



MIT Open Access Articles

Regression-assisted inference for the average treatment effect in paired experiments

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation	Fogarty, Colin B. "Regression-Assisted Inference for the Average Treatment Effect in Paired Experiments." <i>Biometrika</i> 105, 4 (June 2018): 994–1000 © 2018 Biometrika Trust
As Published	http://dx.doi.org/10.1093/BIOMET/ASY034
Publisher	Oxford University Press (OUP)
Version	Author's final manuscript
Citable link	http://hdl.handle.net/1721.1/120748
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

Regression assisted inference for the average treatment effect in paired experiments

Colin B. Fogarty *

Abstract

In paired randomized experiments individuals in a given matched pair may differ on prognostically important covariates despite the best efforts of practitioners. We examine the use of regression adjustment as a way to correct for persistent covariate imbalances after randomization, and present two regression assisted estimators for the sample average treatment effect in paired experiments. Using the potential outcomes framework, we prove that these estimators are consistent for the sample average treatment effect under mild regularity conditions even if the regression model is improperly specified. Further, we describe how asymptotically conservative confidence intervals can be constructed. We demonstrate that the variances of the regression assisted estimators are at least as small as that of the standard difference-in-means estimator asymptotically. Through a simulation study, we illustrate the appropriateness of the proposed methods in small and moderate samples. The analysis does not require a superpopulation model, a constant treatment effect, or the truth of the regression model, and hence provides a mode of inference for the sample average treatment effect with the potential to yield improvements in the power of the resulting analysis over the classical analysis without imposing potentially unrealistic assumptions.

1 Introduction

The use of paired experiments has historically been limited to cases in which pairs are formed using a small number of binary or categorical covariates. In practice, there are often continuous covariates which are also believed to be predictive for the potential outcomes under treatment and control. Greevy et al. (2004) introduced a form of multivariate matching before randomization which assigns a group of $2n$ experimental units to n pairs to minimize the within-pairs covariate distance, hence improving covariate balance for many variables at the same time. Once the pairs are established, exactly one subject in each pair is randomly assigned to the treatment. Fisherian inference for no treatment effect and Neymanian inference for the average treatment effect can then proceed with respect to the randomization distribution generated by this paired design; see, for example, Rosenbaum (2002) and Imai (2008).

While mitigable through the method of Greevy et al. (2004), persistent within-pair discrepancies on the basis of continuous covariates of interest are unavoidable. These discrepancies may,

*Operations Research and Statistics Group, MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge MA 02142 (e-mail: cfogarty@mit.edu)

in turn, yield chance imbalances on the basis of these covariates in any given randomization. A common strategy for accounting for remaining imbalances is covariance adjustment. Freedman (2008) and Lin (2013) investigate the impact of regression adjustment on inference for the sample average treatment effect in completely randomized experiments. Neither correctness of the fitted regression model nor a superpopulation model (that units are drawn independently and identically distributed from a larger population) are assumed in their analysis. Instead, inference proceeds using the physical act of randomization as its sole justification, and asymptotic calculations consider sequences of experiments of increasing size without specifying how the experimental units were sampled. In this context, Lin (2013) discusses how regression adjustment can yield a consistent estimator of the sample average treatment effect whose asymptotic variance is no larger than the standard difference-in-means estimator regardless of the truth of the underlying model. The performed inference is then agnostic to the truth of the underlying model and instead leverages the adjustment as an algorithmic fit in an attempt to yield a more efficient estimator. See also Aronow and Middleton (2013) and Bloniarz et al. (2016) for related work with completely randomized experiments.

In the context of paired experiments, Rosenbaum (2002) describes how covariance adjustment can be utilized to yield exact tests under the assumption of a constant treatment effect. Imbens and Rubin (2015, §10) discuss how under a superpopulation model, linear regression can yield a consistent estimator for the average treatment effect in the superpopulation. In what follows, we show that covariance adjustment can also be leveraged for inference on the sample average treatment effect in a paired experiment in a manner that is agnostic to the truth of the fitted model.

2 A paired randomized experiment with covariates

The i th of n independent pairs contains one individual assigned the treatment, denoted as $Z_{ij} = 1$, and one who receives the control, $Z_{ij'} = 0$, such that $Z_{i1} + Z_{i2} = 1$. Individual j in matched set i has a P -dimensional vector of measured covariates $x_{ij} = (x_{ij1}, \dots, x_{ijP})$. Each individual has a potential outcome under treatment, r_{Tij} , and under control, r_{Cij} , $i = 1, \dots, n; j = 1, 2$. The pair of potential outcomes (r_{Tij}, r_{Cij}) , and with it the treatment effect $\tau_{ij} = r_{Tij} - r_{Cij}$, is not jointly observable for any individual. Instead, we observe the response $R_{ij} = r_{Tij}Z_{ij} + r_{Cij}(1 - Z_{ij})$ for each individual (Neyman, 1923; Rubin, 1974). Quantities dependent on the assignment vector such as $Z = (Z_{11}, Z_{12}, \dots, Z_{n2})$ and $R = (R_{11}, R_{12}, \dots, R_{n2})$ are random, whereas $\mathcal{F} = \{(r_{Tij}, r_{Cij}, x_{ij}), i = 1, \dots, n, j = 1, 2\}$ contains fixed quantities. In a paired experiment, $\text{pr}(Z = z \mid \mathcal{F}) = \text{pr}(Z = z) = 2^{-n}$, and $\text{pr}(Z_{ij} = 1 \mid \mathcal{F}) = \text{pr}(Z_{ij} = 1) = 1/2$ for $i = 1, \dots, n, j = 1, 2$.

Write $d_i = f(x_{i1}) - f(x_{i2})$ for some function $f : \mathbb{R}^P \rightarrow \mathbb{R}^{K_D}$ with K_D fixed as the difference in transformed covariates between unit 1 and unit 2 in matched pair i . Similarly write $m_i = \{g(x_{i1}) + g(x_{i2})\}/2 - n^{-1} \sum_{i'=1}^n \{g(x_{i'1}) + g(x_{i'2})\}/2$ for some function $g : \mathbb{R}^P \rightarrow \mathbb{R}^{K_M}$, as the average of the transformed covariates in matched pair i relative to the mean across matched pairs. For example, if x is scalar then choosing $f(x) = (x, x^2)$, $g(x) = \exp(x)$ yields $K_D = 2$, $K_M = 1$. Guidance for the choice of the functions $f(\cdot)$ and $g(\cdot)$, which may differ, is given in §4. Write D for the $n \times K_D$ matrix whose i th row equals d_i^T , and write M for the $n \times K_M$ matrix whose i th row equals m_i^T . In what follows we require that $K_D + K_M + 1 < n$. Let $H_D = D(D^T D)^{-1} D$

be the hat matrix for D , i.e. the orthogonal projection of \mathbb{R}^K onto the column space of D , and let H_M be that hat matrix for M . Define $V_i = 2Z_{i1} - 1$ such that $E(V_i | \mathcal{F}) = 0$, and let V be the $n \times n$ matrix with V_i on the i th diagonal and zeroes on the off-diagonal. The matrix consisting of the treated-minus-control differences of the f -transformed covariates can be written as VD , with hat matrix $VD(D^T VVD)^{-1}DV = VH_DV$ since $V_i^2 = 1$. Finally, write A for the $n \times (K_D + K_M)$ matrix whose i th row is $a_i^T = (V_i d_{i1}, \dots, V_i d_{iK_D}, m_{i1}, \dots, m_{iK_M})^T$, and write H_A for the corresponding hat matrix.

