outcome observed but the primary outcome unobserved.

In [Section 2](#), we introduce notation and assumptions. In [Section 3](#), we present our formula for precision gain. We show how to approximate the precision gain from using a TMLE in [Section 4](#). In [Section 5](#), we apply our formula to data from the MISTIE-II trial. In [Section 6](#), we conduct simulation studies. Limitations and open problems are discussed in [Section 7](#).

2 Notation and Assumptions

For participant i , let A_i denote the indicator of study arm assignment. We assume A_i is binary-valued with $A_i = 1$ for treatment and $A_i = 0$ for control. Denote by W_i a vector of baseline variables measured before randomization. All variables must be preplanned in the trial protocol. Denote by L_i a short-term outcome, which is observed at time d_L after enrollment. Though all of our results hold if L_i is any prespecified variable measured after randomization, not necessarily the outcome measured at an earlier time, we focus on the case of short-term outcomes. Also, for simplicity, we focus on cases with binary L_i ; the results can be generalized to multidimensional, continuous L_i . Denote by Y_i the primary outcome (continuous or binary-valued), which is observed at time d_Y after enrollment with $d_Y \geq d_L$. The delays d_L and d_Y are prespecified and are common to all participants. When followed up completely, participant i contributes full data (W_i, A_i, L_i, Y_i) .

Let n denote the sample size. We assume that the set of vectors $\{(W_i, A_i, L_i, Y_i)\}_{1 \leq i \leq n}$ are independent and identically distributed draws from the unknown joint distribution P_0 on generic data vector (W, A, L, Y) . The following assumption follows from randomization:

Assumption 1. The study arm assignment A is independent of the baseline variables W and $P(A = 1) = P(A = 0) = 1/2$.

Results for randomization probabilities other than $1/2$ are in Supplementary Material [H.5](#).

Define the average treatment effect as $\Delta = E(Y|A = 1) - E(Y|A = 0)$. The goal is to test the null hypothesis of no average treatment benefit: $H_0 : \Delta \leq 0$ versus $H_1 : \Delta > 0$.

We assume that participants are enrolled at a constant rate. Since not all participants have full data observed at an interim analysis, we introduce indicators C^L and C^Y to denote that L and Y , respectively, have been observed at or before a given analysis time. For a participant, $C^L = 1$ if and only if L is observed, and $C^Y = 1$ if and only if Y is observed. These variables get updated at each analysis. We make the following assumptions:

Assumption 2. (Independent Censoring) (C^L, C^Y) is independent of (W, A, L, Y) .

Assumption 3. (Monotone Censoring) $C^L \geq C^Y$ with probability 1.

An example where Assumptions [2](#) and [3](#) hold is when administrative censoring is the only source of missingness and there are no changes over time in the population who are enrolled; this is what we simulate in [Section 6](#). At any analysis time, each enrolled participant has one of the following missingness patterns:

- (i) $(C^L, C^Y) = (0, 0)$: a pipeline participant with only (W, A) observed;
- (ii) $(C^L, C^Y) = (1, 0)$: a pipeline participant with (W, A, L) observed;
- (iii) $(C^L, C^Y) = (1, 1)$: a participant with (W, A, L, Y) observed.

We assume a nonparametric model for the joint distribution of the variables (W, A, L, Y, C_L, C_Y) , except that we make Assumptions 1-3. The semiparametric efficiency results in the following sections are with respect to this model.

For a group sequential design with K stages, we consider the asymptotic setting where the maximum sample size goes to infinity such that at each interim analysis the proportions

of enrolled participants with final and short-term outcomes observed converge to constants p_y and p_l , respectively (where these limit proportions may differ by analysis time, but for conciseness we suppress their dependence on the analysis time). This asymptotic regime corresponds to fixing the delay times d_L, d_Y and analysis times, while increasing the enrollment rate. We only consider analysis times with $p_y > 0$.

The unadjusted estimator of Δ is the difference between the two arms of the sample mean of the primary outcome Y , using data from all participants with Y observed, i.e., $\sum_i A_i C_i^Y Y_i / \sum_i A_i C_i^Y - \sum_i (1 - A_i) C_i^Y Y_i / \sum_i (1 - A_i) C_i^Y$.

All estimators of the average treatment effect Δ that we consider are regular, asymptotically linear (RAL), defined, e.g., by [Bickel et al. \(1993\)](#). Any such estimator $\widehat{\Delta}$ of Δ is asymptotically normal, i.e., $\sqrt{n}(\widehat{\Delta} - \Delta) \xrightarrow{d} N(0, \sigma^2)$ as sample size n goes to infinity, where \xrightarrow{d} denotes convergence in distribution; the variance σ^2 is called the asymptotic variance of $\widehat{\Delta}$ and is denoted by $\text{AVar}(\widehat{\Delta})$. For example, the asymptotic variance of the unadjusted estimator is $2\{\text{Var}(Y|A=1) + \text{Var}(Y|A=0)\}/p_y$ at an analysis time where p_y fraction of the enrolled participants have their primary outcome observed. The asymptotic relative efficiency (ARE) between two RAL estimators $\widehat{\Delta}_1$ and $\widehat{\Delta}_2$ of Δ is the inverse of the ratio of their asymptotic variances: $\text{ARE}(\widehat{\Delta}_1, \widehat{\Delta}_2) = \text{AVar}(\widehat{\Delta}_2) / \text{AVar}(\widehat{\Delta}_1)$.

For any random vector X and $a \in \{0, 1\}$, let $E_a(Y | X)$ and $\text{Var}_a(Y | X)$ denote $E(Y | X, A = a)$ and $\text{Var}(Y | X, A = a)$, respectively. For example, $E_1(Y | W) = E(Y | W, A = 1)$, and $E_1(Y) = E(Y | A = 1)$. When no subscript is used on E or Var , these refer to expectation and variance, respectively, under the population distribution of the corresponding variables.

3 Best Possible Precision Gain

3.1 Formula for Precision Gain

The following result gives the best possible asymptotic variance that can be achieved by a RAL estimator of the average treatment effect Δ , in terms of the proportions p_l, p_y of enrolled participants with L and Y observed, respectively, at a given analysis time:

Lemma 1. Suppose Assumptions 1, 2, and 3 hold. The asymptotic variance of any RAL

estimator of Δ is at least

$$\begin{aligned} & \text{Var}\{E_1(Y | W) - E_0(Y | W)\} + \sum_{a \in \{0,1\}} \frac{2}{p_l} \text{Var}_a\{E_a(Y | L, W) - E_a(Y | W)\} \\ & + \sum_{a \in \{0,1\}} \frac{2}{p_y} \text{Var}_a\{Y - E_a(Y | L, W)\}. \end{aligned} \quad (1)$$

Lemma 1 is a consequence of the efficient influence function of Δ in our semiparametric model that only makes the assumptions in Section 2. This efficient influence function is given in Section H.5 of the Supplementary Material, and follows from Robins and Rotnitzky (1992); Scharfstein et al. (1999); van der Laan and Robins (2003). The TMLE estimator of van der Laan and Gruber (2012), which will be discussed in Section 4, achieves this variance lower bound when all working models are correct. When certain working models for this estimator are misspecified, one can approximate the corresponding ARE by substituting regression model fits for conditional expectations in (1); see Section 4.2.

The first term in (1) characterizes the variance in the conditional treatment effect across different levels of W . We define the *treatment effect heterogeneity*, denoted by γ , by dividing the variance of the conditional treatment effect by the sum of variances of Y in each arm:

$$\gamma = \frac{\text{Var}\{E_1(Y | W) - E_0(Y | W)\}}{\sum_{a \in \{0,1\}} \text{Var}_a(Y)}. \quad (2)$$

The treatment effect heterogeneity γ is invariant to linear transformations of Y , and is nonnegative. When $\gamma = 0$ there is no treatment effect heterogeneity, i.e., the conditional treatment effect $E_1(Y | W) - E_0(Y | W) = \Delta$ with probability 1.

For each arm $a \in \{0,1\}$, the variance of Y given $A = a$ can be decomposed as follows:

$$\text{Var}_a(Y) = \text{Var}_a\{Y - E_a(Y | L, W)\} + \text{Var}_a\{E_a(Y | L, W) - E_a(Y | W)\} + \text{Var}\{E_a(Y | W)\}, \quad (3)$$

as proved in Section H.2 in the Supplementary Material. The last term in the display above motivates the following definition of the proportion of the variance in Y explained by W (summed across arms):

$$R_W^2 = \frac{\sum_{a \in \{0,1\}} \text{Var}\{E_a(Y | W)\}}{\sum_{a \in \{0,1\}} \text{Var}_a(Y)} = 1 - \frac{\sum_{a \in \{0,1\}} \text{Var}_a\{Y - E_a(Y | W)\}}{\sum_{a \in \{0,1\}} \text{Var}_a(Y)}. \quad (4)$$

Similarly, the middle term in (3) motivates the following definition of the proportion of additional variance in Y explained by L after accounting for W (summed across arms):

$$R_{L|W}^2 = \frac{\sum_{a \in \{0,1\}} \text{Var}_a \{E_a(Y | L, W) - E_a(Y | W)\}}{\sum_{a \in \{0,1\}} \text{Var}_a(Y)}. \quad (5)$$

A RAL estimator of Δ is efficient (at a given data generating distribution) if it achieves the asymptotic variance lower bound (1) in the semiparametric model.

Theorem 2. Suppose Assumptions 1, 2, and 3 hold. Consider any analysis time. The ARE between any efficient RAL estimator of Δ and the unadjusted estimator is

$$\text{ARE} = \frac{1}{1 + (p_y/2)\gamma - R_W^2 - (1 - p_y/p_l)R_{L|W}^2}. \quad (6)$$


The denominator of the right side of the above display never exceeds 1, which follows since $\gamma \leq 2R_W^2$ by the Cauchy-Schwarz inequality.

3.2 Impact of Baseline Variables on Precision Gain

In order to isolate the impact of baseline variables W on the ARE (6), we consider the case where there is no impact of L , i.e., when $R_{L|W}^2 = 0$. This is the case, for example, when L is independent of Y given A and W . Then the ARE from adjusting for baseline variables W depends on three factors: the proportion of variance in Y explained by W (R_W^2), the proportion of participants with Y observed among those enrolled (p_y), and the treatment effect heterogeneity (γ). We plot in Figure 1 the ARE versus p_y , considering different combinations of R_W^2 and γ .

A summary of what is happening in Figure 1 is that with $p_y = 1$ and R_W^2 fixed, greater treatment effect heterogeneity (γ) lessens the precision gain from prognostic baseline variables. However, with fixed $R_W^2 > 0$ and $\gamma > 0$, some precision gain is restored as we decrease p_y from 1 to 0 (i.e., as we move from right to left along any line with $\gamma > 0$ in Figure 1). We describe the intuition for these phenomena below. The key idea is that precision gains result from adjusting for chance imbalances in W both (i) between study arms among the participants with Y observed, and (ii) between the full sample and the subset with Y observed. Whether the adjustments in (i) and (ii) lead to precision gains or not is determined

Figure 1: Asymptotic relative efficiency between an efficient estimator and the unadjusted estimator for estimating Δ , when only the baseline variable W is prognostic ($R_{L|W}^2 = 0$ in all curves).



by the treatment effect heterogeneity γ , which influences whether the estimated means in each arm are adjusted in opposite directions (leading to variance reduction) or in the same direction (leading to cancellation and no impact). This is analogous to constructive versus destructive interference between waves, which can lead to cancellation or amplification as the waves come together.

To give intuition for the phenomena in Figure 1, for the remainder of this subsection we fix R_W^2 and consider the simple case of a single, binary-valued W representing being less than 65 years old at baseline, and primary outcome Y being the indicator of having mRS at most 3 (a good outcome) at 180 days. We focus on the stratum $W = 1$ and assume that the under 65 subset of the study population would have better outcomes on average than the overall study population if all were assigned to the control arm, i.e., $E(Y|A=0, W=1) - E(Y|A=0) > 0$. We next consider $p_y = 1$ and the opposite extremes of no treatment effect heterogeneity (where efficiency gains from adjusting for W are maximal) and maximum treatment effect heterogeneity (where there are no efficiency gains from adjusting for W).

First, consider the case of $p_y = 1$ and no treatment effect heterogeneity ($\gamma = 0$), i.e., $E(Y|A=1, W) - E(Y|A=0, W)$ equals the constant $E(Y|A=1) - E(Y|A=0) = \Delta$. Combining this with the assumption above that those under 65 have better outcomes on average than the overall study population under assignment to the control arm ($E(Y|A=0, W=1) - E(Y|A=0) > 0$), it follows that those under 65 also have better outcomes on average than the overall study population under assignment to the *treatment* arm, i.e., $E(Y|A=1, W=1) - E(Y|A=1) = E(Y|A=0, W=1) - E(Y|A=0)$. If by chance there are proportionally more under 65 participants assigned to the treatment arm than the control arm, then the unadjusted estimator of the mean outcome in the treatment (control) arm is biased upward (downward) conditional on the chance imbalance (Jiang et al., 2019). An efficient estimator adjusts for this by decreasing the unadjusted estimate of $E(Y|A=1)$ and by increasing the unadjusted estimate of $E(Y|A=0)$; the net impact is to decrease the estimate of Δ . Alternatively, if by chance there are proportionally fewer under 65 participants assigned to the treatment arm, then by a symmetric argument an efficient estimator adjusts for this by increasing the unadjusted estimate of Δ . The overall impact of such adjustments across many hypothetical trials, is to remove the variance caused by chance imbalances across

arms in the stratum $W = 1$. This is, intuitively, why there is an efficiency gain when $\gamma = 0$.

Now consider $p_y = 1$ and the opposite extreme of maximum treatment effect heterogeneity ($\gamma = 2R_W^2$), which occurs when $E(Y|A = 1, W) + E(Y|A = 0, W)$ equals the constant $E(Y|A = 1) + E(Y|A = 0)$. An analogous argument (given in Section D of the Supplementary Material) to the previous paragraph shows that adjustment for chance imbalances in W leads to equal changes (in the same direction) in the estimated means in each arm; the net impact is that the difference between estimated means across arms is not changed (due to cancellation). Adjustment has no impact in this case, and leads to no variance reduction.

Above, we considered $p_y = 1$, i.e., no pipeline participants. We now fix $R_W^2 > 0$ and $\gamma > 0$, and consider the impact of decreasing p_y from 1 to 0. This leads to increased ARE because an efficient estimator can extract information from pipeline participants if $\gamma > 0$; in contrast, the unadjusted estimator does not use any information from pipeline participants. Intuition for why $\gamma > 0$ is necessary in order for pipeline participant information to be useful is given in Section D of the Supplementary Material.


3.3 Impact of a Short-term Outcome on Precision Gain

To isolate the impact of a short-term outcome L on precision gain, we set $R_W^2 = 0$ in (6) so that only L is prognostic. The precision gain from adjusting for short-term outcome L depends on two factors: the proportion of variance in Y explained by L after accounting for W ($R_{L|W}^2$), and the proportion of participants with Y observed among those with L observed (p_y/p_l). In Figure 2, we plot the ARE against p_y/p_l , and consider two values of $R_{L|W}^2$.

