

Unifying Design-based Inference: A New Variance Estimation Principle

WORKING PAPER 2 OF 4

Joel A. Middleton*

May 27, 2020

*Charles and Louise Travers Department of Political Science, *University of California, Berkeley*.
email: joel.middleton@gmail.com

1 Introduction

In Unifying paper 1 of 4 a generalized sandwich variance estimators was proposed that is applicable to virtually any experimental design, any linear estimator and any variance bound (defined therein), subsuming other well known sandwich variance estimators (which go by names such as “robust”, “cluster robust”, “heteroskedastic consistent”, “sandwich”, “White”, “Huber-White”, “HC”, “CR”, etc.). This paper presents the Obloženè Chlebižky (OC) variance estimators as a novel alternative to the generalized sandwich in Paper 1 of 4. That the OC concept is unlikely to arise from other, more standard, frameworks is manifestly true in light of the 40 year lacuna since White (1980).

For any member of the generalized sandwich variance estimator class, there is an OC with the same expected value that can be used as a substitute. A key insight in the derivation is that a sandwich may be turned inside-out, written as a quadratic in the (random) outcome vector (or, equivalently in many cases, as a quadratic in residuals). The symmetric matrix at the center of rearranged form is random, but written this way it is free of the outcome. The central idea behind the OC is to identify, by way of computation, a nonrandom substitute for the random matrix such that the expected value of OC is equal to that of the sandwich.

There is heuristic appeal in eliminating variance components attributable to the matrix, because it suggests that the OC may be more precise than the sandwich upon which it is based. While this is a reasonable rule-of-thumb, it need not be true in general, and so a checkable condition is provided. When it holds for a realized data set, the researcher can be certain that OC has lower variance across the randomization distribution. In some simple simulations provided, reductions in variance of variance estimators is on the order of 1/2.

2 Notation

Consider a randomized experiment with k treatment arms. The Neyman causal model (NCM) assumes that the units in the experimental study represent a finite population of size n . For a given outcome measure, call it y , each unit, i , responds with one of k possible values in $\{y_{1i}, y_{2i}, \dots, y_{ki}\}$, depending on their treatment assignment. The possible responses are referred to as the *potential outcomes*. In the NCM these values are considered (nonrandom) constants, which stands in contrast to other, more common, formulations where potential outcomes are assumed to be sampled from some (possibly nonparametric) joint distribution.

The only random elements in the NCM are the treatment assignment indicators $\{R_{1i}, R_{2i}, \dots, R_{ki}\}$, and they determine which potential outcome will be observed by the researcher. Since a unit, i , can only be assigned to one arm of the experiment these random indicators must sum to unity, i.e., $R_{1i} + R_{2i} + \dots + R_{ki} = 1$. A standard representation of the *observed* response for the i^{th} unit under the NCM would be,

$$Y_i^{obs} = y_{1i}R_{1i} + y_{2i}R_{2i} + \dots + y_{ki}R_{ki},$$

which is itself random, due to the assignment indicators. The observed data can then be represented as $\{Y_i^{obs}, R_{1i}, R_{2i}, \dots, R_{ki}, x_i\}_{\forall i}$, where x_i is an additional vector of k covariates. Like the potential outcomes, x_i is considered to be nonrandom, and, unlike the potential outcomes, does not depend on the assignment, which might be ensured, in practice, by collecting the covariates before the assignment.

Ideally, a researcher might like to know the difference between responses under various arms for a given individual, i

for example, $y_{2i} - y_{1i}$, or perhaps $\frac{1}{2}(y_{2i} - y_{1i} + y_{4i} - y_{3i})$ and so on, which are examples of different *contrasts* between potential outcomes.

However, it is clear from the definition of Y_i^{obs} that individual treatment effects are not identified since only one of the two potential outcomes can be observed. This is known as *fundamental problem of causal inference* (Holland, 1986). As a result, we often study averages of these values over the units of study. To continue with the above examples, a researcher might be satisfied to estimate averages over these individual treatment effects. Continuing the example, we have the average treatment effects $n^{-1} \sum_i (y_{2i} - y_{1i})$ and $n^{-1} \sum_i \frac{1}{2} (y_{2i} - y_{1i} + y_{4i} - y_{3i})$, respectively.

To simplify notation, let y_1, y_2, \dots, y_k represent length n vectors of potential outcomes associated with each of the arms, with the i^{th} element of each corresponding to the i^{th} unit. Next, define

$$y := (y'_1 \ y'_2 \ \dots \ y'_k)',$$

which is the length kn , representing all k potential outcomes for each of the n units.

Next, if we let $\mathbf{1}_n$ be a n -length vector of ones, then a $kn \times k$ *intercept matrix* can be defined as,

$$\mathbf{1} := \begin{bmatrix} \mathbf{1}_n & & & \\ & \mathbf{1}_n & & \\ & & \ddots & \\ & & & \mathbf{1}_n \end{bmatrix},$$

which, for example, allows us to express a k -length vector of means of the arms as $n^{-1} \mathbf{1}' y$ or, equivalently, $(\mathbf{1}' \mathbf{1})^{-1} \mathbf{1}' y$.