3 The sample average treatment effect

Let $\ell_{ij} = (r_{Tij} + r_{Cij})/2$ be the average, or level, of the potential outcomes for individual j in pair i , and let $\Delta_i = (\tau_{i1} + \tau_{i2})/2$ be the average of the treatment effects in pair i . The observed treated-minus-control difference in responses in pair i is $Y_i = \Delta_i + V_i(\ell_{i1} - \ell_{i2})$. Write $\ell_j = (\ell_{1j}, \dots, \ell_{nj})$ for $j = 1, 2$, $\Delta = (\Delta_1, \dots, \Delta_n)$, and $Y = (Y_1, \dots, Y_n)$. The sample average treatment effect in a paired experiment with n pairs is

$$\bar{\tau}_n = \frac{1}{2n} \sum_{i=1}^n (\tau_{i1} + \tau_{i2}) = \frac{1}{n} \sum_{i=1}^n \Delta_i.$$

Henceforth we will write $\bar{\tau} = \bar{\tau}_n$, but the dependence of this and other causal estimands on n should be kept in mind. The classical unbiased estimator for $\bar{\tau}$ in a paired experiment is the observed average of the paired differences $\hat{\tau}_C = \bar{Y} = n^{-1} \sum_{i=1}^n Y_i$. The variance for $\hat{\tau}_C$ across randomizations is $\text{var}(\hat{\tau}_C | \mathcal{F}) = n^{-2} \sum_{i=1}^n (\ell_{i1} - \ell_{i2})^2$, which is unknown in practice because it depends on the missing potential outcomes. Imai (2008) shows that the classical estimator of the variance of the difference in means in a paired study, $S_C^2 = \{n(n-1)\}^{-1} \sum_{i=1}^n (Y_i - \hat{\tau}_C)^2$, is always an upper bound on $\text{var}(\hat{\tau}_C | \mathcal{F})$ in expectation. Furthermore, under mild regularity conditions $nS_C^2 - \text{var}(n^{1/2}\hat{\tau}_C | \mathcal{F})$ converges to a nonnegative value in probability, and asymptotically conservative inference can be achieved by employing a normal approximation with S_C^2 in place of $\text{var}(\hat{\tau}_C | \mathcal{F})$.

4 Regression assisted estimators

We consider the intercept coefficients from a regression of Y on VD , $\hat{\tau}_{R1}$, and of Y on VD and M , $\hat{\tau}_{R2}$, as estimators for $\bar{\tau}$. These can be expressed as

$$\hat{\tau}_{R1} = \frac{e^T(I - VH_DV)Y}{e^T(I - VH_DV)e}; \quad \hat{\tau}_{R2} = \frac{e^T(I - H_A)Y}{e^T(I - H_A)e},$$

where I is the identity matrix, and e is the vector containing n ones (i.e., the intercept column). As discussed in Imbens and Rubin (2015, §10), a regression on VD encodes a belief that the difference between f -transformed covariates in matched pair i , d_i , may be predictive of the difference in the levels of the potential outcomes $\ell_{i1} - \ell_{i2}$. A regression including both VD and M encodes a belief that the difference in f -transformed covariates in matched pair i is predictive of the difference in the levels, and that the relative level of the g -transformed covariates m_i is

predictive of the average treatment effect in pair i , Δ_i . As such, the function $f(\cdot)$ giving rise to d_i , should, ideally, be chosen to best reflect the relationship between the transformed differences $d_i = f(x_{i1}) - f(x_{i2})$ and the difference in level of the potential outcomes $\ell_{i1} - \ell_{i2}$. Similarly, the function $g(\cdot)$ should be chosen to best reflect the relationship between the pairwise average of the transformed covariates, $(g(x_{i1}) + g(x_{i2}))/2$, and the pairwise average of the treatment effects Δ_i .

Under a constant effect model $\hat{\tau}_{R1}$ is an unbiased estimator for $\bar{\tau}$, and inference for τ can be conducted using the techniques described in Rosenbaum (2002). Without the assumption of additivity, neither $\hat{\tau}_{R1}$ nor $\hat{\tau}_{R2}$ are unbiased for $\bar{\tau}$. Nonetheless, we now demonstrate that both $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ can be used to facilitate inference on $\bar{\tau}$ without constant effects. Moving forward, we impose two regularity conditions.

Condition 1. (*Bounded Fourth Moments*). There exists a $C < \infty$ such that, for all n , $n^{-1} \sum_{i=1}^n \Delta_i^4 < C$, $n^{-1} \sum_{i=1}^n (\ell_{i1} - \ell_{i2})^4 < C$, $n^{-1} \sum_{i=1}^n d_{ik}^4 < C$, $k = 1, \dots, K_D$, and $n^{-1} \sum_{i=1}^n m_{ik'}^4 < C$, $k' = 1, \dots, K_M$.

Condition 2. (*Cesàro Summability*). $n^{-1} \sum_{i=1}^n \Delta_i$, $n^{-1} \sum_{i=1}^n \Delta_i^2$, and $n^{-1} \sum_{i=1}^n (\ell_{i1} - \ell_{i2})^2$ converge to finite limits as $n \rightarrow \infty$. Further, $n^{-1} \sum_{i=1}^n (\ell_{i1} - \ell_{i2})d_{ik}$ and $n^{-1} \sum_{i=1}^n \Delta_i m_{ik}$ converge to finite limits for $k = 1, \dots, K_D$, $k' = 1, \dots, K_M$ as $n \rightarrow \infty$. Finally, $n^{-1} D^T D$ and $n^{-1} M^T M$ converges to a finite, invertible matrices. Call these limits Σ_D and Σ_M respectively.

Let β_D be the limit of $(D^T D)^{-1} D^T (\ell_1 - \ell_2)$, and let β_M be the limit of $(M^T M)^{-1} M^T \Delta$.

Theorem 1. Define $\hat{\tau}_{R*} = n^{-1} \sum_{i=1}^n (Y_i - V_i d_i^T \beta_D)$, and suppose Conditions 1 and 2 hold. Then, as $n \rightarrow \infty$ with K_D and K_M fixed and conditional upon \mathcal{F} , both $n^{1/2}(\hat{\tau}_{R1} - \hat{\tau}_{R*})$ and $n^{1/2}(\hat{\tau}_{R2} - \hat{\tau}_{R*})$ converge in probability to zero. Furthermore, $n^{1/2}(\hat{\tau}_{R*} - \bar{\tau})$ converges in distribution to a Gaussian random variable with mean zero and variance

$$\sigma_{R*}^2 = \lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n (\ell_{i1} - \ell_{i2})^2 - \beta_D^T \Sigma_D \beta_D.$$

Theorem 1 characterizes several useful properties of the regression assisted estimators $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ under Conditions 1 and 2. First, $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ are consistent estimators of $\bar{\tau}$. Second, $n^{1/2}(\hat{\tau}_{R1} - \bar{\tau})$ and $n^{1/2}(\hat{\tau}_{R2} - \bar{\tau})$ are asymptotically equivalent, with a common Gaussian limiting distribution. Finally, note that the first term in σ_{R*}^2 is precisely the asymptotic variance of $n^{1/2} \hat{\tau}_C$ given \mathcal{F} . Hence, the asymptotic variances of the regression assisted estimators are no larger than that of the classical difference-in-means estimator, as $\beta_D^T \Sigma_D \beta_D \geq 0$ by positive semi-definiteness of Σ_D . Better choices of $f(\cdot)$ impact the magnitude of $\beta_D^T \Sigma_D \beta_D$ and hence the degree of variance reduction.