Smaller p_y/p_l and larger $R_{L|W}^2$ generally increase the precision gain from adjusting for L ; this is because L adds value just for participants with L but not Y observed, and $R_{L|W}^2$ quantifies the prognostic value of L beyond the variance in Y explained by W . When there are no pipeline participants, which would occur at the end of a trial with no early stopping or dropout, we have $p_y/p_l = 1$ and the ARE is 1, i.e., adjusting for L is useless. This is because adjusting for L helps by, in each arm separately, accounting for chance imbalances between the participants with Y observed and the participants with L but not Y observed; when $p_y/p_l = 1$, these groups are identical and no adjustment can be made.

Recall that we are operating in the semiparametric model defined by Assumptions 1-3; in

Figure 2: Asymptotic relative efficiency between an efficient estimator and the unadjusted estimator for estimating Δ , when only the short-term outcome L is prognostic ($R_W^2 = 0$ in both curves). The solid line corresponds to a higher prognostic value $R_{L|W}^2 = 0.25$ and the dashed line corresponds to a lower prognostic value $R_{L|W}^2 = 0.1$.



particular, we are not making any assumptions about the relationships among the variables (Y, A, L) . Therefore, we cannot adjust for observed imbalances in L between arms $A \in \{0, 1\}$, since these imbalances may be due to the impact of arm assignment.

Given fixed $R_{L|W}^2 > 0$, larger p_y/p_l decreases the precision gain from adjusting for L , as seen in Figure 2. In a trial with constant enrollment rate and $d_L < d_Y$, the precision gain from adjusting for L typically attenuates at later stages of a group sequential design due to p_y/p_l being nondecreasing over time. In the case where missingness is only due to administrative censoring, p_y/p_l starts at 0 at time d_L and eventually becomes 1 after enrollment stops and enough time has elapsed for all enrolled participants to have Y observed.

3.4 Comparison of Equally Prognostic Baseline Variable and Short-term Outcome

For any $q: 0 < q \leq 1$, we compare the ARE between two cases: $R_W^2 = q, R_{L|W}^2 = 0$ (only baseline variable prognostic) and $R_W^2 = 0, R_{L|W}^2 = q$ (only short-term outcome prognostic). The ARE in the former case is larger or equal to that in the latter case. Equality occurs if and only if $p_l = 1$ and treatment effect heterogeneity in the first case is the maximum possible ($\gamma = 2R_W^2 = 2q$), as proved in Section H.8 of the Supplementary Material. Intuition for this is given in Section D of the Supplementary Material. The equality can be seen graphically in that the two lines in Figure 2 (corresponding to prognostic value 0.25 and 0.1) are identical to the two corresponding lines in Figure 1 with $\gamma = 2R_W^2$; the horizontal axes of the two figures are identical in the special case of $p_l = 1$.

4 Estimator that Adjusts for Baseline Variables and Short-term Outcomes

4.1 Targeted Minimum Loss-Based Estimator

The targeted minimum loss-based estimator (TMLE) of [van der Laan and Gruber \(2012\)](#), which builds on the ideas of [Robins and Rotnitzky \(1992\)](#); [Scharfstein et al. \(1999\)](#); [Robins](#)

(2000); Bang and Robins (2005); van der Laan and Rubin (2006), is implemented in the R package `ltmle` (Lendle et al., 2017). We use it to estimate Δ and call it the *adjusted estimator*.

Our implementation of this TMLE for a binary outcome Y involves fitting logistic regression working models. Such models are fit first for the conditional probabilities of censoring given the observed history before censoring: $P(C^L = 1 | A, W)$ and $P(C^Y = 1 | L, A, W)$. We also fit such a model for study arm assignment $P(A = 1 | W)$. Lastly, we fit logistic regression models for $E(Y|L, A, W)$ and for $E(Y|A, W)$, given the corresponding observed histories; we refer to these as outcome regression models. The last model is fit using sequential regression, an idea from Robins (2000); Bang and Robins (2005); see our R code (Github link given at end of paper) for the implementation of the software of Lendle et al. (2017) that we used. All models involve an intercept and main terms for each regressor variable. Each regression is performed using all participants for whom the relevant variables are uncensored. For a continuous-valued outcome Y , linear outcome regression models could be used or, if the outcome is bounded then it can be rescaled to the interval $[0, 1]$ and logistic regression can be used as described by Gruber and van der Laan (2010).

The adjusted estimator initially uses the aforementioned regression model fits. It then updates each outcome regression model by adding a new covariate built from the other model fits, as described by van der Laan and Gruber (2012). The final estimator of the average treatment effect Δ is based on the updated outcome regression model fit of $E(Y|A, W)$. It is computed by first generating a prediction $Y_{a,i}$ of the outcome Y under each hypothetical assignment to arm $a \in \{0, 1\}$ for each participant i by substituting $W = W_i$ and $A = a$ into this model fit. The estimator of Δ is $(1/n) \sum_{i=1}^n (Y_{1,i} - Y_{0,i})$.

Assumptions 1-3 from Section 2 imply that the censoring models and the model for $P(A = 1|W)$ are correctly specified. The results of van der Laan and Gruber (2012, Section 4) imply that the adjusted estimator is a consistent estimator of the average treatment effect Δ . This holds regardless of the correctness of the outcome regression models. That is, these models could be arbitrarily misspecified and still the adjusted estimator converges in probability to Δ as sample size goes to infinity. If the outcome regression models are correct, then the adjusted estimator achieves the asymptotic variance lower bound in (1), in which case it is semiparametric, locally efficient (van der Laan and Gruber, 2012, Section 4). Other

estimators with the above properties include the augmented, inverse probability of treatment (and censoring) estimators of [Robins \(2000\)](#) and [Bang and Robins \(2005\)](#).

4.2 Approximating the Relative Efficiency between the Adjusted Estimator and the Unadjusted Estimator

We focused on the asymptotic relative efficiency between an efficient estimator and the unadjusted estimator in Section 3. To connect our results regarding efficient estimators to the adjusted estimator, we use results from [van der Laan and Gruber \(2012, Section 5.3\)](#). These results imply that if all regression models used in the TMLE are correctly specified, then plugging the corresponding model fits into the formulas (2)-(5) and then using these estimates of $\gamma, R_W^2, R_{L|W}^2$ in (6) results in a consistent estimator of the ARE (called the plug-in estimator). This estimator of the ARE is asymptotically conservative, meaning that the AREs may be underestimated but not overestimated, if the models for censoring and for $P(A=1|W)$ are correctly specified; since these models are correct under Assumptions 1-3, it follows that the plug-in estimator can serve as a (possibly conservative) asymptotic approximation to the ARE between the adjusted estimator and the unadjusted estimator.

Denote by $\widehat{E}(Y | X = x)$ the predicted value of $E(Y | X = x)$ from a model fit. Let $\widehat{\text{Var}}(\cdot)$ denote the sample variance over participants from both arms, and $\widehat{\text{Var}}_a(\cdot)$ denote the sample variance over participants from arm a . We estimate $R_W^2, R_{L|W}^2$, and γ , respectively, by

$$\begin{aligned}\widehat{R}_W^2 &= \frac{\sum_{a \in \{0,1\}} \widehat{\text{Var}}\{\widehat{E}(Y | W, A = a)\}}{\sum_{a \in \{0,1\}} \widehat{\text{Var}}_a(Y)}, \\ \widehat{R}_{L|W}^2 &= \frac{\sum_{a \in \{0,1\}} \widehat{\text{Var}}_a\{\widehat{E}(Y | W, L, A = a) - \widehat{E}(Y | W, A = a)\}}{\sum_{a \in \{0,1\}} \widehat{\text{Var}}_a(Y)}, \\ \widehat{\gamma} &= \frac{\widehat{\text{Var}}\{\widehat{E}(Y | W, A = 1) - \widehat{E}(Y | W, A = 0)\}}{\sum_{a \in \{0,1\}} \widehat{\text{Var}}_a(Y)}.\end{aligned}$$

In Sections 5-6 when calculating the conditional expectations \widehat{E} in the formulas for \widehat{R}_W^2 and $\widehat{R}_{L|W}^2$, we used logistic regression model fits with main terms. When calculating $\widehat{\gamma}$, we used similar model fits except including interactions, since we wanted to give an opportunity to detect treatment effect heterogeneity. The estimated R -squared quantities above are relevant

for continuous and binary outcomes in measuring the precision gain from adjustment when estimating Δ ; this generalizes an idea from [Moore and van der Laan \(2009\)](#), who give an R-squared formula related to \widehat{R}_W^2 except that it marginalizes over A .

5 Applying Precision Gain Formula to Trial Data

5.1 Trial Example: MISTIE-II

MISTIE-II is a Phase II randomized clinical trial evaluating a new surgical treatment for intracerebral hemorrhage. The treatment is called Minimally-Invasive Surgery Plus rt-PA for Intracerebral Hemorrhage Evacuation, abbreviated as MISTIE ([Hanley et al., 2016](#)). In the MISTIE-II dataset, the primary and short-term outcomes of each participant correspond to the Modified Rankin Scale (mRS), measured at different times after enrollment. We use the variables in [Table 1](#). The primary outcome Y is a binary indicator of a successful outcome ($\text{mRS} \leq 3$) at 180 days after enrollment. The short-term outcomes $L^{(1)}$ and $L^{(2)}$ are the indicators of mRS no larger than 4 at 30 and 90 days after enrollment, respectively. The cutoff 4 for $L^{(1)}$ and $L^{(2)}$ was chosen because there were relatively few participants with mRS 3 or less at 30 or 90 days after enrollment. The treatment assignment indicator A denotes whether a participant is assigned to the surgical treatment ($A=1$) or to standard of care ($A=0$). Baseline variables $W^{(1)}$ - $W^{(4)}$ are age (dichotomized at 65 years), NIHSS (NIH Stroke Scale total score, quantifying stroke-caused impairment), ICH (intracerebral hemorrhage volume), and GCS (Glasgow Coma Scale), all measured before randomization. The baseline variables (except for age) are treated as continuous variables in the regression models. The dataset has 100 participants.

We let $W^{\text{full}} = (W^{(1)}, W^{(2)}, W^{(3)}, W^{(4)})$ and $L^{\text{full}} = (L^{(1)}, L^{(2)})$. These variables are used in constructing data generating distributions for our simulations in [Section 6](#). However, only the smaller subsets of variables $W = (W^{(1)}, W^{(4)})$ and $L = L^{(1)}$ are made available to the adjusted estimator (both in [Section 5.2](#) and [6](#)). We included the “extra” variables $W^{(2)}, W^{(3)}, L^{(2)}$ (which the adjusted estimator does not have access to) in our data generating distributions for simulations so that the regression models used by the adjusted estimator

Table 1: Variables that we use from the MISTIE-II dataset. Ordinal-valued baseline variables below are treated as continuous in our regression models.

Name	Description
$W^{(1)}$	baseline age (years), dichotomized at ≤ 65
$W^{(2)}$	baseline NIHSS, ordinal
$W^{(3)}$	baseline ICH, ordinal
$W^{(4)}$	baseline GCS, ordinal
A	treatment indicator, 1 being MISTIE
$L^{(1)}$	mRS at 30 days, dichotomized by thresholding at ≤ 4
$L^{(2)}$	mRS at 90 days, dichotomized by thresholding at ≤ 4
Y	mRS at 180 days, dichotomized by thresholding at ≤ 3

will be misspecified; this was done since we expect at least some misspecification to occur in practice.

5.2 Approximate Precision Gain by Using the Adjusted Estimator on MISTIE-II Data

Applying the method in Section 4.2 to the MISTIE-II data, the estimated R-squared quantities are $\widehat{R}_W^2 = 0.36$, $\widehat{R}_{L|W}^2 = 0.08$, and $\widehat{\gamma} = 0.02$. This indicates that the baseline variables are strongly prognostic, that (after accounting for the baseline variables) the short-term outcome L is only mildly prognostic, and that there is little (if any) treatment effect heterogeneity.

At the end of a trial, if every participant has Y observed, we have $p_l = p_y = 1$. Substituting these and the estimated values $\widehat{R}_W^2, \widehat{R}_{L|W}^2, \widehat{\gamma}$ into (6) gives an approximate ARE of 1.53 comparing the TMLE versus the unadjusted estimator. If instead there had been no treatment effect heterogeneity ($\widehat{\gamma} = 0$), then the approximate ARE would be slightly larger (1.55), assuming \widehat{R}_W^2 and $\widehat{R}_{L|W}^2$ remain fixed. Now consider a hypothetical interim analysis where 80% of the enrolled participants have Y observed and 95% of the enrolled participants have L observed (e.g., at the 3rd interim analysis of the group sequential design with $n_{\max} = 480$ considered in Section 6.3; see Table B.1 in the Supplementary Material). Then

we have $p_y = 0.8$, $p_l = 0.95$, and the approximate ARE is 1.56.

Equivalently, in each of the three aforementioned situations, the adjusted estimator approximately reduces the required sample size to achieve a desired power by 35%, 36%, and 36%, respectively, compared to the unadjusted estimator. (The asymptotic sample size reduction is computed as $1 - 1/\text{ARE}$ as described in Section F of the Supplementary Material.) These are all quite similar, which is primarily due to the relatively small values of $\hat{\gamma}$ and $\hat{R}_{L|W}^2$. The magnitudes of the AREs above are not unusual in the MISTIE-II population—see [Optimising the Analysis of Stroke Trials \(OAST\) Collaboration \(2009\)](#)—and result from some baseline variables (especially NIHSS) being strongly prognostic for the outcome.

Approximations of the ARE as above can be applied to trial planning, where one may use previous data to estimate the variance explained by baseline variables and short-term outcomes and the treatment effect heterogeneity, and use (6) to get a rough projection of the precision gain from adjusting for baseline variables and short-term outcomes. If treatment effect heterogeneity is difficult to estimate *a priori*, one could use $\gamma = 0$ as the best case scenario (no treatment effect heterogeneity, corresponding to maximal precision gain), with the expectation that $\gamma > 0$ will result in attenuated precision gain.

A conservative approach is to assume no precision gain when selecting the initial sample size per stage, but then use preplanned sample-size re-estimation (using the estimated variance) to potentially shrink the per-stage sample size to reflect the estimated precision gains based on accrued data. Such adaptations based on estimating a nuisance parameter are generally acceptable to regulators such as the U.S. Food and Drug Administration ([FDA, 2016](#); [FDA, 2018](#)). This approach has important limitations given in Section 7.

6 Simulations of Group Sequential Design

6.1 Data Generating Distributions for Simulated Trials Based on MISTIE-II

We conducted simulation studies to assess how good of an approximation is provided by the precision gain formula (6), and to examine the sample size reduction by using the adjusted

versus unadjusted estimator in a group sequential design. In order to mimic key features of the MISTIE-II trial, we use a resampling-based algorithm to generate participants for our simulated trials. This algorithm generates simulated trials with the following properties:

- (i) The treatment assignment is independent of baseline variables.
- (ii) The relative efficiency between the adjusted estimator and the unadjusted estimator is similar to that calculated directly from the MISTIE-II dataset.