Given a length- k *contrast vector*, c , we can contrast the means for the various arms. In keeping with the running example with four assignment arms, if one is interested in the average difference in responses to the first two treatments, then $c = (-1 \ 1 \ 0 \ 0)'$ and we have $n^{-1} c' \mathbf{1}' y = n^{-1} \sum_i (y_{2i} - y_{1i})$. Likewise, if one chooses $c = (-\frac{1}{2} \ \frac{1}{2} \ -\frac{1}{2} \ \frac{1}{2})'$ then $n^{-1} c' \mathbf{1}' y = n^{-1} \sum_i \frac{1}{2} (y_{2i} - y_{1i} + y_{4i} - y_{3i})$. Sensible contrasts such as these meet the conditions $\sum_{j=1}^k c_j = 0$ and $\sum_{j=1}^k |c_j| = 2$, though this would not be strictly necessary, mathematically speaking.

Next define an $n \times n$ diagonal matrix that has all n assignment indicators for treatment arm 1 on the diagonal,

$$\mathbf{R}_1 := \begin{bmatrix} R_{11} & & & \\ & R_{12} & & \\ & & \ddots & \\ & & & R_{1i} & \\ & & & & \ddots & \\ & & & & & R_{1n} \end{bmatrix},$$

and define $\mathbf{R}_2, \mathbf{R}_3, \dots, \mathbf{R}_k$ analogously. Arrange these matrices to create the diagonal $kn \times kn$ matrix

$$\mathbf{R} := \begin{bmatrix} \mathbf{R}_1 & & & \\ & \mathbf{R}_2 & & \\ & & \ddots & \\ & & & \mathbf{R}_k \end{bmatrix}$$

and note the a $kn \times kn$ diagonal matrix of assignment probabilities can be written as $\boldsymbol{\pi} := \mathbb{E}[\mathbf{R}]$, with the first n diagonal elements representing probabilities of assignment to arm 1, then the next n diagonal elements are probabilities of assignment to arm 2 and so on.

In this alternative notation the researcher can be said to observe the assignment, \mathbf{R} , the observed vector of outcomes, $\mathbf{R}y$, and also a matrix of l pre-treatment covariates, \mathbf{x} , which has size $n \times l$. In a randomized experiment $\boldsymbol{\pi}$ is also observed (known) in many cases. When intractable analytically, however, it might be estimated to arbitrary precision by repeating the original randomization until a target level of precision is achieved.

For covariate adjusted estimators, it will also be convenient to define the $kn \times (k + l)$ matrix,

$$\mathbf{z} := \begin{bmatrix} \mathbf{1}_n & & & \mathbf{x} \\ & \mathbf{1}_n & & \mathbf{x} \\ & & \ddots & \vdots \\ & & & \mathbf{1}_n & \mathbf{x} \end{bmatrix}.$$

Remark 1. For some cases, such adjusting for covariates separately by arm, it might be useful to define \mathbf{z} with \mathbf{x} matrices arranged along a block-diagonal. In that case, it is prudent to stipulate that \mathbf{x} have columns that sum to zero to avoid problems of coefficient interpretation (cf. Lin, 2013; Middleton, 2018). This will be discussed further in paper 3 of 4.

3 The class of WLS estimators

This paper will discuss the class of estimators that can be characterized as weighted least squares (WLS), which includes the difference-of-means, OLS and Hajek estimators as special cases. For the purposes of giving asymptotic variance expressions for this class, the Horvitz-Thompson estimator is also introduced. Paper 3 of 4 will introduce a wider class of estimators that will include these as special cases, and specifically examine their properties as point estimators. This paper is primarily concerned with a new approach to variance bound estimation, irrespective of what the estimator estimates.

As in Paper 1 of 4, *linear estimators* are defined as having the form,

$$\widehat{\delta}_c := c' \mathbf{W} \mathbf{R} y, \tag{1}$$

where \mathbf{W} a matrix with kn columns and k rows if it is an unadjusted estimator and $k + l$ rows if it is a covariate adjusted estimator. The length of the contrast vector, c , is equal to the number of rows in \mathbf{W} . The first k entries of c are the contrast values, followed by l zeros in the case of a covariate adjusted estimator.

Definition 3.1 (Horvitz-Thompson estimator). *The Horvitz-Thompson estimator written as in equation (1) with,*

$$\mathbf{W} = \mathbf{w}^{\text{HT}} := (\mathbf{1}' \mathbf{1})^{-1} \mathbf{1}' \boldsymbol{\pi}^{-1}.$$

Definition 3.2 (WLS estimators). *WLS estimators can be written as in equation (1) with,*

$$\mathbf{W} = \mathbf{W}^{\text{WLS}} := (\mathbf{z}' \mathbf{m} \mathbf{R} \mathbf{z})^{-1} \mathbf{z}' \mathbf{m}$$

where the $kn \times kn$ matrix \mathbf{m} must have nonnegative entries on the diagonal (with some strictly positive) and zeros elsewhere. Also, define

$$\begin{aligned}\mathbf{w}^{\text{WLS}} &:= \mathbf{W}^{\text{WLS}}|_{\mathbf{R}=\boldsymbol{\pi}} \\ &= (\boldsymbol{\varkappa}' \mathbf{m} \boldsymbol{\pi} \boldsymbol{\varkappa})^{-1} \boldsymbol{\varkappa}' \mathbf{m},\end{aligned}$$

which is \mathbf{W}^{WLS} with $\boldsymbol{\pi} = \text{E}[\mathbf{R}]$ replacing \mathbf{R} .