Perhaps surprisingly, $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ have the same asymptotic variance. To see why, the average value of the covariates in a given matched pair does not vary across randomizations; however, the difference between the covariates for the treated and control individuals does, taking values d_i and $-d_i$ with equal probability. If the slope coefficients on the level of covariates m_i have limits as $n \rightarrow \infty$ (as guaranteed by Conditions 1 and 2), then asymptotically the contribution of m_i to the prediction of Y_i also does not vary across randomizations, and hence does not contribute to the

variance of the estimator. The choice of the function $g(\cdot)$ thus plays no role in improving efficiency. In light of this, is there motivation for including the level of the covariates when estimating $\bar{\tau}$? As we now demonstrate, inclusion of M in the regression allows for the construction of variance estimators that are less conservative than those derived from a regression excluding M .

5 Enabling Neyman-style inference

As described in §3 inference using the classical difference-in-means estimator, $\hat{\tau}_C$, typically proceeds using an upper bound on the variance. If the estimator is a consistent upper bound, inference performed is then asymptotically conservative if a Gaussian reference distribution is asymptotically justified. As will be demonstrated, the nominal variance estimators for the intercept coefficients derived from linear model theory with fixed design and homoskedastic errors can be employed towards this end. Under homoskedasticity the classical variance estimators for the intercept coefficients $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ take the form

$$S_{R1}^2 = \frac{(Y - \hat{\tau}_{R1}e)^T(I - VH_DV)(Y - \hat{\tau}_{R1}e)}{(n - K_D - 1)e^T(I - VH_DV)e}; \quad S_{R2}^2 = \frac{(Y - \hat{\tau}_{R2}e)^T(I - H_A)(Y - \hat{\tau}_{R2}e)}{(n - K_D - K_M - 1)e^T(I - H_A)e}.$$

See §D of the appendix for a derivation. The developments that follow also apply if these estimators are replaced with heteroskedasticity consistent standard errors (Long and Ervin, 2000).

Theorem 2. *Under Conditions 1 and 2 and conditional upon \mathcal{F} , $nS_{R1}^2 - \text{var}(n^{1/2}\hat{\tau}_{R*} \mid \mathcal{F})$ converges in probability to*

$$\lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n (\Delta_i - \bar{\tau})^2 \geq 0.$$

Under the same conditions, $nS_{R2}^2 - \text{var}(n^{1/2}\hat{\tau}_{R} \mid \mathcal{F})$ converges in probability to*

$$\lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n (\Delta_i - \bar{\tau})^2 - \beta_M^T \Sigma_M \beta_M = \lim_{n \rightarrow \infty} n^{-1} (\Delta - \bar{\tau}e)^T (I - H_M) (\Delta - \bar{\tau}e) \geq 0.$$

Corollary 1. *Under these conditions, $nS_{R1}^2 - nS_{R2}^2$ converges in probability to*

$$\beta_M^T \Sigma_M \beta_M \geq 0.$$

Theorem 2, along with Corollary 1, indicate that both nS_{R1}^2 and nS_{R2}^2 will be consistent upper bounds for $\text{var}(n^{1/2}\hat{\tau}_{R*} \mid \mathcal{F})$, but that nS_{R2}^2 will be asymptotically no larger than nS_{R1}^2 . Asymptotic equality is attained under an additive treatment effect model, where both nS_{R1}^2 and nS_{R2}^2 are consistent for $\text{var}(n^{1/2}\hat{\tau}_{R*} \mid \mathcal{F})$. Unlike nS_{R1}^2 , nS_{R2}^2 can also be consistent for $\text{var}(n^{1/2}\hat{\tau}_{R*} \mid \mathcal{F})$ if the relative level of the covariates in a pair m_i is perfectly predictive of the pairwise average treatment effects Δ_i , hence highlighting the role of the function $g(\cdot)$ used in defining M . In combination with Theorem 1, Theorem 2 indicates that on asymptotic grounds $\hat{\tau}_{R2}$ should be preferred over $\hat{\tau}_{R1}$. Confidence intervals for $\bar{\tau}$ of the form $\hat{\tau}_{R1} \pm \Phi^{-1}(1 - \alpha/2)S_{R1}$ and $\hat{\tau}_{R2} \pm \Phi^{-1}(1 - \alpha/2)S_{R2}$ will be asymptotically conservative as desired. Confidence intervals

for the sample average treatment effect $\bar{\tau}$ constructed using S_{R2} will be no longer than those constructed using S_{R1} , and hypothesis tests for the null $H_0 : \bar{\tau} = \bar{\tau}_0$ through the test statistic $(\hat{\tau}_{R2} - \bar{\tau}_0)/S_{R2}$ and a Gaussian reference distribution will be more powerful than that using $(\hat{\tau}_{R1} - \bar{\tau}_0)/S_{R1}$ while asymptotically maintaining the desired size.

6 A simulation study

6.1 Linear regression under a nonlinear truth

Our study contains two simulation settings inspired by the functions used in Kang and Schafer (2007). The s th of S samples of n pairs is generated by first sampling covariates w_{ijp} , $p = 1, \dots, 4$ that are unknown to the researcher for each of the $2n$ study participants. For each p , we stipulate that w_{i1p} is Gaussian distributed with mean zero and variance 1, w_{i2p} Gaussian distributed with mean w_{i1p} and variance 1/4, and w_{ijp} is independent of $w_{ijp'}$ for $p' \neq p$. Potential outcomes under treatment and control are simulated through $r_{Tij} = \mu_T(w_{ij}) + \varepsilon_{ij}$ and $r_{Cij} = \mu_C(w_{ij}) + \varepsilon_{ij}$, with ε_{ij} standard normal. The sample average treatment effect for sample s is $\bar{\tau}^{(s)} = (2n)^{-1} \sum_{i=1}^n \sum_{j=1}^2 \{\mu_T(w_{ij}) - \mu_C(w_{ij})\}$. The settings vary in the functions $\mu_T(\cdot)$ and $\mu_C(\cdot)$. The two possibilities for $\mu_T(\cdot)$ and $\mu_C(\cdot)$ are (1) parallel response surfaces with $\mu_T(w_{ij}) = \mu_C(w_{ij}) = 27.4w_{ij1} + 13.7(w_{ij2} + w_{ij3} + w_{ij4})$, implying constant treatment effects; and (2) nonparallel response surfaces with $\mu_T(w_{ij}) = 27.4w_{ij1} + 13.7(w_{ij2} + w_{ij3} + w_{ij4})$ and $\mu_C(w_{ij}) = 13.7(w_{ij1} + w_{ij2}) + 3w_{ij3} + 27.4w_{ij4}$, implying heterogeneous treatment effects.