We define eight data generating distributions, called “settings”. The goal is to consider four situations involving W and/or L being prognostic or not for Y ; these are denoted by $progn_{W,L}$ (both prognostic), $progn_W$ (only W prognostic), $progn_L$ (only L prognostic), and $progn_\emptyset$ (neither prognostic). For each of these four situations (called “prognostic settings”), we construct distributions with the following two average treatment effects: $\Delta = 0$ (no effect) or $\Delta = 0.122$ (benefit equal to the unadjusted estimate from the MISTIE-II dataset). Details of the data generating algorithm for the eight settings are in Section C of the Supplementary Material. Table 2 gives the approximate R -squared and γ values for the eight settings. The only source of missing data in our simulations is administrative censoring, i.e., participants who enrolled but have not been in the trial long enough to have outcomes measured.

Table 2: Approximate R -squared and γ for each of the eight data generating distributions (settings) in our simulation study, computed by simulating 10^6 participants under each distribution and computing $\widehat{R}_W^2, \widehat{R}_{L|W}^2, \widehat{\gamma}$, respectively.

	progn _{W,L}		progn _W		progn _L		progn _∅	
	$\Delta = 0$	$\Delta = 0.122$	$\Delta = 0$	$\Delta = 0.122$	$\Delta = 0$	$\Delta = 0.122$	$\Delta = 0$	$\Delta = 0.122$
R_W^2	0.35	0.36	0.35	0.36	0	0	0	0
$R_{L W}^2$	0.08	0.07	0	0	0.30	0.30	0	0
γ	0	0.01	0	0.01	0	0	0	0

6.2 Group Sequential Trial Design Analysis Timing and Early Stopping Rule

We assume that $d_L = 30$ days and $d_Y = 180$ days. The goal is to control Type I error at level 0.025 and to have at least 80% power to reject the null hypothesis $H_0 : \Delta \leq 0$ when $\Delta = 0.122$.

We used the group sequential design framework in [Hampson and Jennison \(2013\)](#), which involves interim analyses where a choice is made to stop or continue enrollment; if enrollment is stopped, then one waits until all pipeline participants complete the trial and then a hypothesis test (called a decision analysis) is conducted. Full details of the design are in [Supplementary Material B.1](#).

Denote by K the total number of stages. We set $K = 5$ stages in our simulated trials. Participants were enrolled at the constant rate of 140 participants per year based on the projection for the enrollment rate in the MISTIE-III trial ([Hanley, 2016](#)). Given the maximum sample size n_{\max} , the timing of interim analyses is chosen such that for $1 \leq k \leq 4$, at the k -th interim analysis there are $(k/K)n_{\max}$ participants with Y observed.

For either estimator the corresponding Wald statistics at different analysis times have (asymptotically) a multivariate normal distribution, where the information at a given analysis time is the reciprocal of the estimator variance. (This relies on Assumptions 1-3 and holds regardless of the correctness of the outcome regression models used in the adjusted estimator.) The covariance matrix can be estimated, e.g., using the nonparametric bootstrap. To reduce the computational burden in our simulation studies, for any data generating distribution and n_{\max} we precompute the covariance matrix of our statistics at different analysis times, which is then used to determine stopping boundaries. The efficacy/futility stopping boundaries at each stage are calculated using the error spending approach presented in [Supplementary Material B.2](#). We used error spending functions $f(t) = 0.025\min(t^2, 1)$ for Type I error and $g(t) = 0.2\min(t^2, 1)$ for Type II error, where t denotes the information time (observed information divided by maximum information).

For each combination of prognostic setting ($\text{progn}_{W,L}, \text{progn}_W, \text{progn}_L, \text{progn}_\emptyset$) and estimator (unadjusted or adjusted), we use binary search to find the minimum total sample size n_{\max} required to achieve approximately 80% power under $\Delta = 0.122$ at 0.025 Type I error us-

ing the group sequential design. In practice, this could be approximated from accruing data by estimating the variance of the adjusted estimator and then using sample size re-estimation to set n_{\max} . However, we assumed that this was known in order to reduce the computational burden. For the unadjusted estimator, in every setting we let n_{\max} be 480. For the adjusted estimator, n_{\max} is set to be 300 under all settings with $\text{progn}_{W,L}$ and progn_W , and 480 under all settings with progn_L and progn_\emptyset . We simulated 50,000 trials under each of the eight settings. The number of accrued participants, error spending functions, and the stopping boundaries for each analysis are listed in Table B.2 in the Supplementary Material.

6.3 Simulation Results: Sample Size Reduction from Adjustment

Table 3 lists the simulation-based n_{\max} , Type I error (obtained under $\Delta = 0$), power (obtained under $\Delta = 0.122$), and the expected sample size (ESS) for each of the eight settings for each estimator. For each setting, the empirical Type I error rate is controlled at 0.025 and the power is approximately 80%. ESS is calculated as the number of enrolled participants when the trial stops, averaging over 50,000 simulations. The performance of the unadjusted estimator is the same under all prognostic settings and is summarized in the first row.

Compared to the unadjusted estimator, the adjusted estimator substantially reduces the sample size when the baseline variables are prognostic. In particular, comparing the first row (unadjusted) versus the third row (adjusted progn_W), the adjusted estimator reduces the maximum sample size from 480 to 300, the expected sample size under $\Delta = 0$ from 318 to 227, and the expected sample size under $\Delta = 0.122$ from 382 to 260. The sample size reduction due to the prognostic short-term outcome is very small. When neither W nor L is prognostic (progn_\emptyset), performance of the two estimators is similar.

6.4 Simulation Results: Relative Efficiency

In this subsection, we simulated trials with no early stopping, i.e., each simulated trial always enrolls n_{\max} participants and continues follow-up until all participants have Y observed. We then look back at each simulated trial and compute what each estimator’s value would be at each of the interim and decision analysis times. At any analysis time, the number of

Table 3: The maximum sample size (n_{\max}) and empirical Type I error, power, and expected sample size (ESS) under $\Delta = 0$ and $\Delta = 0.122$ for each estimator under each prognostic setting. The n_{\max} was chosen for each prognostic setting in order to achieve approximately 80% power under $\Delta = 0.122$.

Estimator	Progn. set.	n_{\max}	Type I error	Power	ESS($\Delta = 0$)	ESS($\Delta = 0.122$)
unadjusted	-	480	0.0250	0.811	318	382
adjusted	progn $_{W,L}$	300	0.0254	0.791	225	259
	progn $_W$	300	0.0256	0.805	227	260
	progn $_L$	480	0.0253	0.805	309	375
	progn $_{\emptyset}$	480	0.0248	0.811	321	384

participants with Y observed, L observed, and only (A,W) observed, respectively, is fixed; therefore, we can make a direct comparison between the ratio of estimator variances in our simulations and the predicted ARE from our formula (which assumes constant p_l, p_y).

In order to make direct comparisons between estimators at the same sample size, we change (just for this subsection) the setting of n_{\max} . The maximum sample size n_{\max} is set to be 300 under progn $_{W,L}$ and progn $_W$, and is set to 480 under progn $_L$ and progn $_{\emptyset}$. Under each setting the same n_{\max} is used for both the unadjusted estimator and the adjusted estimator.

Table 4 lists the approximate, asymptotic relative efficiency (ARE) computed by evaluating (6) but with conditional expectations replaced by empirical estimates as discussed in Section 4.2, and the relative efficiency (RE) from simulation. Under most settings, the ARE predicted by the theory is similar to the RE computed from the simulation. There are some discrepancies between these at the earlier analysis times. We think that this is due to the relatively small sample sizes at these analysis times compared to the number of variables adjusted for, which can lead to model overfit and so reduced performance of the adjusted estimator. For example, at interim analysis 1 under progn $_{W,L}$ there are only 60 participants with the primary outcome observed but 4 variables get adjusted for in the regression model

for $E(Y|W,A,L)$; this violates the rule of thumb that one should have at least 20 observations per term in the regression model. We discuss possible remedies for this in Section 7.

A comparison between progn_W (when only W is prognostic) and progn_L (when only L is prognostic) shows that there is a much larger precision gain in the former case. This is despite the fact that W and L have roughly similar prognostic values for Y marginally, as shown in Table 2. The reason, as discussed in Sections 3.2-3.4, is that there is relatively little treatment effect heterogeneity, which means that the baseline variables contribute to precision gains in two ways: allowing adjustment for chance imbalances (i) between study arms among the participants with Y observed, and (ii) between the full sample and the subset with Y observed. The short-term outcome can only contribute in the second way, and so may lead to smaller gains. Under progn_L , comparing interim analyses to decision analyses shows that L improves estimation precision only at interim analyses (i.e., when there are pipeline participants), which is in line with our theoretical results.

7 Discussion

We considered independent censoring, and the only source of missing data in our simulations was administrative censoring. If outcomes are missing at random (Rubin, 1976), then the adjusted estimator may still be consistent under correct specification of certain working models, while the unadjusted estimator may be biased. If dropout is missing not at random (Rubin, 1976), then both estimators may be inconsistent and our formulas will not work.

The theoretical ARE can be larger than the simulation-based ARE, especially at the early interim analyses. One reason is that in practice there is a finite sample penalty for adjustment (for each degree of freedom in the models fit) that is not reflected in the theoretical calculations, and that dissipates as sample size grows. This issue may be mitigated by only scheduling interim analyses for times when at least 50% of the participants have primary outcomes observed, and by constraining the number of variables adjusted for such that one has at least 20 participants with primary outcomes observed per variable.

The sample size reductions due to improved precision from adjustment came primarily from selecting smaller n_{\max} . Since the prognostic value of variables is typically not precisely

Table 4: Comparison of the asymptotic relative efficiency (ARE) predicted by Theorem 2 to the relative efficiency (RE) from simulated trials, at each analysis time. Top half is under $\Delta = 0$; bottom half is under $\Delta = 0.122$. The maximum sample size n_{\max} is set to be 300 under $\text{progn}_{W,L}$ and progn_W , and is set to 480 under progn_L and progn_\emptyset . Under each setting the same n_{\max} is used for both the unadjusted estimator and the adjusted estimator. The simulated RE is computed from 50,000 simulated trials.

		ARE from theory				RE from Simulation			
		$\text{progn}_{W,L}$	progn_W	progn_L	progn_\emptyset	$\text{progn}_{W,L}$	progn_W	progn_L	progn_\emptyset
<i>under $\Delta = 0$</i>									
Interim Analysis	1	1.63	1.54	1.13	1.00	1.49	1.41	1.08	0.96
	2	1.59	1.54	1.07	1.00	1.54	1.49	1.06	0.98
	3	1.58	1.54	1.05	1.00	1.54	1.51	1.04	0.99
	4	1.57	1.54	1.04	1.00	1.54	1.51	1.04	0.99
Decision Analysis	1	1.53	1.53	1.00	1.00	1.49	1.51	0.99	0.99
	2	1.53	1.53	1.00	1.00	1.50	1.53	0.99	0.99
	3	1.53	1.53	1.00	1.00	1.51	1.52	0.99	0.99
	4	1.53	1.53	1.00	1.00	1.51	1.54	1.00	1.00
	5	1.53	1.53	1.00	1.00	1.51	1.54	1.00	1.00
<i>under $\Delta = 0.122$</i>									
Interim Analysis	1	1.64	1.56	1.12	1.00	1.49	1.43	1.08	0.96
	2	1.61	1.55	1.07	1.00	1.56	1.51	1.06	0.98
	3	1.59	1.55	1.05	1.00	1.57	1.53	1.04	0.99
	4	1.58	1.55	1.04	1.00	1.57	1.53	1.03	0.99
Decision Analysis	1	1.55	1.55	1.00	1.00	1.51	1.52	0.99	0.99
	2	1.55	1.55	1.00	1.00	1.53	1.53	0.99	0.99
	3	1.55	1.55	1.00	1.00	1.53	1.53	0.99	0.99
	4	1.55	1.55	1.00	1.00	1.54	1.54	1.00	1.00
	5	1.55	1.55	1.00	1.00	1.54	1.54	1.00	1.00

known before the trial starts, one may use a preplanned rule for sample size re-estimation to set n_{\max} based on accruing data. First, the original, maximum sample size is set conservatively assuming no gains from adjustment. During the trial the asymptotic variance of the adjusted estimator is estimated and n_{\max} is set to be the sample size for which the desired power is achieved when Δ equals the a priori specified, clinically meaningful, minimum treatment effect. The sample size re-estimation should only be conducted after a sufficient number of participants have had their primary outcomes measured, e.g., when roughly 50% of the originally planned sample size have this measured.

This approach is only feasible if the delay time d_Y is not too long compared to the enrollment rate (otherwise the originally planned sample size will all be enrolled before a sufficient number of primary outcomes are observed). This was feasible in the MISTIE trial context with $d_Y = 180$ days and enrollment rate 140 per year, since sample size re-estimation could be done at 2.1 years. At this time, 230 (48%) of the originally planned sample size (480) have primary outcomes observed. So restricting to approximately 300 participants (the required n_{\max} when W is as prognostic as in the MISTIE-II data) is possible.

A limitation of our simulation study is that we precomputed n_{\max} for each setting, in order to save computation time. In practice, this would be calculated as a function of the estimated variance using accruing data in the trial. A future area of research is to run simulation studies to evaluate the resulting impact on power and sample size.

R code (R Core Team, 2018) for our simulations can be downloaded at https://github.com/tqian/gst_tmle.

Acknowledgments

This work was supported by the Patient-Centered Outcomes Research Institute (ME-1306-03198), the US Food and Drug Administration (HHSF223201400113C), and NIH grant UL1TR001079. This work is solely the responsibility of the authors and does not represent the views of the above agencies. We thank Mary Joy Argo for helpful comments.

References

- Bang, H. and Robins, J. M. (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics* **61**, 962–973.
- Bickel, P., Klaassen, C., Ritov, Y., and Wellner, J. (1993). *Efficient and Adaptive Estimation for Semiparametric Models*. Springer.
- Davidian, M., Tsiatis, A. A., and Leon, S. (2005). Semiparametric estimation of treatment effect in a pretest–posttest study with missing data. *Statistical Science* **20**, 261.
- EMA (2010). Guideline on missing data in confirmatory clinical trials. *Committee for Medicinal Products for Human Use (CHMP)*. EMA/CPMP/EWP/1776/99 .
- EMA (2015). European Medicines Agency Guideline on Adjustment for Baseline Covariates in Clinical Trials. Reference number EMA/CHMP/295050/2013.
- FDA (2016). Adaptive Designs for Medical Device Clinical Studies: Guidance for Industry. <https://www.fda.gov/downloads/medicaldevices/deviceregulationandguidance/guidancedocuments/ucm446729.pdf>.
- FDA (2018). Draft guidance for industry. Adaptive design clinical trials for drugs and biologics. <https://www.fda.gov/downloads/Drugs/Guidances/UCM201790.pdf>.
- FDA (2019). Adjusting for Covariates in Randomized Clinical Trials for Drugs and Biologics with Continuous Outcomes: Guidance for Industry. <https://www.fda.gov/media/123801/download>.
- FDA and EMA (1998). E9 statistical principles for clinical trials. *U.S. Food and Drug Administration: CDER/CBER. European Medicines Agency: CPMP/ICH/363/96* .
- Freedman, D. A. (2008). On regression adjustments to experimental data. *Advances in Applied Mathematics* **40**, 180–193.
- Gruber, S. and van der Laan, M. (2010). A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *International Journal of Biostatistics* **6**, 26.