Remark 2. WLS gives OLS as a special case when $\mathbf{m} = \mathbf{i}_{kn}$ (\mathbf{i}_{kn} is the identity matrix). If $\mathbf{m} = \mathbf{i}_{kn}$ and, in addition, $\boldsymbol{\varkappa} = \mathbf{1}$ (there are no covariates), WLS is equivalent to the difference-of-means. If $\mathbf{m} = \boldsymbol{\pi}^{-1}$ and $\boldsymbol{\varkappa} = \mathbf{1}$, then it is the Hajek estimator. The covariate adjusted WLS with $\mathbf{m} = \boldsymbol{\pi}^{-1}$ will be discussed further in paper 3 of 4, because it is algebraically equivalent to the generalized regression estimator introduced there.

Theorem 3.3. For a suitably chosen constant, a , and vector of constants, z_c , the first-order Taylor approximation for a WLS can be written as,

$$\widehat{\delta}_c^{\text{T(WLS)}} = a + n \mathbf{1}'_{(k+l)} \mathbf{w}^{\text{HT}} \mathbf{R} z_c,$$

with

$$z_c := \boldsymbol{\pi} \text{diag}(u) \mathbf{w}^{\text{WLS}'},$$

and where

$$u := y - \boldsymbol{\varkappa} b^{\text{WLS}}$$

is the “true” residual and

$$b^{\text{WLS}} := \mathbf{w}^{\text{WLS}} \boldsymbol{\pi} y$$

is the “true” WLS coefficient.

Remark 3. By “true” it is meant to suggest that these are nonrandom quantities that, for large enough samples, represent approximately the centers of the distributions of observed residuals and WLS coefficient, respectively.

Remark 4. While the vectors z_c , u and b^{WLS} are not directly observable, the result is useful because it shows that the first-order Taylor approximation of a linear estimator is a Horvitz-Thompson estimator with contrast vector $c = n \mathbf{1}_{(k+l)}$ and “outcome” vector z_c . Expressing its variance will provide a basis for an asymptotically valid approximation.

Corollary 3.3.1 (The Obloženè Chlebìžky Principle). In Theorem 3.3 we could have also given an equivalent expression for, z_c , as,

$$\begin{aligned}z_c^{\text{equiv}} &:= \boldsymbol{\pi} \text{diag}(c' \mathbf{w}^{\text{WLS}}) u \\ &= z_c.\end{aligned}$$

Remark 5. Corollary 3.3.1 gives the key insight in the development of the Obloženè Chlebìžky variance estimators. Residuals need not be in the center of the sandwich.

4 Variance bounds and their estimation

As in Paper 1 of 4, define $\mathbf{1}_{kn}$ as a kn -length vector of ones and the $kn \times kn$ “design matrix” as,

$$\mathbf{d} := \mathbf{V} (\mathbf{1}'_{kn} \boldsymbol{\pi}^{-1} \mathbf{R}). \quad (2)$$

This is a variance-covariance matrix of weighted treatment assignments.

An exact expression for first-order Taylor approximations of linear estimators can be written as,

$$\mathbf{V} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) = z'_c \mathbf{d} z_c. \quad (3)$$

Definition 4.1 (Bounding Matrix). *Let y be an outcome vector and let the contrast vector be $c = n\mathbf{1}_k$. Then the associated Horvitz-Thompson estimator is $n\mathbf{1}_k \mathbf{w}^{\text{HT}} \mathbf{R} y$ with variance $y' \mathbf{d} y$. Then, the arbitrary $kn \times kn$ matrix, $\tilde{\mathbf{d}}$, is a “bounding matrix” if, for all $y \in \mathbf{R}^{kn}$, $y' \mathbf{d} y \leq y' \tilde{\mathbf{d}} y$.*

Remark 6. *For more on bounding matrices, see Paper 1 of 4. Unless otherwise noted, the Generalized Neyman Bounding matrix, $\tilde{\mathbf{d}} = \tilde{\mathbf{d}}^{\text{N}}$, will be assumed.*

Now, define the $2n \times 2n$ matrix of probabilities and joint probabilities of assignment,

$$\mathbf{p} := \mathbf{E} [\mathbf{R} \mathbf{1}_{2n} \mathbf{1}'_{2n} \mathbf{R}].$$

Next define an inverse probability weighted version of bounding matrix, $\tilde{\mathbf{d}}$, as

$$\tilde{\mathbf{d}}_{/p} := \tilde{\mathbf{d}} / \mathbf{p} \quad (4)$$

with $/$ denoting element-wise division defined such that division by zero equals zero.