The revealed covariates x_{ij} for each individual are complicated functions of w_{ij} ,

$$\begin{aligned} x_{ij1} &= \exp(w_{ij1}/2); & x_{ij2} &= w_{ij2}/\{1 + \exp(w_{ij1})\} + 10 \\ x_{ij3} &= \{(w_{ij1}w_{ij3})/25 + 0.6\}^3; & x_{ij4} &= (w_{ij2} + w_{ij4} + 20)^2. \end{aligned}$$

For each randomization, the researcher computes $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ from a multiple regression using the observed covariates x_{ij} instead of w_{ij} and setting $f(x_{ij}) = g(x_{ij}) = x_{ij}$, in so doing fitting a linear model using covariates that have a highly nonlinear relationship with the potential outcomes. With the experimental units and observed covariates established, we conduct B randomizations wherein we assign the treatment to exactly one unit in each pair. For each randomization, we calculate $\hat{\tau}_C$, $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ based on the observed responses. We then construct normal-based 95% confidence intervals for the experiment-specific sample average treatment effect $\bar{\tau}^{(s)}$ using S_C^2 , S_{R1}^2 , and S_{R2}^2 .

6.2 Results for two sets of experimental units

Figure 1 illustrates the results of $B = 1000$ randomizations for two samples with $n = 500$ individuals, one from each of the specifications for the response functions (parallel and nonparallel). The histograms show the across-randomization distributions of the three estimators minus the true value of $\bar{\tau}^{(s)}$. As predicted by Theorem 1 all three estimators are well approximated by normal distributions. Further, we see that the true dispersions of $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ are closely aligned, but that both estimators have smaller variance than $\hat{\tau}_{RC}$ despite the fact that the regression model was misspecified. The second plots for each setting compare the length of normal-based 95%

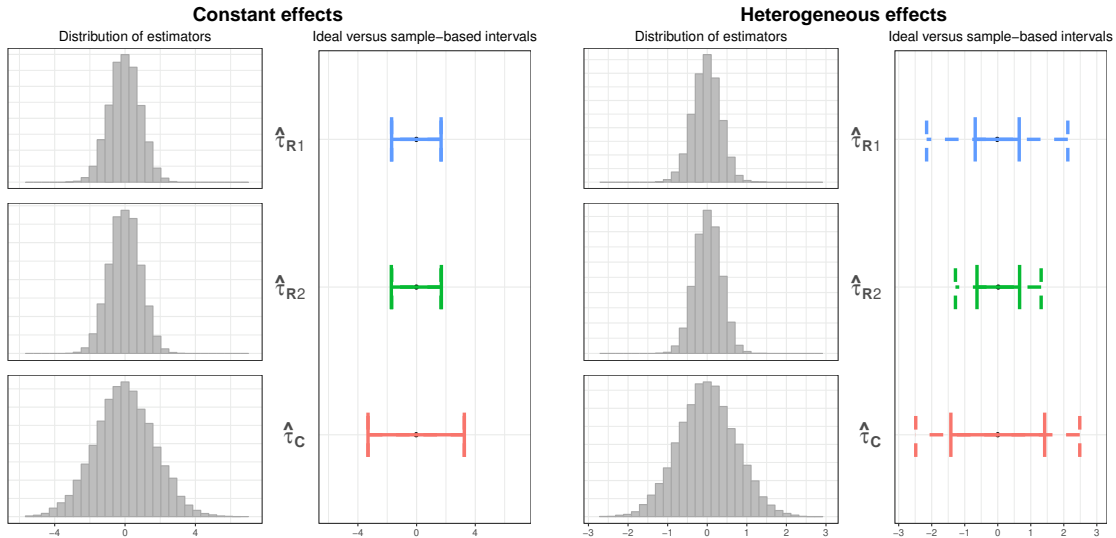


Figure 1: A graph showing, for each estimator, its simulated distribution, normal-based 95% range based on true variances (solid), and the approximation to the 95% range based on estimated variances (dashed) with constant (left) and heterogeneous (right) treatment effects.

confidence intervals constructed using the true variance for the estimators to the typical length of 95% intervals constructed using the sample estimators for those variances, S_{R1}^2 , S_{R2}^2 , and S_C^2 respectively. Here, the conclusions of Theorem 2 come to bear. Under constant effects (left), the expected sample-based interval overlaps with the intervals based on the true variance, as with constant effects the three variance estimators are consistent. Consequently, the corresponding confidence intervals had coverage of roughly 95%. With heterogeneous treatment effects (right) we instead see that typical intervals based on the sample at hand are wider than those based on the true variance, as under effect heterogeneity the estimators are instead consistent upper bounds. This conservativeness led to confidence intervals that, for all estimators, had coverage of 100% in the $B = 1000$ randomizations. Comparing $\hat{\tau}_{R1}$ to $\hat{\tau}_{R2}$, we see that while the ideal confidence intervals based on the true variance are quite similar the corresponding sample-based intervals differ, with those based on $\hat{\tau}_{R2}$ being roughly $2/3$ the length of those based on $\hat{\tau}_{R1}$ while still exceeding their coverage guarantee. Despite the regression model being misspecified, effect modification was leveraged by S_{R2}^2 to yield narrower intervals. In §E.1 of the appendix, we present detailed results for this simulation study with $n = 25, 50, 100$, and 500 . Therein, we see that the predictions of Theorem 1 and 2 provide appropriate guidance even in moderately sized samples.

7 Superpopulation inference after regression adjustment

The improvement in inference by means of $\hat{\tau}_{R2}$ and S_{R2}^2 is not without qualification. Suppose that one instead considered a superpopulation model wherein n pairs of individuals, and hence

their potential outcomes and covariates, were drawn at random from an infinite population and that inference is desired for the expectation of the treatment effect in that superpopulation, call it $\bar{\tau}^{(P)}$. Inference is no longer conducted conditional upon \mathcal{F} as described previously in this article, but instead must account for variation across realizations of \mathcal{F} . In this setting, $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ remain consistent for $\bar{\tau}^{(P)}$, are asymptotically equivalent, and have true variance that is no larger than that of the difference-in-means estimator $\hat{\tau}_C$. The limiting variances $\text{var}(n^{1/2}\hat{\tau}_{R1})$ and $\text{var}(n^{1/2}\hat{\tau}_{R2})$ are increased by $\text{var}(n^{1/2}\bar{\tau})$, as the sample average treatment effect is itself a random variable under this formulation with $\text{var}(\bar{\tau})$ equal to the expectation of $\{n(n-1)\}^{-1} \sum_{i=1}^n (\Delta_i - \bar{\tau})^2$. Arguments parallel to those in §5.3 of a 2017 unpublished paper by the author (arXiv:1706.06469) show that nS_{R1}^2 is a consistent estimator of $\text{var}(n^{1/2}\hat{\tau}_{R1})$, while nS_{R2}^2 (which achieves its lower probability limit by exploiting effect modification) is an underestimate of $\text{var}(n^{1/2}\hat{\tau}_{R2})$, in turn yielding anti-conservative inference. For superpopulation inference using $\hat{\tau}_{R2}$, one should instead employ the corrected variance estimator $S_{R2,P}^2 = S_{R2}^2 + n^{-1}\hat{\beta}_M^T \hat{\Sigma}_M \hat{\beta}_M$, where $\hat{\Sigma}_M = (n-1)^{-1}M^T M$ is the sample covariance matrix for the centered covariates M and $\hat{\beta}_M$ is the vector of regression coefficients corresponding to the columns of M from a regression of Y on $A = (VD, M)$ along with an intercept. Defined in this manner, $nS_{R2,P}^2$ is consistent for $\text{var}(n^{1/2}\hat{\tau}_{R2})$. The importance of this correction is highlighted in a simulation study presented in §E.2 of the appendix. In light of Corollary 1, we see that asymptotically $nS_{R2,P}^2$ is increased by precisely the discrepancy between nS_{R2}^2 and nS_{R1}^2 . In fact, $nS_{R1}^2 - nS_{R2,P}^2$ converges in probability to zero, hence eliminating the inferential advantage held by $\hat{\tau}_{R2}$ over $\hat{\tau}_{R1}$ asymptotically. Whether finite-sample properties of inference based on $\hat{\tau}_{R2}$ versus $\hat{\tau}_{R1}$ would lead one estimator to be preferred in a superpopulation setting remains an area for further research.