- Hampson, L. V. and Jennison, C. (2013). Group sequential tests for delayed responses (with discussion). *Journal of the Royal Statistical Society: Series B* **75**, 3–54.
- Hanley, D. (2016). Minimally Invasive Surgery Plus Rt-PA for ICH evacuation Phase III (MISTIE III). In: ClinicalTrials.gov [Internet]. Bethesda (MD): National Library of Medicine (US). 2000- [cited 2016 September 30]. Available from <https://clinicaltrials.gov/ct2/show/NCT01827046>.
- Hanley, D. F., Thompson, R. E., Muschelli, J., Rosenblum, M., McBee, N., Lane, K., Bistran-Hall, A. J., Mayo, S. W., Keyl, P., Gandhi, D., et al. (2016). Safety and efficacy of minimally invasive surgery plus recombinant tissue plasminogen activator in intracerebral haemorrhage evacuation (MISTIE): a randomised, phase 2 trial. *The Lancet Neurology* **15**, 1228.
- Jennison, C. and Turnbull, B. W. (1999). *Group sequential methods with applications to clinical trials*. CRC Press.
- Jiang, F., Tian, L., Fu, H., Hasegawa, T., and Wei, L. (2019). Robust alternatives to ANCOVA for estimating the treatment effect via a randomized comparative study. *J Am Stat Assoc* pages 1–21.
- Lendle, S. D., Schwab, J., Petersen, M. L., and van der Laan, M. J. (2017). ltmle: An R package implementing targeted minimum loss-based estimation for longitudinal data. *Journal of Statistical Software* **81**, 1–21.
- Leon, S., Tsiatis, A. A., and Davidian, M. (2003). Semiparametric estimation of treatment effect in a pretest-posttest study. *Biometrics* **59**, 1046–1055.
- Lin, W. et al. (2013). Agnostic notes on regression adjustments to experimental data: Re-examining Freedman’s critique. *The Annals of Applied Statistics* **7**, 295–318.
- Moore, K. L., Neugebauer, R., Valappil, T., and Laan, M. J. (2011). Robust extraction of covariate information to improve estimation efficiency in randomized trials. *Stat Med* **30**, 2389–2408.

- Moore, K. L. and van der Laan, M. J. (2009). Covariate adjustment in randomized trials with binary outcomes: Targeted maximum likelihood estimation. *Stat Med* **28**, 39–64.
- National Research Council (2010). The Prevention and Treatment of Missing Data in Clinical Trials: Panel on Handling Missing Data in Clinical Trials. National Academies Press. <https://doi.org/10.17226/12955>.
- Optimising the Analysis of Stroke Trials (OAST) Collaboration (2009). Should stroke trials adjust functional outcome for baseline prognostic factors? *Stroke* **40**, 888–894.
- Pocock, S. J., Assmann, S. E., Enos, L. E., and Kasten, L. E. (2002). Subgroup analysis, covariate adjustment and baseline comparisons in clinical trial reporting: current practice and problems. *Stat Med* **21**, 2917–2930.
- R Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Robins, J. M. (2000). Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*, volume 1999, pages 6–10.
- Robins, J. M. and Rotnitzky, A. (1992). Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS Epidemiology*, pages 297–331. Springer.
- Rosenblum, M. and van der Laan, M. J. (2011). Simple examples of estimating causal effects using targeted maximum likelihood estimation. Available at http://works.bepress.com/mark_van_der_laan/223 .
- Rubin, D. B. (1976). Inference and missing data. *Biometrika* **63**, 581–592.
- Rubin, D. B. and van der Laan, M. J. (2008). Covariate adjustment for the intention-to-treat parameter with empirical efficiency maximization. *U.C. Berkeley Division of Biostatistics Working Paper Series*. Available at <https://biostats.bepress.com/ucbbiostat/paper229> .

- Scharfstein, D. O., Rotnitzky, A., and Robins, J. M. (1999). Adjusting for nonignorable drop-out using semiparametric nonresponse models. *J Am Stat Assoc* **94**, 1096–1120.
- Scharfstein, D. O., Tsiatis, A. A., and Robins, J. M. (1997). Semiparametric efficiency and its implication on the design and analysis of group-sequential studies. *J Am Stat Assoc* **92**, 1342–1350.
- Tsiatis, A. A., Davidian, M., Zhang, M., and Lu, X. (2008). Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Stat Med* **27**, 4658–4677.
- van der Laan, M. J. (2010). Targeted maximum likelihood based causal inference: Part I. *The International Journal of Biostatistics* **6**,.
- van der Laan, M. J. and Gruber, S. (2012). Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The International Journal of Biostatistics* **8**,.
- van der Laan, M. J. and Robins, J. M. (2003). *Unified methods for censored longitudinal data and causality*. Springer Science & Business Media.
- van der Laan, M. J. and Rubin, D. (2006). Targeted maximum likelihood learning. *The International Journal of Biostatistics* **2**,.
- van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- Yang, L. and Tsiatis, A. A. (2001). Efficiency study of estimators for a treatment effect in a pretest–posttest trial. *The American Statistician* **55**, 314–321.

Supplementary Material

In Section [A](#), we discuss the relationship between asymptotic relative efficiency and sample size reduction for a Wald test of a single stage trial. In Section [B](#), we present the group sequential design used in the simulation studies in Section [6](#). In Section [C](#), we specify the data generating distribution used in the simulation studies in Section [6](#). Intuition for the impact of treatment effect heterogeneity on precision gains is given in Section [D](#). In Section [E](#), we provide additional theoretical results regarding the precision gain from adjusting for prognostic baseline variables, when the baseline variable is either purely predictive or purely prognostic. In Section [F](#), we provide theoretical results regarding the precision gain from adjusting for prognostic baseline variables and the short-term outcome, when the parameter is $E(Y | A = a)$ for each $a \in \{0,1\}$ (rather than the average treatment effect). Simulations to evaluate the theory in Section [F](#) are provided in Section [G](#). Section [H](#) includes proofs of the results in the main paper as well as results in Section [E](#) and Section [F](#). Auxiliary lemmas that are used in the proofs in Section [H](#) are themselves proven in Section [I](#).

A Relationship between asymptotic relative efficiency and sample size reduction for a Wald test of a single stage trial

Consider one-sided Wald tests of the null hypothesis $H_0 : \Delta \leq 0$ versus the alternative $H_1 : \Delta > 0$, using RAL estimators $\hat{\Delta}_1$ and $\hat{\Delta}_2$, respectively (where the test statistics are the estimators divided by their standard errors). The asymptotic relative efficiency can be interpreted as the limit as sample size goes to infinity of the inverse of the ratio of the required sample sizes for the two estimators to achieve a given power at local alternatives ([van der Vaart, 2000](#), Section 8.2). For example, $\text{ARE}(\hat{\Delta}_1, \hat{\Delta}_2) = 1.2$ means that by using $\hat{\Delta}_1$ instead of $\hat{\Delta}_2$, the required sample size is reduced by $1 - 1/1.2 \approx 17\%$ asymptotically.

B Full Description of Group Sequential Design Used in Simulation

Components of the group sequential design used in simulation are presented in the following order: the multiple testing procedure (Section B.1), the computation of the test boundaries (Section B.2), the sample size at each analysis for the simulated trials (Section B.3), and the value of Type I error and Type II error spent and the testing boundary at each analysis (Section B.4).

B.1 Multiple Testing Procedure

In the simulation studies in Section 6 we use the group sequential test from [Hampson and Jennison \(2013\)](#) with K stages. Given $\alpha, \beta \in (0, 1)$, the design goal is to control the Type I error rate at level α and have power $1 - \beta$ at alternative $\Delta = \delta > 0$. Such a group sequential test can terminate enrollment at an interim analysis, and if such early stopping happens follow-up continues until all pipeline participants have Y observed before conducting a decision analysis to reject or accept H_0 . For each stage k , denote by S_k and \tilde{S}_k the test statistics at the k th interim analysis (where the decision to stop or continue enrollment occurs) and the k th decision analysis (where the hypothesis test is conducted), respectively; let u_k and l_k denote the efficacy and futility boundaries for terminating enrollment at interim analysis k , and let c_k denote the critical value for the corresponding decision analysis. These are used in the group sequential testing procedure below, reproduced from [Hampson and Jennison \(2013, Figure 1\)](#):

{	At interim analysis $k = 1, \dots, K - 2$,	
	if $S_k \leq l_k$ or $S_k \geq u_k$	stop recruitment and proceed to decision analysis k ;
	otherwise	continue recruitment and proceed to interim analysis $k + 1$.
{	At interim analysis $K - 1$,	
	if $S_{K-1} \leq l_{K-1}$ or $S_{K-1} \geq u_{K-1}$	stop recruitment and proceed to decision analysis $K - 1$;
	otherwise	complete recruitment and proceed to decision analysis K .
{	At decision analysis $k = 1, \dots, K$,	
	if $\tilde{S}_k \geq c_k$	reject H_0 ;
	if $\tilde{S}_k < c_k$	accept H_0 .

B.2 Computation of Test Boundaries

Following [Hampson and Jennison \(2013\)](#), consider a Type I error spending function $f(\cdot)$ and Type II error spending function $g(\cdot)$, which are non-decreasing with $f(0) = g(0) = 0$ and $f(t) = \alpha$ and $g(t) = \beta$ for $t \geq 1$. The maximum information level \mathcal{I}_{\max} is chosen depending on the power goal and the error spending functions. Denote by \mathcal{I}_k and $\tilde{\mathcal{I}}_k$ the information levels at the k -th interim analysis and decision analysis, respectively. Denote by $\mathcal{C}_k = (l_k, u_k)$ the critical region at interim analysis k , $1 \leq k \leq K$. The test boundaries u_k , l_k , and c_k are calculated by (12)-(15) in [Hampson and Jennison \(2013, Section 4.1.1\)](#), and we paraphrase as follows. Let u_1 and l_1 be the solutions to

$$P(S_1 \geq u_1; \Delta = 0) = f(\mathcal{I}_1/\mathcal{I}_{\max}) \quad \text{and} \quad P(S_1 \leq l_1; \Delta = \delta) = g(\mathcal{I}_1/\mathcal{I}_{\max}).$$

For $2 \leq k \leq K - 1$, u_k is the solution to

$$P(S_1 \in \mathcal{C}_1, \dots, S_{k-1} \in \mathcal{C}_{k-1}, S_k \geq u_k; \Delta = 0) = f(\mathcal{I}_k/\mathcal{I}_{\max}) - f(\mathcal{I}_{k-1}/\mathcal{I}_{\max}), \quad (\text{B.1})$$

and l_k is the solution to

$$P(S_1 \in \mathcal{C}_1, \dots, S_{k-1} \in \mathcal{C}_{k-1}, S_k \leq l_k; \Delta = \delta) = g(\mathcal{I}_k/\mathcal{I}_{\max}) - g(\mathcal{I}_{k-1}/\mathcal{I}_{\max}). \quad (\text{B.2})$$

For $1 \leq k \leq K-1$, the critical value c_k is the solution to

$$\begin{aligned} P(S_1 \in \mathcal{C}_1, \dots, S_{k-1} \in \mathcal{C}_{k-1}, S_k \geq u_k, \tilde{S}_k < c_k; \Delta = 0) \\ = P(S_1 \in \mathcal{C}_1, \dots, S_{k-1} \in \mathcal{C}_{k-1}, S_k \leq l_k, \tilde{S}_k \geq c_k; \Delta = 0). \end{aligned}$$

And the critical value c_K for the last stage is the solution to

$$P(S_1 \in \mathcal{C}_1, \dots, S_{K-1} \in \mathcal{C}_{K-1}, \tilde{S}_K \geq c_K; \Delta = 0) = \alpha - f(\mathcal{I}_{K-1}/\mathcal{I}_{\max}).$$

B.3 Sample Size at Each Analysis

Table B.1 lists the sample size and analysis timing of the group sequential designs with $n_{\max} = 480$ and $n_{\max} = 300$ used in Section 6.4. For $1 \leq k \leq 4$, Analysis k indicates interim analysis at stage k and k^* indicates the corresponding decision analysis if enrollment is early stopped at that stage. Analysis 5^* indicates the final decision analysis. There is not any interim analysis for the final stage. Fully observed participants are those with W, L, Y observed; partially observed participants are those with W, L but not Y observed. The three groups of participants listed in Table B.1 are inclusive of all enrollees and mutually exclusive.

B.4 Type I Error and Type II Error Spent and Testing Boundary

Table B.2 lists the Type I error per stage $f(\mathcal{I}_k/\mathcal{I}_{\max}) - f(\mathcal{I}_{k-1}/\mathcal{I}_{\max})$, Type II error per stage $g(\mathcal{I}_k/\mathcal{I}_{\max}) - g(\mathcal{I}_{k-1}/\mathcal{I}_{\max})$, and the testing boundaries at each stage under different settings used in the simulation studies. Given any prognostic setting $\text{progn}_{W,L}, \text{progn}_W, \text{progn}_L, \text{progn}_\emptyset$ and estimator pair, the designs are the same for $\Delta = 0$ and $\Delta = 0.122$.

C Resampling-based Algorithm to Simulate Trial Data

For the data generating distribution to have the two properties in Section 6.1, we design the algorithm to generate a simulated trial of n independent and identically distributed samples from the 100 participants in MISTIE-II dataset. Recall that the three properties are:

- (i) The treatment assignment is independent of baseline variables.

Table B.1: Analysis time and sample size at each interim and decision analysis for group sequential designs with $n_{\max} = 480$ and $n_{\max} = 300$. For $1 \leq k \leq 4$, Analysis k indicates interim analysis and k^* indicates the corresponding decision analysis if enrollment is early stopped. Analysis 5^* indicates the final decision analysis.

Analysis	1	1*	2	2*	3	3*	4	4*	5*
Design with $n_{\max} = 480$									
Time (year)	1.2	1.7	1.9	2.4	2.6	3.0	3.2	3.7	3.9
# Fully observed (W, L, Y)	96	165	192	261	288	357	384	453	480
# Partially observed (W, L) only	57	0	57	0	57	0	57	0	0
# Pipeline with only W observed	12	0	12	0	12	0	12	0	0
Design with $n_{\max} = 300$									
Time (year)	0.9	1.4	1.4	1.8	1.8	2.3	2.2	2.6	2.6
# Fully observed (W, L, Y)	60	129	120	189	180	249	240	300	300
# Partially observed (W, L) only	57	0	57	0	57	0	57	0	0
# Pipeline with only W observed	12	0	12	0	12	0	3	0	0

Table B.2: Type I error per stage $f(\mathcal{I}_k/\mathcal{I}_{\max}) - f(\mathcal{I}_{k-1}/\mathcal{I}_{\max})$, Type II error per stage $g(\mathcal{I}_k/\mathcal{I}_{\max}) - g(\mathcal{I}_{k-1}/\mathcal{I}_{\max})$, and boundaries for the designs in the different settings used in Section 6.3.