Hypothetically, if the observed outcome vector was $\mathbf{R} z_c$, then the Horvitz-Thompson estimator with $c = n\mathbf{1}_k$ has variance $z'_c \tilde{\mathbf{d}} z_c$, which could be estimated unbiasedly by

$$\widehat{\mathbf{V}} \left(\widehat{\delta}_{n\mathbf{1}_k}^{\text{HT}} \right) := z'_c \mathbf{R} \tilde{\mathbf{d}}_{/p} \mathbf{R} z_c. \quad (5)$$

Of course, the vector z_c is defined by quantities which are not directly observed. However, we might identify random quantities that approximate z_c and u as

$$\begin{aligned} Z_c &:= \boldsymbol{\pi} \text{diag}(U) \mathbf{W}^{\text{WLS}' c} \\ &= \boldsymbol{\pi} \text{diag}(c' \mathbf{W}^{\text{WLS}}) U \end{aligned}$$

and

$$U := y - \mathbf{x} \mathbf{W}^{\text{WLS}} \mathbf{R} y,$$

respectively.

Next, consider the generalized sandwich,

$$\begin{aligned} \widehat{\mathbf{V}}^{\text{GS}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) &:= Z'_c \mathbf{R} \tilde{\mathbf{d}}_{/p} \mathbf{R} Z_c \\ &= c' \mathbf{W}^{\text{WLS}} \text{diag}(\mathbf{R} U) \left(\boldsymbol{\pi} \tilde{\mathbf{d}}_{/p} \boldsymbol{\pi} \right) \text{diag}(\mathbf{R} U) \mathbf{W}^{\text{WLS}' c} \\ &= y' \left\{ \mathbf{H}' \text{diag}(c' \mathbf{W}^{\text{WLS}}) \left(\boldsymbol{\pi} \tilde{\mathbf{d}}_{/p} \boldsymbol{\pi} \right) \text{diag}(c' \mathbf{W}^{\text{WLS}}) \mathbf{H} \right\} y \\ &= y' \mathbf{O} y \end{aligned} \quad (6)$$

where $\mathbf{H} := (\mathbf{R} - \mathbf{R}\mathbf{x}\mathbf{W}^{\text{WLS}}\mathbf{R})$ is a “residual maker” matrix and \mathbf{O} denotes the random matrix in curly brackets in the line just above. The second line gives a form that may be most recognizable in canonical sandwich. The third line uses the Obloženè Chlebìžky principle from Corollary 3.3.1 and the equality $\mathbf{R}\mathbf{U} = \mathbf{H}\mathbf{y}$.

Remark 7. *The generalized sandwich is equivalent to Eicker-Huber-White variance estimator under Bernoulli designs and $\mathbf{m} = \mathbf{i}_{kn}$. Likewise it is algebraically equivalent to “cluster robust” standard errors (CR0) under Bernoulli assignment of clusters and $\mathbf{m} = \mathbf{i}_{kn}$ (see Paper 1 of 4 for more details). Refinements for degrees-of-freedom and leverage are easily accommodated, and simulations will compare sandwiches HC0, HC1, and HC2 to the parallel OC estimators.*

5 Obloženè Chlebìžky variance bound estimators

The last line of equation (6) has used the Obloženè Chlebìžky principle from Corollary 3.3.1 to rearrange terms in the generalized sandwich. Written like this, it is easy to see that the expectation of the generalized sandwich is,

$$\begin{aligned} \mathbb{E} \left[\widehat{\mathbf{V}}^{\text{GS}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) \right] &= \mathbf{y}' \mathbb{E}[\mathbf{O}] \mathbf{y} \\ &= \mathbf{y}' \mathbf{o} \mathbf{y}, \end{aligned} \tag{7}$$

with $\mathbf{o} := \mathbb{E}[\mathbf{O}]$. The matrix \mathbf{o} may be computable directly for small samples. For larger samples, it may be simulated to arbitrary precision by drawing from the randomization distribution.

5.1 OC0: An alternative variance estimator

Next, an estimator with the same expected value as the generalized sandwich, given in equation 6, is the OC0 estimator,

$$\widehat{\mathbf{V}}^{\text{OC0}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) := \mathbf{y}' \mathbf{R} \mathbf{o}_{\mathbf{p}} \mathbf{R} \mathbf{y},$$

with $\mathbf{o}_{\mathbf{p}} := \mathbf{o}/\mathbf{p}$ where, as above, “/” is element-wise division with division by zero resolving to zero.

OC0 is a Horvitz-Thompson estimator, and therefore lacks invariance to location shifts in the outcome variable. Likewise, its variance may depend on the scaling of y and so it may also be less precise than covariate-adjusted alternatives. A useful heuristic is that Horvitz-Thompson estimators may have higher variance when the outcomes have a mean that is far from zero, though this is not always true. See, for example, in the case of probability proportional to size sampling.