8 Discussion

This work has focused on the use of agnostic linear regression (i) to yield an estimator of the sample average treatment effect with improved asymptotic efficiency over the difference-in-means, and (ii) to furnish conservative standard error estimators for regression-adjusted estimators which dominate the conventional standard error for the difference-in-means. A natural and important extension of this work would be to consider other forms of regression adjustment in paired experiments by leveraging semiparametric theory, aligning with the approach taken for completely randomized designs by Zhang et al. (2008) and Tsiatis et al. (2008). The focus on linear regression serves, in part, to provide justification for what remains the most common form of covariance adjustment utilized by practitioners under minimal assumptions about the manner in which the data were generated.

In completely randomized experiments, Aronow et al. (2014) presents sharp variance estimators for the difference-in-means estimator of the sample average treatment effect without covariance adjustment. These results may furnish improvements over the variance estimators considered by Lin (2013) for regression adjustment in completely randomized designs; however, the natural extension to block-randomized designs requires at least two treated and two control individuals in each block, a feature absent in paired experiments. Nonetheless, the variance estimator S_{R2}^2 provides a means of improving the estimator S_{R1}^2 when the target of estimation is the sample average treatment effect.

A Facts and Lemmas

The following fact about the matrix M , true by construction, will be useful moving forward.

Fact 1. $e^T m_{k'} = 0$ and $(I - H_M)e = e$ due to centering of the covariates $m_{k'}, k' = 1, \dots, K_M$.

We now prove three lemmas which will be utilized in the proofs of our main results.

Lemma 1. $n^{-1}A^T A$ converges in probability to a block-diagonal matrix with $\lim_{n \rightarrow \infty} n^{-1}D^T D$ in the upper-left block of size $K_D \times K_D$, $\lim_{n \rightarrow \infty} n^{-1}M^T M$ in the lower-right block of size $K_M \times K_M$, and zeroes in the remaining entries.

Proof. It suffices to show that $n^{-1} \sum_{i=1}^n V_i d_{ik} m_{ik'}$ converges in probability to zero for any $k = 1, \dots, K_D$, $k' = 1, \dots, K_M$. Trivially, $E(n^{-1} \sum_{i=1}^n V_i d_{ik} m_{ik'} | \mathcal{F}) = 0$. Through Chebyshev's inequality, we can show that $\text{var}(n^{-1} \sum_{i=1}^n V_i d_{ik} m_{ik'} | \mathcal{F}) \rightarrow 0$ to complete the proof. The proof utilizes Condition 1, along with the Cauchy-Schwarz inequality.

$$\begin{aligned} \text{var} \left(n^{-1} \sum_{i=1}^n V_i d_{ik} m_{ik'} | \mathcal{F} \right) &= n^{-2} \sum_{i=1}^n (d_{ik})^2 (m_{ik'}^2) \\ &\leq n^{-2} \left(\sum_{i=1}^n d_{ik}^4 \right)^{1/2} \left(\sum_{i=1}^n m_{ik'}^4 \right)^{1/2} \leq C/n \end{aligned}$$

by Condition 1, which tends to zero as $n \rightarrow \infty$.

Lemma 2. $n^{-1} \sum_{i=1}^n V_i Y_i d_{ik}$ converges in probability to $\lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n (\ell_{i1} - \ell_{i2}) d_{ik}$ for any $k = 1, \dots, K_D$. Similarly, $n^{-1} \sum_{i=1}^n Y_i m_{ik'}$ converges in probability to $\lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n \Delta_i m_{ik'}$ for any $k' = 1, \dots, K_M$, and $n^{-1} \sum_{i=1}^n V_i d_{ik}$ converges in probability to zero for any $k = 1, \dots, K$.

Proof. We prove the result for $n^{-1} \sum_{i=1}^n V_i Y_i d_{ik}$; the remaining proofs are analogous. First, $E(n^{-1} \sum_{i=1}^n V_i Y_i d_{ik} | \mathcal{F}) = n^{-1} (\ell_{i1} - \ell_{i2}) d_{ik}$. Similar derivations to those in Lemma 1 yield by Condition 1 that $\text{var}(n^{-1} \sum_{i=1}^n V_i Y_i d_{ik} | \mathcal{F})$ converges to zero, which along with Chebyshev's inequality and existence of the limiting value by Condition 2 completes the proof.

Lemma 3. $n^{-1}e^T(I - VH_DV)e$ and $n^{-1}e^T(I - H_A)e$ converge in probability to zero.

Proof. We prove the result for $n^{-1}e^T(I - H_A)e$, and in the process we also prove the result for $n^{-1}e^T(I - VH_DV)e$.

$$\begin{aligned} n^{-1}e^T(I - H_A)e &= 1 - n^{-1} \sum_{i=1}^n a_i (A^T A)^{-1} A^T e \\ &= 1 - n^{-1} \sum_{i=1}^n V_i d_i \{ (n^{-1}D^T D)^{-1} n^{-1}D^T V e \} + o_p(1), \end{aligned}$$

where the $o_p(1)$ term comes from Fact 1, Condition 2, and Lemma 1. By Condition 2 and Chebyshev's inequality, $n^{-1}D^T V e = n^{-1} \sum_{i=1}^n V_i d_i$ converges to a K_D dimensional vector with zeroes in every entry, while $n^{-1}D^T D$ converges to a finite limit, thus completing the proof.

B Proof of Theorem 1

Recall the definition of $\hat{\tau}_{R^*}$ as

$$\hat{\tau}_{R^*} = n^{-1} \sum_{i=1}^n (Y_i - V_i d_i^T \beta_D).$$

We first prove the result for $n^{1/2} \hat{\tau}_{R2}$. Let $\hat{\beta}_D$ and $\hat{\beta}_M$ be vectors of length K_D and K_M corresponding to the sample slopes for the covariates VD and M respectively from a regression of Y on (e, A) . By Lemmas 1 and 2, $\hat{\beta}_D$ and $\hat{\beta}_M$ converge in probability to β_D and β_M . Along with Lemmas 2 and 3, the discrepancy $n^{1/2}(\hat{\tau}_{R2} - \hat{\tau}_{R^*})$ is then of the form

$$n^{1/2}(\hat{\tau}_{R2} - \hat{\tau}_{R^*}) = n^{-1/2} \sum_{i=1}^n m_i^T \hat{\beta}_M + n^{-1/2} \sum_{i=1}^n V_i d_i^T (\hat{\beta}_D - \beta_D) + o_p(1) = o_p(1)$$

since by construction $\sum_{i=1}^n m_i^T \hat{\beta}_M = 0$ by Fact 1, $n^{1/2} \sum_{i=1}^n V_i d_i$ is $O_p(1)$ under Conditions 1 and 2, and $(\hat{\beta}_D - \beta_D)$ is $o_p(1)$. Hence, $n^{1/2}(\hat{\tau}_{R2} - \hat{\tau}_{R^*})$ converges in probability to zero as desired. By Conditions 1 and 2, Lyapunov's central limit theorem holds for $n^{1/2} \hat{\tau}_{R^*}$ at $\delta = 2$, whose asymptotic variance follows from a straightforward calculation. To prove the result for $\hat{\tau}_{R1}$, simply disregard mention of β_M and $\hat{\beta}_M$.