Analysis (k)	1	2	3	4	5
TMLE under progn _{W,L} , $n_{\max} = 300$					
Type I error per stage $\times 10^{-3}$	0.9	3.2	5.1	7.2	8.6
Type II error per stage $\times 10^{-3}$	7.5	25.4	41.1	57.3	68.6
Efficacy boundary at interim analysis (u_k)	3.11	2.71	2.47	2.27	-
Futility boundary at interim analysis (l_k)	-1.20	-0.08	0.70	1.36	-
Critical value at decision analysis (c_k)	1.30	1.54	1.74	1.91	2.07
TMLE under progn _{W} , $n_{\max} = 300$					
Type I error per stage $\times 10^{-3}$	0.8	2.9	5.0	7.0	9.3
Type II error per stage $\times 10^{-3}$	6.5	23.4	40.1	55.8	74.1
Efficacy boundary at interim analysis (u_k)	3.16	2.73	2.48	2.28	-
Futility boundary at interim analysis (l_k)	-1.26	-0.14	0.66	1.33	-
Critical value at decision analysis (c_k)	1.32	1.55	1.75	1.91	2.06
TMLE under progn _{L} , $n_{\max} = 480$					
Type I error per stage $\times 10^{-3}$	1.2	3.2	5.3	7.5	7.8
Type II error per stage $\times 10^{-3}$	9.3	25.8	42.6	60.2	62.1
Efficacy boundary at interim analysis (u_k)	3.05	2.68	2.44	2.24	-
Futility boundary at interim analysis (l_k)	-0.99	0.04	0.80	1.47	-
Critical value at decision analysis (c_k)	1.24	1.50	1.72	1.91	2.06
TMLE under progn _{\emptyset} , $n_{\max} = 480$					
Type I error per stage $\times 10^{-3}$	0.9	2.9	5.1	6.9	9.3
Type II error per stage $\times 10^{-3}$	7.3	22.9	40.6	54.9	74.3
Efficacy boundary at interim analysis (u_k)	3.12	2.73	2.47	2.28	-
Futility boundary at interim analysis (l_k)	-1.16	-0.08	0.72	1.37	-
Critical value at decision analysis (c_k)	1.24	1.50	1.72	1.91	2.04
Unadjusted estimator, $n_{\max} = 480$					
Type I error per stage $\times 10^{-3}$	1.0	3.0	4.9	6.9	9.2
Type II error per stage $\times 10^{-3}$	7.8	24.1	39.6	54.9	73.6
Efficacy boundary at interim analysis (u_k)	3.10	2.72	2.48	2.29	-
Futility boundary at interim analysis (l_k)	-1.37	-0.43	0.24	0.82	-
Critical value at decision analysis (c_k)	1.18	1.36	1.53	1.67	2.11

- (ii) The relative efficiency between the adjusted estimator and the unadjusted estimator in the simulated data is similar to that calculated directly from the MISTIE-II dataset.

Briefly, the algorithm ensures property (i) by adding a “twin” with identical baseline variables and opposite treatment assignment to each participant in the MISTIE-II dataset. The variables L^{full} and Y for each “twin” are then generated using regression model fits. The original MISTIE-II data and the set of “twins” results in a 200 participant data set. This is done once, before any of our simulations are conducted. To generate each simulated data set, we resample participant vectors from this 200 participant data set with replacement and then make modifications to the replicate data set depending on the desired simulation setting. Details of the algorithm are given below.

Step 1: Construct a set of 100 pairs of “twins”. We start with the 100-participant MISTIE-II dataset, and we augment the data with a hypothetical “twin” for each participant. A “twin” has identical baseline variables as the original participant, but opposite treatment assignment. We fit logistic regression models for $L^{(1)}$ on (W^{full}, A) , for $L^{(2)}$ on $(W^{\text{full}}, A, L^{(1)})$, and for Y on $(W^{\text{full}}, A, L^{(1)}, L^{(2)})$, using the original 100 participants in MISTIE-II dataset. The preliminary L_i and Y_i of each newly added twin are then predicted based on these logistic regression models by rounding the predicted success probability to be 0 or 1. The indicator of whether a participant is an original participant in the MISTIE-II data set or a hypothetical twin is included as a variable in this augmented data set of 200 participants, which will be used in Step 2. Step 1 is only done once at the beginning of the simulation.

Step 2: Sample n participants from the augmented data set with 200 participants and calibrate the treatment effect on the sampled participants. We sample n participants uniformly with replacement from the augmented dataset constructed in Step 1. Then, for each participant in the n participants whose “hypothetical twin indicator” is true, with probability 0.03 we reset its Y_i to equal A_i . This resetting step increases the treatment effect of the augmented data to 0.122, matching that of the original data.

Step 3: Calibrate relative efficiency between the adjusted estimator and the unadjusted estimator. For $a \in \{0, 1\}$, we empirically estimate the marginal distribution $p_{Y;a} = P(Y = 1 | A = a)$ using the corresponding sample proportions in the original data. Then, for each

participant in the n participants whose “hypothetical twin indicator” is true, with probability 0.164 we reset Y_i by a realization of an independent Bernoulli random draw with success probability $p_{Y;A_i}$. This resetting step adds random noise to reduce the prognostic value in W and L , so that the relative efficiency between the adjusted estimator and the unadjusted estimator for the simulated data is comparable to the estimated ARE from the original MISTIE-II dataset.

We then make modifications that are described next to assess validity of the theory under various settings. We define eight data generating distributions, called “settings”, each based on modifying the aforementioned data generating process. The goal is to consider four situations involving W and/or L being prognostic or not for Y ; these are denoted by $progn_{W,L}$ (both prognostic), $progn_W$ (only W prognostic), $progn_L$ (only L prognostic), and $progn_{\emptyset}$ (neither prognostic). For each of these four situations, we consider the following two average treatment effects: $\Delta = 0$ (no effect) or $\Delta = 0.122$ (benefit).

We first describe how we generated simulated trials with $\Delta = 0.122$. The aforementioned data generating algorithm corresponds to $progn_{W,L}$. To create data generating distributions with $progn_W$, we first generate each participant’s data as in $progn_{W,L}$ and then replace L by an independent draw from the marginal distribution in the MISTIE-II data. Similarly, in order to make only L prognostic for Y ($progn_L$), we first generate each participant’s data as in $progn_{W,L}$ and then replace W by an independent draw from the marginal distribution in the MISTIE-II data. Lastly, to make neither W nor L prognostic, we first generate each participant’s data as in $progn_{W,L}$ and then replace each of W and L by independent draws from the corresponding marginal distributions in the MISTIE-II data, respectively.

To generate participants under $\Delta = 0$, the data generating process above is followed, and then each study arm assignment A is replaced by an independent Bernoulli draw with probability 1/2. Under the data generating process with $\Delta = 0.122$, there is slight treatment effect heterogeneity $\gamma \approx 0.01$. By construction, under $\Delta = 0$ there is no treatment effect heterogeneity ($\gamma = 0$).

D Intuition for Impact of Treatment Effect Heterogeneity on Precision Gains

To give intuition for the phenomena in Figure 1 in Section 3.2, for the remainder of this appendix we fix R_W^2 and consider the simple case of a single, binary-valued W representing being less than 65 years old at baseline, and primary outcome Y being the indicator of having mRS at most 3 (a good outcome) at 180 days. We focus on the stratum $W = 1$ and assume that the under 65 subset of the study population would have better outcomes on average than the overall study population if all were assigned to the control arm, i.e., $E(Y|A=0, W=1) - E(Y|A=0) > 0$. We next consider $p_y = 1$ and the extreme of maximum treatment effect heterogeneity (where there are no efficiency gains from adjusting for W). (The opposite extreme of $p_y = 1$ and no treatment effect heterogeneity is discussed in detail in Section 3.2.)

Consider $p_y = 1$ and with maximum treatment effect heterogeneity ($\gamma = 2R_W^2$), which occurs when $E(Y|A=1, W) + E(Y|A=0, W)$ equals the constant $E(Y|A=1) + E(Y|A=0)$. It follows from the assumption above that those under 65 have better outcomes on average than the overall study population in the control arm ($E(Y|A=0, W=1) - E(Y|A=0) > 0$), that those under 65 have worse outcomes on average than the overall study population under assignment to the treatment arm (and by precisely the same magnitude), which follows since $E(Y|A=1, W=1) - E(Y|A=1) = -\{E(Y|A=0, W=1) - E(Y|A=0)\}$. If by chance there are proportionally more under 65 participants assigned to the treatment arm than the control arm, then the unadjusted estimator of the mean outcome in each arm is biased downward, and by the same amount, conditional on the chance imbalance. Since this bias cancels out when estimating Δ , an efficient estimator of Δ makes no adjustment. Alternatively, if by chance there are proportionally fewer under 65 participants assigned to the treatment arm, then by a symmetric argument an efficient estimator of Δ makes no adjustment. The overall impact is no adjustment due to chance imbalances in the stratum $W = 1$, which means that no improvement in variance is made compared to the unadjusted estimator. This is, intuitively, why there is no efficiency gain under maximum treatment effect heterogeneity ($\gamma = 2R_W^2$) when there are no pipeline participants ($p_y = 1$).

We next give intuition for why one needs $\gamma > 0$ in order for pipeline participants to contribute useful information toward estimating Δ (when $p_y < 1$). Consider the case where those in stratum $W = 1$ benefit more than the overall population, i.e., $E(Y|A = 1, W = 1) - E(Y|A = 0, W = 1) > \Delta$; this is only possible if $\gamma > 0$. Assume that adjustment for chance imbalance in the stratum $W = 1$ between arms has already been done to the unadjusted estimator (using only those with Y observed), as described above. Consider the proportion of participants having $W = 1$ among all participants and also among the subset of participants with Y observed. If the former proportion is larger than the latter, then the stratum $W = 1$ (of participants who benefit more than average) is underrepresented among those with outcomes observed. An efficient estimator adjusts the estimate of Δ upward to compensate. Alternatively, if the former proportion is smaller than the latter, by a symmetric argument an efficient estimator adjusts the estimate of Δ downward. In this way, variance due to imbalance in the proportion with $W = 1$ between the overall sample (including pipeline participants) and those with Y observed is removed. The only cases where adjusting for the aforementioned imbalance at each level of W has no impact are when $E(Y|A = 1, W) - E(Y|A = 0, W) = \Delta$, i.e., zero treatment effect heterogeneity, or $p_y = 1$ (no pipeline participants).

E Additional Results: When Baseline Variable is Purely Predictive or Purely Prognostic.

The following theorem examines the impact of treatment effect heterogeneity on precision gain from adjusting for prognostic baseline variables under two extreme cases. Define $\text{ARE}(\text{efficient}, \text{unadjusted})$ to be the asymptotic relative efficiency in (6). More generally, $\text{ARE}(\hat{\Delta}_1, \hat{\Delta}_2)$ denotes the asymptotic relative efficiency of the estimators $\hat{\Delta}_1, \hat{\Delta}_2$.

Corollary 1. Suppose Assumptions 1, 2, and 3 hold. In addition, assume $R_{L|W}^2 = 0$. For the estimand $E(Y | A = 1) - E(Y | A = 0)$, we have the following result regarding the asymptotic relative efficiency between an efficient RAL estimator and the unadjusted estimator.

- (i) (Maximal treatment effect heterogeneity.)

If $\text{Var}\{E_1(Y | W) - E_0(Y | W)\} > 0$ and $\text{Var}\{E(Y | W)\} = 0$, then $\gamma = 2R_W^2$ and

$$\text{ARE}(\text{efficient, unadjusted}) = \frac{1}{1 - (1 - p_y)R_W^2}. \quad (\text{E.1})$$

(ii) (No treatment effect heterogeneity.)

If $\text{Var}\{E_1(Y | W) - E_0(Y | W)\} = 0$ and $\text{Var}\{E(Y | W)\} > 0$, then $\gamma = 0$ and

$$\text{ARE}(\text{efficient, unadjusted}) = \frac{1}{1 - R_W^2}. \quad (\text{E.2})$$

We call a baseline variable W *purely predictive* if it satisfies the conditions in Corollary 1(i), as it predicts the conditional average treatment effect (i.e., it identifies which strata of W benefit from the treatment on average), but it does not explain the variance in Y marginally (i.e., margining out A). For such a baseline variable, $\text{ARE} = 1$ in (E.1) when $p_y = 1$. This means that adjusting for a purely predictive baseline variable results in no precision gain when there are no missing primary outcomes. An extreme case of this, for illustration only, is the following data generating distribution where the baseline variable is perfectly correlated with the outcome within each arm but still contributes nothing to the precision gain: W, A are independently distributed as Bernoulli(0.5) and $Y = \mathbb{1}(A = W)$, where $\mathbb{1}(\cdot)$ is the indicator function. In this case it is straightforward to verify that $R_W^2 = 1$, but adjusting for W would not impact asymptotic variance compared to the unadjusted estimator if all participants have Y observed ($p_y = 1$) since $\gamma = 2$. Though we don't expect to encounter a purely predictive variable in practice, the above thought experiment shows how W could explain treatment effect heterogeneity (which is useful on its own) while being useless for improving precision in estimating the average treatment effect.

We call a baseline variable W *purely prognostic* if it satisfies the conditions in Corollary 1(ii), as the treatment effect is constant across strata of W , and W explains some of the marginal variance in Y . For such a baseline variable, the ARE in (E.2) does not depend on p_y . An intuition for why this holds, as discussed in Section 3.2, is that adjusting for such a baseline variable reduces variance by correcting for chance imbalances between arms among those who have Y observed, while the baseline variables for the pipeline participants contain no information about Δ since $\gamma = 0$ (and we are ignoring L here).

F Precision Gain When Estimating the Treatment Specific Mean $E(Y|A = a)$

In this section we focus on estimating the population mean outcome under assignment to a single arm (called the treatment specific mean): $E(Y|A = a)$ for each $a \in \{0,1\}$ separately. The following theorem gives the lower bound on the asymptotic variance for all regular asymptotically linear (RAL) estimators of $E(Y|A = a)$.