5.2 OC0.B: Adjustment for precision achieves invariance, introduces bias

To improve upon OC0, consider the covariate adjusted alternative,

$$\widehat{\mathbf{V}}^{\text{OC0.B}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) := \left(\mathbf{y} - \mathbf{x} \widehat{\mathbf{b}}^{\text{WLS}} \right)' \mathbf{R} \mathbf{o}_{\mathbf{p}} \mathbf{R} \left(\mathbf{y} - \mathbf{x} \widehat{\mathbf{b}}^{\text{WLS}} \right),$$

with $\widehat{b}^{\text{WLS}} := \mathbf{W}^{\text{WLS}} \mathbf{R} y$ giving the WLS coefficient. The covariate adjustment mean that OC0.B is more precise than OC0, and OC0.B is also invariant to location shifts in y . It is, nonetheless, biased.

The bias relative to the expected value of the generalized sandwich given in equation (7) is

$$\begin{aligned} \mathbf{B} \left[\widehat{\mathbf{V}}^{\text{OC0.B}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) \right] &= \mathbb{E} \left[\widehat{\mathbf{V}}^{\text{OC0.B}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) - y' \mathbf{o} y \right] \\ &= y' \left\{ \mathbb{E} \left[\mathbf{R} \mathbf{W}^{\text{WLS}'} \mathbf{z} \left(\mathbf{R} \mathbf{o}_{\mathbf{p}} \mathbf{R} \right) \mathbf{z} \mathbf{W}^{\text{WLS}} \mathbf{R} \right. \right. \\ &\quad \left. \left. - 2 \left(\mathbf{R} \mathbf{o}_{\mathbf{p}} \mathbf{R} \right) \mathbf{z} \mathbf{W}^{\text{WLS}} \mathbf{R} \right] \right\} y \\ &= y' \mathbf{b}^{(0)} y \end{aligned}$$

with $\mathbf{b}^{(0)}$ in the last line defined as the $kn \times kn$ matrix given inside the curly brackets in the line before. Similar to \mathbf{o} , the matrix $\mathbf{b}^{(0)}$ may be computed exactly for small experiments, or to arbitrary precision by repeating random draws from the assignment distribution.

5.3 OC1: Bias correction reintroduces non-invariance

Note that OC0.B is now invariant and we might expect it to have a lower variance, but it also biased. Therefore consider,

$$\begin{aligned} \widehat{\mathbf{V}}^{\text{OC1}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) &:= \widehat{\mathbf{V}}^{\text{OC0.B}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) \\ &\quad - y' \mathbf{R} \mathbf{b}_{\mathbf{p}}^{(0)} \mathbf{R} y, \end{aligned}$$

where the second term is a bias correction for OC0.B.

Noting that the matrix $\mathbf{b}_{\mathbf{p}}^{(0)}$ is an inverse-probability weighted version of $\mathbf{b}^{(0)}$, we again have $\mathbb{E} \left[\mathbf{R} \mathbf{b}_{\mathbf{p}}^{(0)} \mathbf{R} \right] = \mathbf{b}^{(0)}$, so that the second term in OC1 is an unbiased estimator of the bias of OC0.B. Therefore, OC1 is an unbiased estimator of the expectation of the generalized sandwich variance given in equation (6). However, the second term reintroduces non-invariance.

5.4 OC1.B: Reestablishing invariance reintroduces bias

Once again iterating to achieve invariance through covariate adjustment we have,

$$\begin{aligned} \widehat{\mathbf{V}}^{\text{OC1.B}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) &:= \widehat{\mathbf{V}}^{\text{OC0.B}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) \\ &\quad - \left(y - \mathbf{z} \widehat{b}^{\text{WLS}} \right) \mathbf{R} \mathbf{b}_{\mathbf{p}}^{(0)} \mathbf{R} \left(y - \mathbf{z} \widehat{b}^{\text{WLS}} \right) \\ &= \left(y - \mathbf{z} \widehat{b}^{\text{WLS}} \right) \mathbf{R} \left(\mathbf{o}_{\mathbf{p}} - \mathbf{b}_{\mathbf{p}}^{(0)} \right) \mathbf{R} \left(y - \mathbf{z} \widehat{b}^{\text{WLS}} \right), \end{aligned}$$

which, again, has some bias relative to the generalized sandwich given in equation (6).

As before, write,

$$\begin{aligned}
\mathbf{B} \left[\widehat{\mathbf{V}}^{\text{OC1.B}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) \right] &= \mathbf{E} \left[\widehat{\mathbf{V}}^{\text{OC1.B}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) - \widehat{\mathbf{V}}^{\text{OC1}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) \right] \\
&= y' \left\{ \mathbf{E} \left[-\mathbf{R}\mathbf{W}^{\text{WLS}'\mathbf{x}} \left(\mathbf{R}\mathbf{b}_{\text{p}}^{(0)}\mathbf{R} \right) \mathbf{x}\mathbf{W}^{\text{WLS}}\mathbf{R} \right. \right. \\
&\quad \left. \left. + 2 \left(\mathbf{R}\mathbf{b}_{\text{p}}^{(0)}\mathbf{R} \right) \mathbf{x}\mathbf{W}^{\text{WLS}}\mathbf{R} \right] \right\} y \\
&= y'\mathbf{b}^{(1)}y,
\end{aligned}$$

and we see that we could continue this process of iteration to achieve convergence, i.e., a variance estimator that has the same expected value as the generalized sandwich but which is potentially more precise.