C Proof of Theorem 2

Recall the definitions for S_{R1}^2 and S_{R2}^2 :

$$S_{R1}^2 = (n - K_D - 1)^{-1} \frac{(Y - \hat{\tau}_{R1}e)^T (I - V H_D V) (Y - \hat{\tau}_{R1}e)}{e^T (I - V H_D V) e},$$

$$S_{R2}^2 = (n - K_D - K_M - 1)^{-1} \frac{(Y - \hat{\tau}_{R2}e)^T (I - H_A) (Y - \hat{\tau}_{R2}e)}{e^T (I - H_A) e}.$$

We prove the remark for S_{R2}^2 . Note first that by Lemma 3, $n\{e^T (I - H_A) e\}^{-1}$ converges in probability to 1. Re-expressing S_{R2}^2 using standard identities for residuals from linear regression

$$n S_{R2}^2 = n^{-1} \{1 + o_p(1)\} (Y - \hat{\tau}_{R2}e)^T (I - H_M) (Y - \hat{\tau}_{R2}e) \quad (1)$$

$$- n^{-1} \{1 + o_p(1)\} (Y - \hat{\tau}_{R2}e)^T (I - H_M) V D \{D^T V (I - H_M) V D\}^{-1} D V (I - H_M) (Y - \hat{\tau}_{R2}e). \quad (2)$$

We begin by assessing term (1). By Condition 2, Lemma 2 and Theorem 1, (1) converges in probability to

$$\lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n (\Delta_i - \bar{\tau})^2 + n^{-1} \sum_{i=1}^n (\ell_{i1} - \ell_{i2})^2 - n^{-1} (\Delta - \bar{\tau}e)^T H_M (\Delta - \bar{\tau}e).$$

We now assess term (2). Recall that any term of the form $n^{-1} M^T V D$ converges in probability to a matrix of all zeroes as described in the proof of Lemma 1, and recall that $H_M =$

$M(M^T M)^{-1}M^T$. Hence, by Condition 2, Lemma 2, and Theorem 1, (2) converges in probability to

$$-\lim_{n \rightarrow \infty} n^{-1}(\ell_1 - \ell_2)^T H_D(\ell_1 - \ell_2) = -\beta_D^T \Sigma_D \beta_D.$$

Meanwhile, $\text{var}(n^{1/2}\hat{\tau}_{R2} \mid \mathcal{F}) \rightarrow \lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n (\ell_{i1} - \ell_{i2})^2 - \beta_D^T \Sigma_D \beta_D$ as given in Theorem 1. Subtracting these two quantities, $nS_{R2}^2 - \text{var}(n^{1/2}\hat{\tau}_{R2} \mid \mathcal{F})$ converges in probability to

$$\lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n (\Delta_i - \bar{\tau})^2 - n^{-1}(\Delta - \bar{\tau}e)^T H_M(\Delta - \bar{\tau}e) = \lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n (\Delta_i - \bar{\tau})^2 - \beta_M^T \Sigma_M \beta_M.$$

The proof for S_{R1}^2 follows by replacing $I - H_M$ with I and $\hat{\tau}_{R2}$ with $\hat{\tau}_{R1}$ in (1) and (2) above.

D Deriving the Expressions for the Variance Estimators

We now derive the form of the variance estimators presented in Section 5. We do so for S_{R2}^2 ; the derivation for S_{R1}^2 is analogous. In the classical homoskedastic fixed- X regression setting, a regression of Y on a fixed matrix $X = (e, A)$ yields a true variance for the intercept coefficient

$$\text{var}(\hat{\beta}_0) = \sigma^2 \frac{e^T(I - H_A)e}{(e^T(I - H_A)e)^2} = \sigma^2 \frac{1}{e^T(I - H_A)e}.$$

One then replaces σ^2 with the mean squared error, $\hat{\sigma}^2$, to yield the variance estimator. Let H_X be the projection matrix associated with $X = (e, A)$. $I - H_X$ can be attained by iterative projections, that is

$$I - H_X = I - \frac{(I - H_A)ee^T(I - H_A)}{e^T(I - H_A)e}.$$

Recalling that $\hat{\tau}_{R2} = Y^T(I - H_A)e / (e^T(I - H_A)e)$, the sum of square errors is then of the form

$$\begin{aligned} SSE &= Y^T \left(I - \frac{(I - H_A)ee^T(I - H_A)}{e^T(I - H_A)e} \right) Y \\ &= Y^T Y - \hat{\tau}_{R2} e^T (I - H_A) e \hat{\tau}_{R2}. \end{aligned}$$

Further, $\hat{\tau}_{R2} e^T (I - H_A) e \hat{\tau}_{R2} = e^T (I - H_A) e e^T (I - H_A) Y / (e^T (I - H_A) e) = e^T (I - H_A) Y$. Hence

$$\begin{aligned} SSE &= Y^T Y - \hat{\tau}_{R2} e^T (I - H_A) e \hat{\tau}_{R2} \\ &= (Y - \hat{\tau}_{R2} e)(I - H_A)(Y - \hat{\tau}_{R2} e). \end{aligned}$$

The form for $S_{R2}^2 = SSE / ((n - K_D - K_M - 1)e^T(I - H_A)e)$ then follows.

E Additional simulation results

E.1 Confidence intervals and standard errors for the sample average treatment effect

In each of the settings described in §6.1 of the article, we simulate $S = 1000$ sets of experimental units each containing n matched pairs. Within each sample, we conduct $B = 1000$ randomizations. We conducted our simulation study with experiments of size $n = 25, 50, 100$ and $n = 500$. Given the experiment sizes considered, the extent to which the results presented in this work are applicable for experiments of these sizes is far but certain at the onset of this investigation.

The first three columns characterize the actual coverage properties of 95% confidence intervals for the sample average treatment effect constructed using the standard formulae for the standard errors of $\hat{\tau}_C$, $\hat{\tau}_{R1}$, and $\hat{\tau}_{R2}$ respectively. The second three columns compare the relative values of these standard errors based on these three estimators (and hence, the relative width of confidence intervals across experiments). The final two columns show the relative values for the true root mean squared errors of the regression estimators $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ relative to the classical estimator $\hat{\tau}_C$. The initial number provided in each entry is the median value across the 1000 sets of experimental units, and the endpoints of the interval provided in parenthesis are the 2.5% and 97.5% quantiles across the 1000 samples.