Theorem F.1. Denote by $p_a = P(A = a)$ for $a \in \{0,1\}$. Assume $p_y > 0$ and $0 < p_a < 1$. Under Assumptions 1, 2, and 3 (except allowing $p_a \neq 1/2$) the semiparametric lower bound on the asymptotic variance of all RAL estimators for $E(Y|A = a)$ is

$$\text{Var}_a\{E_a(Y|W)\} + \frac{1}{p_a p_l} \text{Var}_a\{E_a(Y|L,W) - E_a(Y|W)\} + \frac{1}{p_a p_y} \text{Var}_a\{Y - E_a(Y|L,W)\}. \quad (\text{F.1})$$

Analogous to the R -squared quantities defined in Section 3, we define new R -squared quantities relevant to estimating the population mean of the primary outcome Y in each arm separately. For each $a \in \{0,1\}$, define the proportion of variance in Y explained by W in arm a as $R_{W;a}^2 = \text{Var}\{E_a(Y|W)\}/\text{Var}_a(Y)$; the proportion of additional variance in Y explained by L after accounting for W in arm a as $R_{L|W;a}^2 = \text{Var}_a\{E_a(Y|L,W) - E_a(Y|W)\}/\text{Var}_a(Y)$. We also will refer to the proportion of variance in Y explained by L alone in arm a , defined as $R_{L;a}^2 = \text{Var}_a\{E_a(Y|L)\}/\text{Var}_a(Y)$.

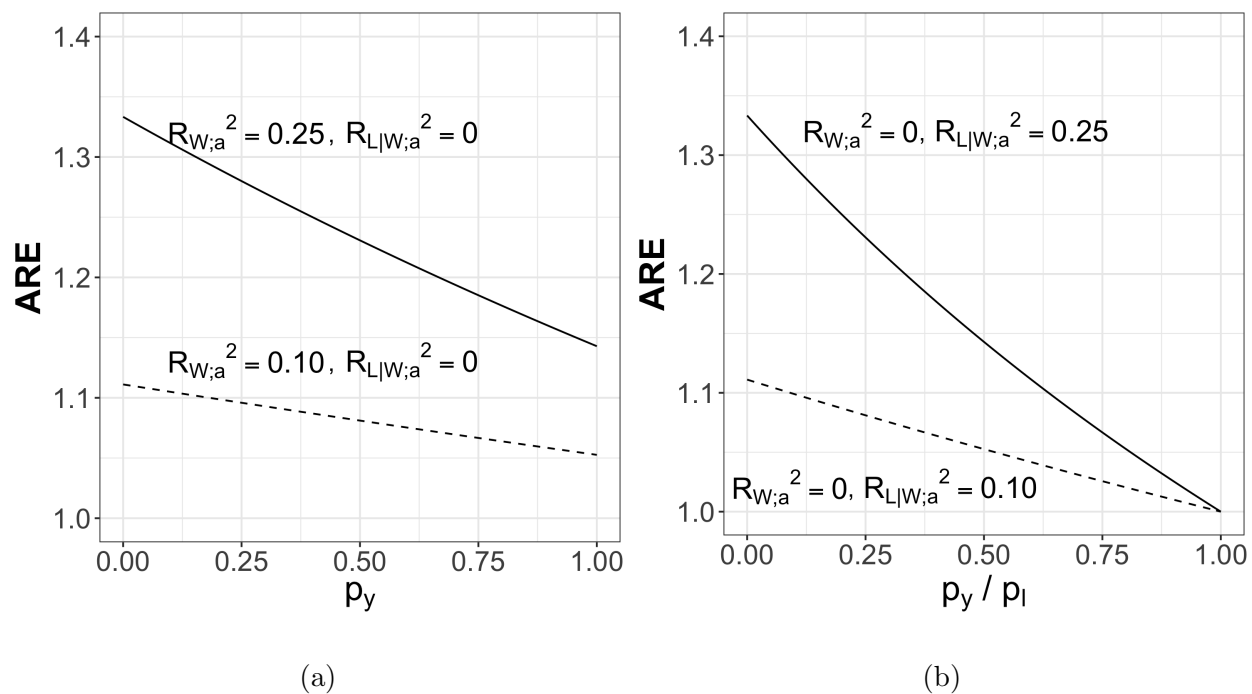
The following result gives the ARE between an efficient RAL estimator and the unadjusted estimator for $E(Y|A = a)$.

Theorem F.2. Assume all conditions in Theorem F.1 hold. For each arm $a \in \{0,1\}$, the asymptotic relative efficiency between an efficient RAL estimator and the unadjusted estimator of the treatment specific mean $E(Y|A = a)$ is

$$\frac{1}{1 - (1 - p_a p_y) R_{W;a}^2 - (1 - p_y/p_l) R_{L|W;a}^2}. \quad (\text{F.2})$$

When only W is prognostic ($R_{L|W;a}^2 = 0$), the asymptotic relative efficiency between an efficient RAL estimator and the unadjusted estimator is equal to when the L variable is not

Figure F.1: Asymptotic relative efficiency between an efficient RAL estimator and the unadjusted estimator for estimating $E(Y | A = a)$. Plot (a) is when only the baseline variable W is prognostic; plot (b) is when only the short-term outcome L is prognostic.



available to the efficient RAL estimator (i.e., when the efficient RAL estimator adjusts for W alone). Similarly, when only L is prognostic ($R_{W;a}^2 = 0$), the asymptotic relative efficiency is equal to when the W variable is not available to the efficient RAL estimator (i.e., when only adjusts for L alone). In Figure F.1a we illustrate the implication of Theorem F.2 under those special cases. We plot ARE against p_y when only W is prognostic ($R_{L|W;a}^2 = 0$), and in Figure F.1b we plot ARE against p_y/p_l when only L is prognostic ($R_{W;a}^2 = 0$). In each plot, we separately consider prognostic value being 0.1 and 0.25. We consider three implications of the curves in Figure F.1. First, the precision gain (ARE) from adjusting for a prognostic baseline variable increases with a greater proportion of pipeline participants (i.e., smaller p_y). The reason is that each additional pipeline participant (who has observed baseline variables but missing outcome) contributes some information to the adjusted estimator (through the baseline variables) but no information to the unadjusted estimator. For a similar reason, the precision gain from adjusting for a prognostic short-term outcome increases with a greater proportion of participants with W, L but not Y observed (i.e., smaller p_y/p_l).

Second, when every participant has their primary outcome observed (e.g., at the final analysis time of a group sequential design with no dropouts), adjusting for the prognostic baseline variables still improves estimation precision as long as $R_{W;a}^2 > 0$, but adjusting for prognostic short-term outcome no longer does.

Third, for any given (p_y, p_l) , adjusting for a prognostic baseline variable alone always leads to larger precision gain than adjusting for an equally prognostic short-term outcome alone.

Following the discussion in Section A, we define the asymptotic equivalent reduction in sample size (AERSS) of $\hat{\Delta}_1$ compared to $\hat{\Delta}_2$ as


$$\text{AERSS}(\hat{\Delta}_1, \hat{\Delta}_2) = 1 - \text{ARE}(\hat{\Delta}_1, \hat{\Delta}_2)^{-1}. \quad (\text{F.3})$$

Adjusting for prognostic W alone with $R_{W;a}^2 = q$ (for $0 < q < 1$) yields $\text{AERSS} = q(1 - p_y/2)$, and adjusting for prognostic L alone with $R_{L;a}^2 = q$ yields $\text{AERSS} = q(1 - p_y/p_l)$. The ratio of the two AERSSs equals (assuming $q > 0$)

$$r = \frac{\text{AERSS from adjusting only for prognostic } W \text{ with } R_{W;a}^2 = q}{\text{AERSS from adjusting only for prognostic } L \text{ with } R_{L;a}^2 = q} = \frac{1 - p_y/2}{1 - p_y/p_l}. \quad (\text{F.4})$$

$r > 1$ means that the sample size reduction from adjusting for a prognostic W is larger than that from adjusting for an equally prognostic L ; $r < 1$ means the opposite. Figure F.2 plots r against (p_l, p_y) under the constraint $0 < p_y \leq p_l \leq 1$. For all such (p_l, p_y) , $r > 1$. In addition, we have $r \geq 2$ if $p_y \geq 2/3$. This means that if at most 1/3 of the enrollees are in the pipeline, the sample size reduction from adjusting for a prognostic short-term outcome alone is at most half of that from adjusting for an equally prognostic baseline variable alone. Roughly speaking, small p_y makes r close to 1, meaning that adjusting for W or adjusting for L results in a similar sample size reduction when there are relatively few with Y observed. This may occur at early stages of a group sequential design, if the delay of the primary outcome is long relative to the enrollment rate.

Figure F.2: Contour plot of $r(p_l, p_y)$, where r is the ratio of the reductions in sample size when only W is prognostic with $R_{W;a}^2 = q$ and when only L is prognostic with $R_{L;a}^2 = q$, for any fixed $q \in (0, 1)$. The estimand is the treatment specific mean $E(Y | A = a)$.



G Simulation Results for Relative Efficiency When Estimating the Treatment Specific Mean

The simulation setup, including the data generating distributions (settings), are the same as Section 6. Table G.1 gives the values of $R_{W;a}^2$ and $R_{L|W;a}^2$ for $a \in \{0,1\}$, which are defined in the paragraph after Theorem F.1. Tables G.2 and G.3 list the asymptotic relative efficiency (ARE) approximated by substituting model-based estimates of the conditional expectations in the quantities in (F.2).

Table G.1: $R_{W;a}^2$ and $R_{L|W;a}^2$ for $a \in \{0,1\}$ for the data generating distributions from the settings in Section 6, approximated based on a simulated dataset with 1,000,000 participants under each setting.

	progn _{W,L}		progn _W		progn _L		progn _∅	
	Δ = 0	Δ = 0.122	Δ = 0	Δ = 0.122	Δ = 0	Δ = 0.122	Δ = 0	Δ = 0.122
$R_{W;0}^2$	0.35	0.33	0.35	0.33	0	0	0	0
$R_{L W;0}^2$	0.08	0.06	0	0	0.30	0.28	0	0
$R_{W;1}^2$	0.35	0.38	0.35	0.38	0	0	0	0
$R_{L W;1}^2$	0.08	0.08	0	0	0.30	0.31	0	0

H Proofs

We prove the results in the main paper as well as results in Section E and Section F. Auxiliary lemmas (used in these proofs) are proved Section I.

H.1 Identification of the average treatment effect

Under Assumptions 1-3, the average treatment effect can be expressed as follows:

$$\begin{aligned}
 E(Y | A = 1) - E(Y | A = 0) = & E[E\{E(Y | L, W, A = 1, C^Y = 1) | W, C^L = 1, A = 1\}] \\
 & - E[E\{E(Y | L, W, A = 0, C^Y = 1) | W, C^L = 1, A = 0\}] \quad (\text{H.1})
 \end{aligned}$$

Table G.2: Comparison of the asymptotic relative efficiency (ARE) predicted by Theorem F.2 and the relative efficiency computed from simulated trials under $\Delta = 0$. The maximum sample size n_{\max} is set to be 300 under $\text{progn}_{W,L}$ and progn_W , and is set to 480 under progn_L and progn_\emptyset . Under each setting the same n_{\max} is used for both the unadjusted estimator and the adjusted estimator. The simulated RE is based on 50,000 simulated trials.

Under $\Delta = 0$		ARE from theory				RE from Simulation			
		$\text{progn}_{W,L}$	progn_W	progn_L	progn_\emptyset	$\text{progn}_{W,L}$	progn_W	progn_L	progn_\emptyset
		<i>Estimand: $E(Y A = 0)$</i>							
Interim Analysis	1	1.44	1.36	1.13	1.00	1.34	1.29	1.09	0.97
	2	1.35	1.31	1.07	1.00	1.32	1.28	1.06	0.99
	3	1.32	1.29	1.05	1.00	1.29	1.27	1.04	0.99
	4	1.29	1.26	1.04	1.00	1.27	1.25	1.03	0.99
Decision Analysis	1	1.21	1.21	1.00	1.00	1.19	1.20	0.99	0.99
	2	1.21	1.21	1.00	1.00	1.19	1.20	0.99	0.99
	3	1.21	1.21	1.00	1.00	1.20	1.20	0.99	0.99
	4	1.21	1.21	1.00	1.00	1.20	1.21	1.00	1.00
	5	1.21	1.21	1.00	1.00	1.20	1.21	1.00	1.00
		<i>Estimand: $E(Y A = 1)$</i>							
Interim Analysis	1	1.44	1.37	1.13	1.00	1.34	1.28	1.10	0.97
	2	1.36	1.31	1.07	1.00	1.33	1.29	1.06	0.99
	3	1.32	1.29	1.05	1.00	1.31	1.27	1.04	0.99
	4	1.29	1.26	1.04	1.00	1.28	1.26	1.03	0.99
Decision Analysis	1	1.21	1.21	1.00	1.00	1.20	1.20	0.99	0.99
	2	1.21	1.21	1.00	1.00	1.20	1.21	0.99	0.99
	3	1.21	1.21	1.00	1.00	1.20	1.21	0.99	0.99
	4	1.21	1.21	1.00	1.00	1.21	1.21	1.00	1.00
	5	1.21	1.21	1.00	1.00	1.21	1.21	1.00	1.00

Table G.3: Comparison of the asymptotic relative efficiency (ARE) predicted by Theorem F.2 and the relative efficiency computed from simulated trials under $\Delta = 0.122$. The maximum sample size n_{\max} is set to be 300 under $\text{progn}_{W,L}$ and progn_W , and is set to 480 under progn_L and progn_\emptyset . Under each setting the same n_{\max} is used for both the unadjusted estimator and the adjusted estimator. The simulated RE is based on 50,000 simulated trials.

Under $\Delta = 0.122$		RE approximated by Theory				RE from Simulation			
		$\text{progn}_{W,L}$	progn_W	progn_L	progn_\emptyset	$\text{progn}_{W,L}$	progn_W	progn_L	progn_\emptyset
		<i>Estimand: $E(Y A = 0)$</i>							
Interim Analysis	1	1.59	1.54	1.14	1.00	1.49	1.43	1.10	0.97
	2	1.48	1.45	1.08	1.00	1.46	1.41	1.07	0.99
	3	1.43	1.41	1.06	1.00	1.44	1.40	1.05	0.99
	4	1.39	1.38	1.04	1.00	1.40	1.38	1.03	0.99
Decision Analysis	1	1.30	1.30	1.00	1.00	1.28	1.28	0.99	0.99
	2	1.30	1.30	1.00	1.00	1.29	1.29	0.99	0.99
	3	1.30	1.30	1.00	1.00	1.30	1.30	0.99	0.99
	4	1.30	1.30	1.00	1.00	1.30	1.30	1.00	1.00
	5	1.30	1.30	1.00	1.00	1.30	1.30	1.00	1.00
		<i>Estimand: $E(Y A = 1)$</i>							
Interim Analysis	1	1.37	1.30	1.11	1.00	1.26	1.21	1.08	0.97
	2	1.30	1.26	1.07	1.00	1.26	1.22	1.05	0.99
	3	1.27	1.24	1.05	1.00	1.24	1.21	1.04	0.99
	4	1.24	1.22	1.04	1.00	1.23	1.19	1.03	0.99
Decision Analysis	1	1.18	1.18	1.00	1.00	1.16	1.16	0.99	0.99
	2	1.18	1.18	1.00	1.00	1.16	1.16	0.99	0.99
	3	1.18	1.18	1.00	1.00	1.17	1.16	0.99	0.99
	4	1.18	1.18	1.00	1.00	1.17	1.16	1.00	1.00
	5	1.18	1.18	1.00	1.00	1.17	1.16	1.00	1.00

The proof of (H.1) is given as follows.