5.5 OCJ and OCJ.B: Iterating $J + 1$ times

In the manner above, iterate $J + 1$ times to obtain,

$$\begin{aligned}
\widehat{\mathbf{V}}^{\text{OCJ}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) &:= \left(y - \mathbf{x}\widehat{\mathbf{b}}^{\text{WLS}} \right) \mathbf{R} \left(\mathbf{o}_{\text{p}} - \sum_{j=0}^{J-1} \mathbf{b}_{\text{p}}^{(j)} \right) \mathbf{R} \left(y - \mathbf{x}\widehat{\mathbf{b}}^{\text{WLS}} \right), \\
&\quad - y\mathbf{R}\mathbf{b}_{\text{p}}^{(J)}\mathbf{R}y
\end{aligned}$$

where the last term introduces some degree of invariance but insures unbiasedness, or

$$\widehat{\mathbf{V}}^{\text{OCJ.B}} \left(\widehat{\delta}_c^{\text{T(WLS)}} \right) := \left(y - \mathbf{x}\widehat{\mathbf{b}}^{\text{WLS}} \right) \mathbf{R} \left(\mathbf{o}_{\text{p}} - \sum_{j=0}^J \mathbf{b}_{\text{p}}^{(j)} \right) \mathbf{R} \left(y - \mathbf{x}\widehat{\mathbf{b}}^{\text{WLS}} \right)$$

has a bias but ensures invariance. OCJ and OCJ.B differ only in whether the term associated with $\mathbf{b}_{\text{p}}^{(J)}$ is a quadratic in y or in the regression residuals, $y - \mathbf{x}\widehat{\mathbf{b}}^{\text{WLS}}$.

In practice one can check the convergence of the sequence $\sum_{j=0}^J \mathbf{b}_{\text{p}}^{(j)}$. Upon convergence, OCJ and OCJ.B will be negligibly different.

6 Guaranteed conservative variance estimator

From first principles, the exact variance of a linear estimators of the form given in equation (1) is,

$$\begin{aligned}
\mathbf{V} \left(\widehat{\delta}_c \right) &= \mathbf{E} [y'\mathbf{R}\mathbf{W}'c c'\mathbf{W}\mathbf{R}y] - \mathbf{E} [y'\mathbf{R}\mathbf{W}'c] \mathbf{E} [c'\mathbf{W}\mathbf{R}y] \\
&= y' \left\{ \mathbf{E} [\mathbf{R}\mathbf{W}'c c'\mathbf{W}\mathbf{R}] - \mathbf{E} [\mathbf{R}\mathbf{W}'c] \mathbf{E} [c'\mathbf{W}\mathbf{R}] \right\} y \\
&= y'\mathbf{g}y
\end{aligned}$$

where \mathbf{g} is defined as the matrix inside the curly brackets in the line above. The matrix, \mathbf{g} , may be computed exactly for small studies or to arbitrary precision by repeating the randomization until desired precision is achieved.

As before, some terms in this quadratic are not observed, and a suitable bound can be found generalizing one of the methods from Paper 1 of 4 (Neyman, Aronow-Samii or Middleton). Algorithm 4.7 from Paper 1 of 4 will be slightly modified and applied to the simulations below.

7 Data from Paluck and Green (2009)

Paluck and Green (2009) pair-randomized 14 villages (7 pairs) in post-genocide Rwanda to receive one of two possible radio programs. Half of the villages were exposed to a program aimed at “discouraging blind obedience and reliance on direction from authorities and promoting independent thought and collective action in problem solving” (treatment). The other half were exposed to a radio program about health (control).

The analysis of the properties of the OC and GS variance estimators are compared in the spirit of (permutation/rerandomization-based) simulation. The analysis is simulation in the sense that it makes untestable assumptions about the outcomes that would have been observed under alternative assignments (e.g., the values of missing potential outcomes). The appeal of using real data, rather than generating data from a contrived DGP, however, is that it may provide a more realistic data set and the analysis can reflect the design actually used.

7.1 Randomization method

The original paired-cluster design was simulated and the estimator is OLS as in the original paper. With 7 pairs this results in $2^7 = 128$ possible randomizations.

7.2 Data

Analysis used the same 497 cases analyzed in model 1 of Table 4 in Paluck and Green (2009). Number of units by village pair and assignment arm are given in Table 1. The average cluster size was 35.5. The minimum and maximum cluster sizes were 20 and 43, respectively.

Five cases had missing values for *age*, which were mean-imputed for this analysis so that the same units to be used in both specifications. By contrast, the original analysis dropped these five cases when *age* was included in the OLS.

Table 1: Number of units by village pair and original assignment

	Village Pair						
	1	2	3	4	5	6	7
Control	37	39	39	39	33	37	43
Treatment	33	37	36	37	38	20	29

7.3 Outcome measure

The outcome is a post-treatment measure of social distance. The original effect estimates using OLS were presented in Table 4 of Paluck and Green (2009). The social distance measure was an index created by combining four survey items, each of which could take on integer values of 1,2,3 or 4. The resulting measure, an average of the items, also ranged from 1 to 4, but could take on 13 possible values.