In the parallel response setting with $\mu_T(\cdot) = \mu_C(\cdot)$, the performed inference will be asymptotically exact rather conservative since here the treatment effects are constant. As Table 1 indicates, for all values of n the median coverage of intervals constructed using $\hat{\tau}_C$ and $\hat{\tau}_{R1}$ for $\bar{\tau}^{(s)}$ over the $S = 1000$ sets of experimental units was 95%, with coverage falling in the range (94%, 96%) for 95% of the simulated experiments. For $\hat{\tau}_{R2}$, the median coverage was 94% at $n = 25$, with a 95% range of (89%, 97%) indicating a more drastic deviation from claimed coverage across randomizations in some sets of experimental units with n small. The reason for this is not small sample bias in the estimator $\hat{\tau}_{R2}$ itself, but rather a systematic downward bias of S_{R2} relative to the true value of $\text{var}(\hat{\tau}_{R2} | \mathcal{F})^{1/2}$. In fact, the 95% range for $E(S_{R2} | \mathcal{F}) / \text{var}(\hat{\tau}_{R2} | \mathcal{F})^{1/2}$ across samples was (0.81, 1.05) with $n = 25$, which may justify the use of heteroskedasticity-consistent standard errors such as HC3 in smaller samples (Long and Ervin, 2000). The coverage improves quickly as n increases, and presents little reason for alarm once $n = 100$. The median ratio of both the standard errors and the root mean squared errors from $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ relative to $\hat{\tau}_C$ are slightly over 1/2 across experiments, indicating the potential benefits from estimation and inference using regression adjustment in this setting. As predicted by Theorem 2, the ratios of the standard error estimates based on $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ are clustered about 1 due to the treatment effect being additive, indicating that even in finite samples there is little to be gained from using $\hat{\tau}_{R2}$ over $\hat{\tau}_{R1}$ under additivity.

For the settings with a nonparallel response surface, intervals constructed will be conservative since the treatment effect is no longer additive. Indeed, we see that for all values of n the median coverage across experiments was either 99% or 100% for confidence intervals based on all three estimators. In these settings, we also observe the potential benefits of using the estimator $\hat{\tau}_{R2}$. For $\hat{\tau}_{R1}$, the standard errors used to construct confidence intervals were only slightly smaller than those constructed using $\hat{\tau}_C$, with the median ratio of the two across experiments decreasing from 0.9 for $n = 25$ to 0.85 at $n = 500$. On the other hand, the median ratio of interval widths

Table 1: Classical and regression assisted average treatment effect estimators with a misspecified response function, sample average treatment effect

	True Coverage of 95% CIs			100 × Stand. Error Ratios			100 × RMSE Ratios	
	<i>C</i>	<i>R1</i>	<i>R2</i>	<i>R1 : C</i>	<i>R2 : C</i>	<i>R2 : R1</i>	<i>R1 : C</i>	<i>R2 : C</i>
Parallel								
<i>n</i> = 25	95 (94, 96)	95 (94, 96)	94 (89, 97)	55 (37, 77)	55 (37, 77)	100 (91, 107)	54 (37, 78)	55 (38, 85)
<i>n</i> = 50	95 (94, 96)	95 (94, 96)	94 (92, 96)	54 (41, 69)	53 (41, 67)	99 (94, 101)	54 (41, 69)	55 (41, 72)
<i>n</i> = 100	95 (94, 96)	95 (94, 96)	95 (93, 96)	54 (45, 65)	54 (44, 64)	99 (96, 100)	54 (44, 65)	55 (45, 67)
<i>n</i> = 500	95 (94, 96)	95 (94, 96)	95 (94, 96)	54 (50, 60)	54 (50, 60)	100 (99, 100)	54 (50, 61)	55 (50, 61)
Nonparallel								
<i>n</i> = 25	100 (99, 100)	100 (99, 100)	99 (96, 100)	90 (77, 90)	54 (40, 71)	60 (43, 83)	81 (55, 125)	63 (43, 98)
<i>n</i> = 50	100 (99, 100)	100 (99, 100)	100 (99, 100)	86 (78, 94)	53 (43, 63)	61 (50, 74)	64 (48, 88)	53 (39, 74)
<i>n</i> = 100	100 (100, 100)	100 (100, 100)	100 (100, 100)	86 (80, 90)	53 (46, 60)	62 (53, 71)	55 (45, 69)	48 (39, 61)
<i>n</i> = 500	100 (100, 100)	100 (100, 100)	100 (100, 100)	85 (82, 97)	53 (50, 56)	62 (53, 71)	45 (40, 51)	44 (39, 50)

constructed using $\hat{\tau}_{R2}$ relative to those using $\hat{\tau}_C$ was slightly over 1/2, while the median ratio relative to those constructed based on $\hat{\tau}_{R1}$ was 0.6. To emphasize, this result does not stem from an inherently lower variance for the estimator $\hat{\tau}_{R2}$, as reflected in the ratios of the true root mean square errors for estimation using $\hat{\tau}_C$ relative to $\hat{\tau}_{R1}$ and relative to $\hat{\tau}_{R2}$ shown in Table 1. Indeed, the mean squared errors for $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ are equal in the limit by Theorem 1. Rather, as Corollary 1 describes, this stems from S_{R2}^2 being a tighter upper bound for $\text{var}(\hat{\tau}_{R2} | \mathcal{F})$ than S_{R1}^2 is for $\text{var}(\hat{\tau}_{R1} | \mathcal{F})$. Hence, intervals based on $\hat{\tau}_{R2}$ are shorter than those constructed using $\hat{\tau}_{R1}$ while maintaining the proper coverage.

E.2 Coverage in a superpopulation simulation

In this section, we modify the simulation setting described in §6.1 of the manuscript to instead assess coverage for the population average treatment effect. We keep the same functional forms for the response functions under treatment and control, and maintain the same misspecified form for the covariates available to the practitioner. In the sth of $S = 10000$ sets of experimental units, we simply conduct one random allocation to treatment or control in accordance with the paired design. We then compute the estimators in question, corresponding standard errors, and confidence intervals. Rather than attempting to cover $\bar{\tau}^{(s)}$, the intervals are now assessed based on covering $E\{\bar{\tau}^{(s)}\} = 0$. For $\hat{\tau}_{R2}$, we compute intervals based on both S_{R2}^2 and its superpopulation modification $S_{R2,P}^2 = S_{R2}^2 + n^{-1}\hat{\beta}_M^T \hat{\Sigma}_M \hat{\beta}_M$. We do so for $n = 25$ and $n = 1000$.

Table 2 present the results. In the parallel response setting, the results are quite similar to those presented in §E.1, where we instead constructed confidence intervals for the sample average treatment effect. The reason is that the treatment effect was constant for all individuals, such that $\text{var}\{\bar{\tau}^{(c)}\} = 0$ in the superpopulation simulation. In the case of constant effects S_{R1}^2 , S_{R2}^2 , and $S_{R2,P}^2$ are asymptotically equivalent, and S_{R2}^2 can be used to yield valid confidence intervals in this superpopulation setting. The estimators $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ remain asymptotically equivalent, and continue to improve efficiency over $\hat{\tau}_C$.

As the nonparallel case illustrates, validity of S_{R2}^2 in the superpopulation context described is unique to the case of constant treatment effects. In this simulation, effect heterogeneity was present. S_{R2}^2 attains its smaller value relative to S_{R1}^2 by means of exploiting effect heterogeneity. Yet as we see, S_{R1}^2 yields confidence intervals that meet their coverage guarantee, while S_{R2}^2 creates confidence intervals that fall well short. As the table presents, the ratio of the average standard error S_{R2} to the true standard deviation of $\hat{\tau}_{R2}$ is far below 1, further indicative of S_{R2}^2 being anticonservative. The downwards bias does not disappear asymptotically, as the simulation study with $n = 1000$ seeks to emphasize. The corrected version, $S_{R2,P}^2$, adds back a consistent estimate of the downwards bias in S_{R2}^2 , which reestablishes proper coverage for intervals using $\hat{\tau}_{R2}$. Once again, $\hat{\tau}_{R1}$ and $\hat{\tau}_{R2}$ behave quite similarly in terms of estimator variance (indicative of their asymptotic equivalence), and both dominate the classical difference-in-means $\hat{\tau}_C$.