Proof. It suffices to show that for each $a \in \{0,1\}$ the following holds:

$$E(Y | A = a) = E[E\{E(Y | L, W, A = a, C^Y = 1) | W, A = a, C^L = 1\}]. \quad (\text{H.2})$$

Using the law of iterated expectation twice, we deduce

$$E(Y | A = a) = E[E\{E(Y | L, W, A = a) | W, A = a\}]. \quad (\text{H.3})$$

Let $f(L, W) = E(Y | L, W, A = a)$. By Assumption 2 we have

$$E(f(L, W) | A = a) = E(f(L, W) | A = a, C^L = 1), \quad (\text{H.4})$$

and

$$E(Y | L, W, A = a) = E(Y | L, W, A = a, C^L = 1, C^Y = 1). \quad (\text{H.5})$$

Equations (H.3)-(H.5) together yield

$$E(Y | A = a) = E[E\{E(Y | L, W, A = a, C^L = 1, C^Y = 1) | W, A = a, C^L = 1\}]. \quad (\text{H.6})$$

Equation (H.6) and Assumption 3 yield (H.2). This completes the proof. □

H.2 Lemma on Variance Decomposition

Lemma H.1. For study arm $a \in \{0,1\}$, we have the following decomposition of the variance of Y in that arm:

$$\text{Var}_a(Y) = \text{Var}_a\{E_a(Y | W)\} + \text{Var}_a\{E_a(Y | L, W) - E_a(Y | W)\} + \text{Var}_a\{Y - E_a(Y | L, W)\}. \quad (\text{H.7})$$

In addition, we have

$$R_{W;a}^2 + R_{L|W;a}^2 + R_{r;a}^2 = 1, \quad \text{for each } a \in \{0,1\}. \quad (\text{H.8})$$

Proof. Adding and subtracting terms, we have

$$\text{Var}_a(Y) = \text{Var}_a[\{Y - E_a(Y | L, W)\} + \{E_a(Y | L, W) - E_a(Y | W)\} + E_a(Y | W)]. \quad (\text{H.9})$$

So for proving (H.7), it suffices to establish the following:

$$\text{Cov}_a\{Y - E_a(Y | L, W), E_a(Y | L, W) - E_a(Y | W)\} = 0, \quad (\text{H.10})$$

$$\text{Cov}_a\{Y - E_a(Y | L, W), E_a(Y | W)\} = 0, \quad (\text{H.11})$$

$$\text{Cov}_a\{E_a(Y | L, W) - E_a(Y | W), E_a(Y | W)\} = 0, \quad (\text{H.12})$$

where Cov_a denotes the conditional covariance given $A = a$.

First, we have

$$\begin{aligned} & \text{Cov}_a\{Y - E_a(Y | L, W), E_a(Y | L, W) - E_a(Y | W)\} \\ &= E_a\{Y E_a(Y | L, W) - E_a(Y | L, W)^2 - Y E_a(Y | W) + E_a(Y | L, W) E_a(Y | W)\}. \end{aligned} \quad (\text{H.13})$$

By Lemma I.1 with $X = Z = (L, W)$, we have

$$E_a\{Y E_a(Y | L, W)\} = E_a\{E_a(Y | L, W)^2\}. \quad (\text{H.14})$$

By Lemma I.1 with $X = W$ and $Z = (L, W)$, we have

$$E_a\{Y E_a(Y | W)\} = E_a\{E_a(Y | L, W) E_a(Y | W)\}. \quad (\text{H.15})$$

Equations (H.13), (H.14), and (H.15) imply (H.10).

Second, we have

$$\text{Cov}_a\{Y - E_a(Y | L, W), E_a(Y | W)\} = E_a\{Y E_a(Y | W) - E_a(Y | L, W) E_a(Y | W)\}. \quad (\text{H.16})$$

Equations (H.16) and (H.15) imply (H.11).

Third, since $E_a(Y | L, W) - E_a(Y | W)$ has expectation zero, we have

$$\text{Cov}_a\{E_a(Y | L, W) - E_a(Y | W), E_a(Y | W)\} = E_a\{E_a(Y | L, W) E_a(Y | W) - E_a(Y | W)^2\}. \quad (\text{H.17})$$

In Lemma I.1, letting $X = Z = W$ and replacing Y in the lemma by $E(Y | L, W)$ implies

$$E_a\{E_a(Y | L, W) E_a(Y | W)\} = E_a\{E_a(Y | W)^2\}. \quad (\text{H.18})$$

Equations (H.17) and (H.18) imply (H.12).

This proves (H.7). Equation (H.8) follows immediately from (H.7) and the definition of $R_{W;a}^2$, $R_{L|W;a}^2$ and $R_{r;a}^2$. This completes the proof for Lemma H.1.

□

H.3 Proof of Theorem F.1

Proof. In the proof, we will use equation (H.6) derived earlier; we rewrite it below:

$$E(Y | A = a) = E[E\{E(Y | L, W, A = a, C^L = 1, C^Y = 1) | W, A = a, C^L = 1\}]. \quad (\text{H.19})$$

Treating the missingness indicators C^L and C^Y as binary treatments, the right-hand side of (H.19) becomes the average of outcome Y under time dependent treatment assignment: $A = a, C^L = 1, C^Y = 1$. Because there is no measurement made between A and C^L , we can combine the two as a single treatment \tilde{A} , with $\tilde{A} = 1$ if and only if $A = a$ and $C^L = 1$. Equation (H.19) becomes

$$E(Y | A = a) = E[E\{E(Y | L, W, \tilde{A} = 1, C^Y = 1) | W, \tilde{A} = 1\}]. \quad (\text{H.20})$$

Using the fact that L is binary-valued, by equations (24) and (28) in Rosenblum and van der Laan (2011) or Theorem 1 in van der Laan (2010), the efficient influence function for (H.20) is

$$D(W, \tilde{A}, L, C^Y, Y) = D_0(W) + D_1(W, \tilde{A}, L) + D_2(W, \tilde{A}, L, C^Y, Y), \quad (\text{H.21})$$

where

$$D_0(W) = E(Y | W, \tilde{A} = 1, C^Y = 1) - E(Y | \tilde{A} = 1), \quad (\text{H.22})$$

$$D_1(W, \tilde{A}, L) = \frac{\mathbf{1}(\tilde{A} = 1)\{E(Y | L, W, \tilde{A} = 1, C^Y = 1) - E(Y | W, \tilde{A} = 1, C^Y = 1)\}}{P(\tilde{A} = 1 | W)}, \quad (\text{H.23})$$

and

$$D_2(W, \tilde{A}, L, C^Y, Y) = \frac{\mathbf{1}(C^Y = 1)\mathbf{1}(\tilde{A} = 1)\{Y - E(Y | L, W, \tilde{A} = 1, C^Y = 1)\}}{P(C^Y = 1 | L, W, \tilde{A} = 1)P(\tilde{A} = 1 | W)}. \quad (\text{H.24})$$

By randomization and independent censoring assumptions, (H.22) simplifies to

$$D_0(W) = E(Y | W, A = a) - E(Y | A = a); \quad (\text{H.25})$$

equation (H.23) simplifies to

$$D_1(W, \tilde{A}, L) = \frac{\mathbf{1}(\tilde{A} = 1)\{E(Y | L, W, A = a) - E(Y | W, A = a)\}}{P(\tilde{A} = 1)}; \quad (\text{H.26})$$

equation (H.24) simplifies to

$$D_2(W, \tilde{A}, L, C^Y, Y) = \frac{\mathbf{1}(C^Y = 1)\mathbf{1}(\tilde{A} = 1)\{Y - E(Y | L, W, A = a)\}}{P(C^Y = 1 | C^L = 1)P(\tilde{A} = 1)}. \quad (\text{H.27})$$

The following lemma states that D_0 , D_1 , and D_2 are pairwise uncorrelated.

Lemma H.2. We have

$$\text{Cov}\{D_0(W), D_1(W, \tilde{A}, L)\} = 0, \quad (\text{H.28})$$

$$\text{Cov}\{D_0(W), D_2(W, \tilde{A}, L, C^Y, Y)\} = 0, \quad (\text{H.29})$$

$$\text{Cov}\{D_1(W, \tilde{A}, L), D_2(W, \tilde{A}, L, C^Y, Y)\} = 0. \quad (\text{H.30})$$

Lemma H.2 implies

$$\text{Var}\{D(W, \tilde{A}, L, C^Y, Y)\} = \text{Var}\{D_0(W)\} + \text{Var}\{D_1(W, \tilde{A}, L)\} + \text{Var}\{D_2(W, \tilde{A}, L, C^Y, Y)\}. \quad (\text{H.31})$$

By (H.25) we have

$$\text{Var}\{D_0(W)\} = \text{Var}\{E_a(Y | W)\} = \text{Var}_a\{E_a(Y | W)\}, \quad (\text{H.32})$$

where the last equality follows from randomization assumption. By (H.26) we have $E\{D_1(W, \tilde{A}, L)\} = 0$, so it follows from randomization and independent censoring that

$$\begin{aligned} \text{Var}\{D_1(W, \tilde{A}, L)\} &= E\left[\frac{\mathbb{1}(\tilde{A} = 1)^2 \{E(Y | L, W, A = a) - E(Y | W, A = a)\}^2}{P(\tilde{A} = 1)^2}\right] \\ &= \frac{E\{\mathbb{1}(\tilde{A} = 1)\}}{P(\tilde{A} = 1)^2} E[\{E(Y | L, W, A = a) - E(Y | W, A = a)\}^2 | A = a]. \end{aligned} \quad (\text{H.33})$$

By independent censoring we have $P(\tilde{A} = 1) = p_a p_l$. It then follows from (H.33) and randomization assumption that

$$\text{Var}\{D_1(W, \tilde{A}, L)\} = \frac{1}{p_a p_l} \text{Var}_a\{E_a(Y | L, W) - E_a(Y | W)\}. \quad (\text{H.34})$$

Similarly, (H.27) together with randomization and monotone censoring imply

$$\text{Var}\{D_2(W, \tilde{A}, L, C^Y, Y)\} = \frac{1}{p_a p_y} \text{Var}_a\{Y - E_a(Y | L, W)\}. \quad (\text{H.35})$$

Because the semiparametric lower bound on the asymptotic variance for an estimand equals the variance of the efficient influence function, by (H.31), (H.32), (H.34), and (H.35) we proved (F.1). □

H.4 Proof of Theorem F.2

Proof. The unadjusted estimator $\hat{\tau}$ for $E(Y | A = a)$ is defined as

$$\hat{\tau} = \frac{\sum_{i=1}^n Y_i \mathbf{1}(A_i = a, C_i^Y = 1)}{\sum_{i=1}^n \mathbf{1}(A_i = a, C_i^Y = 1)}. \quad (\text{H.36})$$

Under Assumptions 1 and 2, $\hat{\tau}$ is unbiased:

$$E(\hat{\tau}) = E \left[E \left\{ \frac{\sum_{i=1}^n Y_i \mathbf{1}(A_i = a, C_i^Y = 1)}{\sum_{i=1}^n \mathbf{1}(A_i = a, C_i^Y = 1)} \middle| A_1, \dots, A_n, C_{Y_1}, \dots, C_{Y_n} \right\} \right] = E(Y | A = a). \quad (\text{H.37})$$

In the following we calculate the asymptotic variance of $\hat{\tau}$.

$$\begin{aligned} \sqrt{n}\{\hat{\tau} - E(Y | A = a)\} &= \frac{\frac{1}{\sqrt{n}} \sum_{i=1}^n Y_i \mathbf{1}(A_i = a, C_i^Y = 1)}{\frac{1}{n} \sum_{i=1}^n \mathbf{1}(A_i = a, C_i^Y = 1)} - \sqrt{n}E(Y | A = a) \\ &= \frac{\frac{1}{\sqrt{n}} \sum_{i=1}^n \{Y_i - E(Y | A = a)\} \mathbf{1}(A_i = a, C_i^Y = 1)}{\frac{1}{n} \sum_{i=1}^n \mathbf{1}(A_i = a, C_i^Y = 1)}. \end{aligned} \quad (\text{H.38})$$

By Weak Law of Large Numbers and the independent censoring assumption,

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1}(A_i = a, C_i^Y = 1) \xrightarrow{P} p_a p_y, \quad (\text{H.39})$$

where \xrightarrow{P} denotes convergence in probability. By Central Limit Theorem,

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n \{Y_i - E(Y | A = a)\} \mathbf{1}(A_i = a, C_i^Y = 1) \xrightarrow{d} N(0, \sigma^2), \quad (\text{H.40})$$

where by randomization and independent censoring we have

$$\begin{aligned} \sigma^2 &= \text{Var}[\{Y - E(Y | A = a)\} \mathbf{1}(A = a, C^Y = 1)] = E[\{Y - E(Y | A = a)\}^2 \mathbf{1}(A = a, C^Y = 1)^2] \\ &= p_a p_y \text{Var}(Y | A = a). \end{aligned} \quad (\text{H.41})$$

Combining (H.38)-(H.41), it follows from Slutsky's theorem that

$$\sqrt{n}\{\hat{\tau} - E(Y | A = a)\} \xrightarrow{d} N(0, (p_a p_y)^{-1} \text{Var}(Y | A = a)).$$

So the asymptotic variance of $\hat{\tau}$ is $(p_a p_y)^{-1} \text{Var}(Y | A = a)$, which by randomization yields (recall that by definition $\text{Var}_a(Y) = \text{Var}(Y | A = a)$)

$$\text{AVar}(\text{unadjusted}) = \frac{1}{p_a p_y} \text{Var}_a(Y). \quad (\text{H.42})$$

Equations (H.42) and (H.8) imply

$$\text{AVar}(\text{unadjusted}) = \frac{1}{p_a p_y} \text{Var}_a(Y) (R_{W;a}^2 + R_{L|W;a}^2 + R_{r;a}^2). \quad (\text{H.43})$$

On the other hand, Theorem F.1 and the definition of R_W^2 and R_r^2 imply

$$\text{AVar}(\text{efficient}) = \text{Var}_a(Y) (R_{W;a}^2 + \frac{1}{p_a p_l} R_{L|W;a}^2 + \frac{1}{p_a p_y} R_{r;a}^2). \quad (\text{H.44})$$

Equations (H.43), (H.44), and (H.8) yield (F.2). The proof is thus finished. \square

H.5 Generalization and proof of Lemma 1

We provide proof for a generalization of Lemma 1, which allows for constant randomization probability other than 1/2.