Additionally the outcome, y , was location shifted by subtracting the mid-point on the scale, $(\max(y) - \min(y))/2$, such that the scale ranged from -1.5 to 1.5 rather than 1 to 4. One motivation for this is that non-invariant estimators such as *OCJ* and *GCJ* can, themselves, have high variance when the mean of y is shifted away from zero.

Another motivation is that when scale is bounded, and thus $\max(y)$ and $\min(y)$ are known, shifting by the scale midpoint introduces an invariance property. For example, subtracting off the scale midpoint results in the same observed values whether a researcher records the original outcomes on a scale of integers $\{0, 1, 2, 3\}$ or $\{1, 2, 3, 4\}$.

Another option might be to subtract the observed sample mean from each outcome. However, this leads to a shift that may depend on the randomization, when there are treatment effects, and so the practice can introduce biases.

7.4 Variance estimators compared

The Generalized Neyman bounding matrix, given in section 4.1 of Paper 1 of 4 (Middleton, 2020), \mathbf{d}^N , is not applicable in this case because of unobservable cells in main block diagonal of the design matrix, \mathbf{d} . Therefore simulations use the bounding matrix, \mathbf{d}^M , which is based on Algorithm 4.7 in Paper 1 of 4. It provides a tighter bound than the Aronow-Samii bounding matrix because clusters were pair-randomized, a result that is similar to the one given for pair-randomization of units in the example provided in section 4.4 of that paper.

The variance estimators are based upon the Generalized Sandwich (GS) given in equation (6) in Section 4, the Obloženè Chlebižky-based, given in Section 5.5 ($J = 10$), and the Guaranteed Conservative, given in Section 6 ($J = 10$).

7.5 Covariate specifications

The OC and GS variance estimators were examined with and without covariates. The “no covariates” OLS included only intercepts, so that $\mathbf{x} = \mathbf{1}$. This is equivalently the difference-of-means. The second specification included the additional covariates in column 3 of Table 4 in the original: *displaced by violence*, *sex*, *age*, and *radio lisening habits*.

For both of these specifications, classic Cluster Robust (CR) variance estimators also included fixed-effects for cluster-pairs as in columns 2 and 3 of Table 4 of Paluck and Green (2009). The typical motivation for including the fixed effects for blocks/pairs is to obtain standard errors that reflect this design feature. These are unnecessary in the case of OC and GC variance estimators because the paired-cluster design is already accounted for in the design matrix, \mathbf{d} .

7.6 Simulation assumptions

The assumptions of NCM apply, for example, assuming that potential outcomes are fixed (non-stochastic) and depend only on assignment. Also the sharp null, $y_{0i} = y_{1i}$ for all i , is invoked in order to impute missing potential outcomes.¹ The simulation also proceeds as if

¹Simulating under the sharp null seems reasonable for two reasons. First, the estimated average effects (-.029 to -.041, depending on specification) are small relative to the range of the measure (1.0% to 1.4%, respectively) and also to its standard deviation (3.6% to 5.1%, respectively). Second, and more importantly, simulating under the sharp null (or constant treatment effects more generally) is useful because the variance and its (identified) bound will be equal. The implication is that any *downward* bias of the variance estimator that is attributable to borrowing the variance of the first-order approximation, and/or to using a “plug-in” estimator thereof, will be evident in simulation. By contrast, downward bias can be masked (i.e., offset) under heterogeneous effects because the *bound* will tend to be conservative in that case. So, simulating under the sharp null gives a “worst-case” assessment of anti-conservative bias.

the survey response is not stochastic and whether a unit participates in the survey is not dependent on treatment assignments.²

7.7 Simulation results

Tables 3 and 2 give simulation results for OLS with and without covariates, respectively, using the data originally reported in Paluck and Green (2009). Comparing rMSE of variance estimators:

Comparisons:

- Generalized Sandwich is less biased and has a smaller $SE[\widehat{V}]$ compared to the classical Cluster Robust variance estimators
- OC estimators have smaller $SE[\widehat{V}]$ compared to GS, while having the same expected value as GS by construction
- Only GCJ is unbiased for the variance of the OLS coefficient. It has slightly higher $SE[\widehat{V}]$ than the GS, though the rMSE is lower because GS is biased.