Table 2: Classical and regression assisted average treatment effect estimators with a misspecified response function, population average treatment effect. The table shows coverage of 95% normal-based intervals, the ratio of the average estimated standard error to the true standard deviation of the estimator (in parentheses), and the ratio of the standard deviations of the estimators.

	True Coverage of 95% CIs				100 × RMSE Ratios		
	(100 × Estimated SE / True SE)				$R2 : R1$	$R1 : C$	$R2 : C$
	C	$R1$	$R2$	$R2, P$			
Parallel							
$n = 25$	95.0 (98.0)	95.2 (99.7)	93.8 (92.7)	96.8 (104)	107	54.1	58.1
$n = 1000$	94.8 (99.3)	95.0 (99.5)	94.9 (99.2)	95.0 (99.5)	100	54.1	54.3
Nonparallel							
$n = 25$	95.3 (99.0)	94.3 (96.3)	77.2 (59.3)	95.6 (99.8)	97.9	92.2	90.2
$n = 1000$	95.1 (100)	95.2 (100)	78.2 (62.8)	95.1 (101)	99.9	84.7	84.6

E.3 A numerical example

The following R code simulates an experiment in the nonparallel simulation setting from §6.1 of the manuscript with $I = 25$, showing the numerical values for the treatment effect estimators and the corresponding standard errors.

```

library(MASS)
#Nonparallel Setting, n=25 Pairs
set.seed(01562)
#Set number of pairs
n = 25
#Simulate (unknown) covariates W, error, and potential outcomes
w1 = mvrnorm(n, c(0,0,0,0), diag(4))
w2 = w1 + mvrnorm(n, c(0,0,0,0), .25*diag(4))
e1 = rnorm(n, 0, 1)
e2 = rnorm(n, 0, 1)
rt1 = w1%*(c(27.4,13.7,13.7, 13.7)) + e1
rc1 = w1%*(c(13.7,13.7,3, 27.4)) + e1
rt2= w2%*(c(27.4,13.7,13.7, 13.7)) + e2
rc2= w2%*(c(13.7,13.7,3, 27.4)) + e2

```



```

#Form Delta (pairwise average of treatment effects)
tau1 = rt1-rc1
tau2 = rt2-rc2
Delta = (tau1+tau2)/2

#Form eta (difference in levels)
l1 = (rt1+rc1)/2
l2 = (rt2+rc2)/2
eta = l1-l2

#Form the observed covariates as nonlinear functions of W
X1 = cbind(exp(w1[,1]/2), w1[,2]/(1+exp(w1[,1])) + 10,
            (w1[,1]*w1[,3]/25 + .6)^3, (w1[,2] + w1[,4] + 20)^2)
X2 = cbind(exp(w2[,1]/2), w2[,2]/(1+exp(w2[,1])) + 10,
            (w2[,1]*w2[,3]/25 + .6)^3, (w2[,2] + w2[,4] + 20)^2)

#Form D and M based on X
D = X1-X2
D = as.matrix(D)
M = (X1+X2)/2 - (rep(1, n))%*%t(colMeans(rbind(X1, X2)))
KD = ncol(D)
KM = ncol(M)

#Sample and Population ATE in simulation
SATE = (sum(rt1-rc1) + sum(rt2-rc2))/(2*I)
PATE = 0

#Randomly assign treatment across pairs; form V
Z = runif(n) < .5
V = 2*Z-1

#Here are the treated-minus-control response differences
Y = Delta + V*eta

#Here are the treated-minus-control covariate differences
VD = diag(V)%*%D

#Form the classical estimator and its standard error
tauhat_C = mean(Y)
SE_C = sd(Y)/sqrt(n)

#Now, form tauhat_R1 with its standard error
reg1 = lm(Y~VD)

```

```

tauhat_R1 = reg1$coef[1]
SE_R1 = summary(reg1)$coef[1,2]

#Now, tauhat_R2. The difference is the inclusion of M in
#the regression
reg2 = lm(Y~VD+M)
tauhat_R2 = reg2$coef[1]
SE_R2 = summary(reg2)$coef[1,2]

#Form the superpopulation version of SE_R2
betaM = reg2$coef[6:9]
SE_R2P = sqrt(SE_R2^2 + t(betaM)%*%cov(M)%*%betaM/n)

#Compare the point estimates and standard errors
compare25 = rbind(c(tauhat_C, tauhat_R1, tauhat_R2, tauhat_R2),
                  c(SE_C, SE_R1, SE_R2, SE_R2P))
colnames(compare25) = c("tauhat_C", "tauhat_R1", "tauhat_R2",
                        "tauhat_R2P")
rownames(compare25) = c("Estimate", "Standard Error")
compare25

#Output
#           tauhat_C tauhat_R1 tauhat_R2 tauhat_R2P
# Estimate      3.640938 -1.884071 -2.728688 -2.728688
# Standard Error 5.483459  3.935346  2.966589  4.077766

```

References

- Aronow, P. M., Green, D. P., and Lee, D. K. (2014). Sharp bounds on the variance in randomized experiments. *The Annals of Statistics*, 42(3):850–871.
- Aronow, P. M. and Middleton, J. A. (2013). A class of unbiased estimators of the average treatment effect in randomized experiments. *Journal of Causal Inference*, 1(1):135–154.
- Bloniarz, A., Liu, H., Zhang, C.-H., Sekhon, J., and Yu, B. (2016). Lasso adjustments of treatment effect estimates in randomized experiments. *Proceedings of the National Academy of Sciences*, 113(27):7383–7390.
- Freedman, D. A. (2008). On regression adjustments to experimental data. *Advances in Applied Mathematics*, 40(2):180–193.
- Greevy, R., Lu, B., Silber, J. H., and Rosenbaum, P. (2004). Optimal multivariate matching before randomization. *Biostatistics*, 5(2):263–275.

- Imai, K. (2008). Variance identification and efficiency analysis in randomized experiments under the matched-pair design. *Statistics in Medicine*, 27(24):4857–4873.
- Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- Kang, J. D. and Schafer, J. L. (2007). Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science*, pages 523–539.
- Lin, W. (2013). Agnostic notes on regression adjustments to experimental data: Reexamining Freedman’s critique. *The Annals of Applied Statistics*, 7(1):295–318.
- Long, J. S. and Ervin, L. H. (2000). Using heteroscedasticity consistent standard errors in the linear regression model. *The American Statistician*, 54(3):217–224.
- Neyman, J. (1923). On the application of probability theory to agricultural experiments. Essay on principles. Section 9 (in Polish). *Roczniki Nauk Rolniczych*, X:1–51. Reprinted in *Statistical Science*, 1990, 5(4):463–480.
- Rosenbaum, P. R. (2002). Covariance adjustment in randomized experiments and observational studies. *Statistical Science*, 17(3):286–327.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701.
- 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. *Statistics in medicine*, 27(23):4658–4677.
- Zhang, M., Tsiatis, A. A., and Davidian, M. (2008). Improving efficiency of inferences in randomized clinical trials using auxiliary covariates. *Biometrics*, 64(3):707–715.