Lemma H.3. Suppose Assumptions 1, 2, and 3 hold. Define $P(A = a) = p_a$ for $a \in \{0, 1\}$. The lower bound on the asymptotic variance of all RAL estimators of $E(Y | A = 1) - E(Y | A = 0)$ is

$$\begin{aligned} & \text{Var}\{E_1(Y | W) - E_0(Y | W)\} + \sum_{a \in \{0, 1\}} \frac{1}{p_a p_l} \text{Var}_a\{E_a(Y | L, W) - E_a(Y | W)\} \\ & + \sum_{a \in \{0, 1\}} \frac{1}{p_a p_y} \text{Var}_a\{Y - E_a(Y | L, W)\}. \end{aligned} \quad (\text{H.45})$$

Proof. For notation simplicity, denote by $Q(\cdot)$ the conditional expectation $E(Y | \cdot)$. Using the derivation in (H.19)-(H.27) twice for $A = 1$ and $A = 0$, we get the efficient influence function $D(W, A, C^L, L, C^Y, Y)$ for $E(Y | A = 1) - E(Y | A = 0)$:

$$D(W, A, C^L, L, C^Y, Y) = D_0(W) + D_1(W, A, C^L, L) + D_2(W, A, C^L, L, C^Y, Y), \quad (\text{H.46})$$

where

$$D_0(W) = \{Q(W, A = 1) - Q(A = 1)\} - \{Q(W, A = 0) - Q(A = 0)\}, \quad (\text{H.47})$$

$$\begin{aligned} D_1(W, A, C^L, L) &= \frac{AC^L}{p_1 p_l} \{Q(W, L, A = 1) - Q(W, A = 1)\} \\ &\quad - \frac{(1-A)C^L}{p_0 p_l} \{Q(W, L, A = 0) - Q(W, A = 0)\}, \end{aligned} \quad (\text{H.48})$$

and

$$D_2(W, A, C^L, L, C^Y, Y) = \frac{AC^Y}{p_1 p_y} \{Y - Q(W, L, A = 1)\} - \frac{(1 - A)C^Y}{p_0 p_y} \{Y - Q(W, L, A = 0)\}. \quad (\text{H.49})$$

The following lemma states that D_0 , D_1 , and D_2 are pairwise uncorrelated.

Lemma H.4. We have

$$\text{Cov}\{D_0(W), D_1(W, A, C^L, L)\} = 0, \quad (\text{H.50})$$

$$\text{Cov}\{D_0(W), D_2(W, A, C^L, L, C^Y, Y)\} = 0, \quad (\text{H.51})$$

$$\text{Cov}\{D_1(W, A, C^L, L), D_2(W, A, C^L, L, C^Y, Y)\} = 0. \quad (\text{H.52})$$

Lemma H.4 implies

$$\begin{aligned} \text{Var}\{D(W, A, C^L, L, C^Y, Y)\} &= \text{Var}\{D_0(W)\} + \text{Var}\{D_1(W, A, C^L, L)\} \\ &\quad + \text{Var}\{D_2(W, A, C^L, L, C^Y, Y)\}. \end{aligned} \quad (\text{H.53})$$

By H.47 we have

$$\text{Var}\{D_0(W)\} = \text{Var}\{Q(W, A = 1) - Q(W, A = 0)\}. \quad (\text{H.54})$$

By (H.48) we have

$$\begin{aligned} \text{Var}\{D_1(W, A, C^L, L)\} &= E\left[\frac{AC^L}{p_1^2 p_l^2} \{Q(W, L, A = 1) - Q(W, A = 1)\}^2\right] \\ &\quad + E\left[\frac{(1 - A)C^L}{p_0^2 p_l^2} \{Q(W, L, A = 0) - Q(W, A = 0)\}^2\right] \end{aligned} \quad (\text{H.55})$$

$$= \sum_{a \in \{0,1\}} \frac{1}{p_a p_l} E\left[\{Q(W, L, A = a) - Q(W, A = a)\}^2 \mid A = a\right] \quad (\text{H.56})$$

$$= \sum_{a \in \{0,1\}} \frac{1}{p_a p_l} \text{Var}\{Q(W, L, A = a) - Q(W, A = a) \mid A = a\}. \quad (\text{H.57})$$

The step from (H.55) to (H.56) utilizes the independent censoring and randomization assumptions. Similarly, (H.49) together with randomization and independent censoring imply

$$\text{Var}\{D_2(W, A, C^L, L, C^Y, Y)\} = \sum_{a \in \{0,1\}} \frac{1}{p_a p_l} \text{Var}\{Y - Q(W, L, A = a) \mid A = a\}. \quad (\text{H.58})$$

Because the semiparametric lower bound on the asymptotic variance for an estimand equals the variance of the efficient influence function, by (H.53), (H.54), (H.57), and (H.58) we proved Lemma 1. □

H.6 Proof of Theorem 2

Proof. The unadjusted estimator for the average treatment effect is

$$\hat{\tau} = \frac{\sum_{i=1}^n Y_i \mathbf{1}(A_i = 1, C_i^Y = 1)}{\sum_{i=1}^n \mathbf{1}(A_i = 1, C_i^Y = 1)} - \frac{\sum_{i=1}^n Y_i \mathbf{1}(A_i = 0, C_i^Y = 1)}{\sum_{i=1}^n \mathbf{1}(A_i = 0, C_i^Y = 1)}.$$

Similar to the derivation from (H.36) to (H.42), when estimating the average treatment effect with $P(A = 1) = P(A = 0) = 1/2$, we have

$$\text{AVar}(\text{unadjusted}) = \frac{2}{p_y} \sum_{a \in \{0,1\}} \text{Var}_a(Y). \quad (\text{H.59})$$

The result in Theorem 2 then follows immediately from Lemma 1. □

H.7 Proof of Corollary 1

Proof. (i) Because

$$\text{Var}\{E(Y | W)\} = \text{Var}\{0.5E(Y | W, A = 1) + 0.5E(Y | W, A = 0)\},$$

$\text{Var}\{E(Y | W)\} = 0$ implies

$$\text{Var}\{E(Y | W, A = 1)\} + \text{Var}\{E(Y | W, A = 0)\} = -2\text{Cov}\{E(Y | W, A = 1), E(Y | W, A = 0)\}. \quad (\text{H.60})$$

By (H.60) and the definition of γ and R_W^2 , we have $\gamma = 2R_W^2$. This combined with (6) with $R_{L|W}^2 = 0$ imply (E.1).

(ii) By the definition of γ and $\text{Var}\{E(Y | W, A = 1) - E(Y | W, A = 0)\} = 0$, (E.2) follows immediately from (6) with $R_{L|W}^2 = 0$. □

H.8 Proof of statements in Section 3.4

For ease of reading we restate the claims to be proven in Section 3.4: We compare the ARE between two cases: $R_W^2 = q > 0, R_{L|W}^2 = 0$ (only baseline variable prognostic) and $R_W^2 = 0, R_{L|W}^2 = q > 0$ (only short-term outcome prognostic). Regardless of the value of $q > 0$, the ARE in the former case is larger or equal to that in the latter case. Equality occurs if and only if $p_l = 1$ and W is uncorrelated with Y (marginally) in the former case. The latter condition is equivalent to the treatment effect heterogeneity being the maximum possible $\gamma = 2R_W^2 = 2q$.

Proof. We first show that $\gamma \leq 2R_W^2$ with equality holds only when $E_1(Y | W) = -E_0(Y | W)$ almost surely. By Cauchy-Schwarz inequality we have

$$\begin{aligned} & \text{Var}\{E_1(Y | W) - E_0(Y | W)\} \\ &= \text{Var}\{E_1(Y | W)\} + \text{Var}\{E_0(Y | W)\} - 2\text{Cov}\{E_1(Y | W), E_0(Y | W)\} \\ &\leq \text{Var}\{E_1(Y | W)\} + \text{Var}\{E_0(Y | W)\} + 2[\text{Var}\{E_1(Y | W)\}\text{Var}\{E_0(Y | W)\}]^{1/2} \\ &\leq 2\text{Var}\{E_1(Y | W)\} + 2\text{Var}\{E_0(Y | W)\}, \end{aligned}$$

where both inequalities becomes equality if and only if $E_1(Y | W) = -E_0(Y | W)$ almost surely. This proves the claim.

Theorem 2 states that the ARE equals (where we indicate dependence on the arguments by explicitly writing them out)

$$\text{ARE}(R_W^2, R_{L|W}^2, \gamma, p_y, p_l) = \frac{1}{1 + (p_y/2)\gamma - R_W^2 - (1 - p_y/p_l)R_{L|W}^2}.$$

When $R_W^2 = q > 0, R_{L|W}^2 = 0$,

$$\text{ARE}(q, 0, \gamma, p_y, p_l) = \frac{1}{1 + (p_y/2)\gamma - q} \geq \frac{1}{1 - (1 - p_y)q},$$

with equality holds if and only if $\gamma = 2q$. When $R_W^2 = 0, R_{L|W}^2 = q > 0$ (which implies $\gamma = 0$ because $\gamma \leq 2R_W^2$),

$$\text{ARE}(0, R_{L|W}^2, 0, p_y, p_l) = \frac{1}{1 - (1 - p_y/p_l)q}.$$

Because $p_l \leq 1$, $\text{ARE}(q, 0, \gamma, p_y, p_l) \geq \text{ARE}(0, R_{L|W}^2, 0, p_y, p_l)$ with equality holds if and only if $p_l = 1$ and $\gamma = 2q$, i.e., $E_1(Y | W) = -E_0(Y | W)$ almost surely.

Lastly, we show that $E_1(Y | W) = -E_0(Y | W)$ almost surely implies that W is uncorrelated with Y (marginally):

$$\begin{aligned} \text{Cov}(Y, W) &= E\{Y(W - E(W))\} = E\{E(Y | W)(W - E(W))\} \\ &= E\{E(Y | W)(W - E(W))\} = E[\{0.5E_1(Y | W) + 0.5E_0(Y | W)\}(W - E(W))] = 0. \end{aligned}$$

This completes the proof. □

I Proof of Auxiliary Lemmas

I.1 Additional Supporting Lemmas

Lemma I.1. Consider three random variables X , Y , and Z . Denote by $\sigma(Z)$ the σ -field generated by Z . If $X \in \sigma(Z)$, then

$$E\{YE(Y | X)\} = E\{E(Y | Z)E(Y | X)\}. \quad (\text{I.1})$$

Proof. By the law of iterated expectation, we have

$$E\{YE(Y | X)\} = E[E\{YE(Y | X) | Z\}]. \quad (\text{I.2})$$

Because $X \in \sigma(Z)$, we have $E(Y | X) \in \sigma(X) \subset \sigma(Z)$. This implies

$$E\{YE(Y | X) | Z\} = E(Y | X)E(Y | Z). \quad (\text{I.3})$$

Equations (I.2) and (I.3) imply (I.1). This completes the proof. □

Lemma I.2. Consider three random variables W , L , and Y . For any measurable functions $f(W)$ and $g(W, L)$, we have

$$E[f(W)\{E(Y | W, L) - E(Y | W)\}] = 0, \quad (\text{I.4})$$

$$E[g(W, L)\{Y - E(Y | W, L)\}] = 0. \quad (\text{I.5})$$

Proof. By the law of iterated expectation, we have

$$\begin{aligned} E\{g(W)E(Y | W, L)\} &= E[E\{g(W)E(Y | W, L) | W\}] \\ &= E[g(W)E\{E(Y | W, L) | W\}] \\ &= E\{g(W)E(Y | W)\}, \end{aligned}$$

which proves (I.4).

Similarly, we have

$$E\{g(W, L)Y\} = E[E\{g(W, L)Y | W, L\}] = E\{g(W, L)E(Y | W, L)\},$$

which proves (I.5). □

I.2 Proof of Lemma H.2

Proof. By (H.25) and (H.26) we have

$$\text{Cov}(D_0, D_1) \propto E[\mathbf{1}(\tilde{A} = 1)\{E_a(Y | W) - E_a(Y)\}\{E_a(Y | L, W) - E_a(Y | W)\}]. \quad (\text{I.6})$$

Combining (H.12), (I.6), and Assumptions 1 and 2, we derive (H.28).

By (H.25) and (H.27) we have

$$\text{Cov}(D_0, D_2) \propto E[\mathbf{1}(\tilde{A} = 1, C^Y = 1)\{E_a(Y | W) - E_a(Y)\}\{Y - E_a(Y | L, W)\}]. \quad (\text{I.7})$$

Combining (H.11), (I.7), and Assumptions 1 and 2, we derive (H.29).

By (H.26) and (H.27) we have

$$\text{Cov}(D_1, D_2) \propto E[\mathbf{1}(\tilde{A} = 1, C^Y = 1)\{E_a(Y | L, W) - E_a(Y | W)\}\{Y - E_a(Y | L, W)\}]. \quad (\text{I.8})$$

Combining (H.10), (I.8), and Assumptions 1 and 2, we derive (H.30).

This completes the proof. □

I.3 Proof of Lemma H.4

For notation simplicity, we use $E_1(\cdot)$ and $E_0(\cdot)$ to denote $E(\cdot | A=1)$ and $E(\cdot | A=0)$, respectively. By (H.47) and (H.48) we have

$$\begin{aligned}
\text{Cov}(D_0, D_1) &\propto E[A\{Q(W, A=1) - Q(W, A=0)\}\{Q(W, L, A=1) - Q(W, A=1)\}] \\
&\quad - E[(1-A)\{Q(W, A=1) - Q(W, A=0)\}\{Q(W, L, A=0) - Q(W, A=0)\}] \\
&= \sum_{a \in \{0,1\}} E_a[\{E_1(Y | W) - E_0(Y | W)\}\{E_a(Y | W, L) - E_a(Y | W)\}] \times P(A=a).
\end{aligned} \tag{I.9}$$

Both terms in (I.9) equals 0 by (I.4) in Lemma I.2 with $f(W) = E_1(Y | W) - E_0(Y | W)$. This yields (H.50).

By (H.47) and (H.49) we have

$$\begin{aligned}
\text{Cov}(D_0, D_2) &\propto E[A\{Q(W, A=1) - Q(W, A=0)\}\{Y - Q(W, L, A=1)\}] \\
&\quad - E[(1-A)\{Q(W, A=1) - Q(W, A=0)\}\{Y - Q(W, L, A=0)\}] \\
&= \sum_{a \in \{0,1\}} E_a[\{E_1(Y | W) - E_0(Y | W)\}\{Y - E_a(Y | W, L)\}] \times P(A=a).
\end{aligned} \tag{I.10}$$

Both terms in (I.10) equals 0 by (I.5) in Lemma I.2 with $g(W, L) = E_1(Y | W) - E_0(Y | W)$. This yields (H.51).

By (H.48) and (H.49) we have

$$\begin{aligned}
\text{Cov}(D_1, D_2) &\propto E[A\{Q(W, L, A=1) - Q(W, A=1)\}\{Y - Q(W, L, A=1)\}] \\
&\quad - E[(1-A)\{Q(W, L, A=0) - Q(W, A=0)\}\{Y - Q(W, L, A=0)\}] \\
&= \sum_{a \in \{0,1\}} E_a[\{E_a(Y | W, L) - E_a(Y | W)\}\{Y - E_a(Y | W, L)\}] \times P(A=a).
\end{aligned} \tag{I.11}$$

Both terms in (I.11) equals 0 by (I.5) in Lemma I.2 with $g(W, L) = E_a(Y | W, L) - E_a(Y | W)$. This yields (H.52).

This completes the proof.