Table 2: Comparing Variance Estimators for OLS, No Covariates

	Classical Cluster Robust			Gen. Sandwich	Obloženě Chlebižky			Guaranteed Conservative		
	CR0	CR1	CR2	GS0	OC0.B	OCJ.B	OCJ	GC0.B	GCJ.B	GCJ
$E[\widehat{V}/V]$	0.429	0.462	0.818	0.842	0.841	0.842	0.842	0.858	0.860	1.000
$SE[\widehat{V}/V]$	0.097	0.104	0.167	0.144	0.123	0.123	0.123	0.128	0.129	0.132
$Bias[\widehat{V}/V]$	-0.571	-0.538	-0.182	-0.158	-0.159	-0.158	-0.158	-0.142	-0.140	0.000
$rMSE[\widehat{V}/V]$	0.579	0.548	0.247	0.214	0.201	0.200	0.200	0.191	0.190	0.132
$CV[\widehat{V}]$	0.225	0.225	0.204	0.171	0.146	0.147	0.146	0.149	0.150	0.132
rMSE ratio:										
vs. GS0	2.712	2.566	1.155	1.000						
vs. OCJ	2.892	2.737	1.232	1.066	1.005	1.000	1.000			
vs. GCJ	4.392	4.156	1.870	1.619	1.526	1.519	1.519	1.451	1.442	1.000
CV ratio:										
vs. GS0	1.317	1.317	1.191	1.000						
vs. OCJ	1.537	1.537	1.389	1.167	0.999	1.000	1.000			
vs. GCJ	1.707	1.707	1.543	1.296	1.110	1.111	1.111	1.132	1.134	1.000

²Simulating as if survey response is non-stochastic and not affected by treatment assignments also seems reasonable for two reasons. First, an assumption about the independence of assignment and response would also be required for consistency in the original analysis given by Paluck and Green (2009). Second, similar to the case of the sharp null, if survey response was independent of assignment but stochastic, then variance estimators would be conservativeness relative to the assumption of response being non-random. Again, with this assumption, simulation results give a “worst-case” assessment of anti-conservative bias.

Table 3: Comparing Variance Estimators for OLS, With Covariates

	Classic Cluster Robust			Gen. Sandwich	Obloženè Chlebičky			Guaranteed Conservative		
	CR0	CR1	CR2	GS0	OC0.B	OCJ.B	OCJ	GC0.B	GCJ.B	GCJ
$E[\widehat{V}/V]$	0.431	0.464	0.819	0.802	0.797	0.802	0.802	0.853	0.861	1.000
$SE[\widehat{V}/V]$	0.113	0.121	0.187	0.131	0.094	0.095	0.095	0.118	0.121	0.125
$Bias[\widehat{V}/V]$	-0.569	-0.536	-0.181	-0.198	-0.203	-0.198	-0.198	-0.147	-0.139	0.000
$rMSE[\widehat{V}/V]$	0.580	0.550	0.260	0.237	0.224	0.220	0.220	0.189	0.184	0.125
$CV[\widehat{V}]$	0.261	0.261	0.228	0.163	0.118	0.119	0.118	0.138	0.140	0.125
rMSE ratio:										
vs. GS0	2.447	2.318	1.095	1.000						
vs. OCJ	2.640	2.500	1.181	1.079	1.019	1.000	1.000			
vs. GCJ	4.638	4.393	2.075	1.895	1.790	1.758	1.757	1.506	1.472	1.000
CV ratio:										
vs. GS0	1.606	1.606	1.400	1.000						
vs. OCJ	2.207	2.207	1.923	1.374	0.998	1.003	1.000			
vs. GCJ	2.089	2.089	1.820	1.300	0.944	0.949	0.946	1.106	1.121	1.000

References

- Aronow, Peter M. and Cyrus Samii. 2012. Conservative variance estimation for sampling designs with zero pairwise inclusion probabilities. *Survey Methodology* **39**(1): 231-241.
- Aronow, Peter M. and Cyrus Samii. 2017. Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics* .
- Aronow, Peter M. and Joel A. Middleton. 2015. A class of unbiased estimators of average treatment effect in randomized experiments. *Journal of Causal Inference* **1**(1): 135-154.
- Holland, P.W. 1986. Statistics and Causal Inference, *Journal of the American Statistical Association*, vol. 81, no. 396: 945-968.
- Horvitz, D.G. and Thompson, D.J. 1952. A generalization of sampling without replacement from a finite universe. *J. Amer. Statist. Assoc.* **47**: 663-684.
- Lin, Winston. 2013. Agnostic Notes on Regression Adjustments to Experimental Data: Reexamining Freedman’s Critique. *Annals of Applied Statistics* **7**(1): 295-318
- Middleton, Joel A. 2018. A unified theory of regression adjustment for design-based inference. *arXiv:1803.06011*
- Middleton, Joel A. 2020. Unifying Design-based Inference: On bounding and estimating the variance of any linear estimator in any experimental design. *Working paper at www.joelmiddleton.com*.
- Middleton, Joel A. and Peter M. Aronow. 2015. Unbiased Estimation of the Average Treatment Effect in Cluster-Randomized Experiments. *Statistics, Politics and Policy* **1**:

Neyman, Jerzy Splawa, D. M. Dabrowska, and T. P. Speed. [1923.] 1990. On the application of probability theory to agricultural experiments: Essay on principles, section 9. *Statistical Science* **5**: 465-480.

Paluck, Elizabeth Levy, and Donald P. Green. 2009. Deference, Dissent, and Dispute Resolution: An Experimental Intervention Using Mass Media to Change Norms and Behavior in Rawanda. *American Political Science Review* **103**(4): 622-644.