Unifying Design-based Inference: On Bounding and Estimating the Variance of any Linear Estimator in any Experimental Design

WORKING PAPER 1 OF 4

Joel A. Middleton*

May 18, 2020

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

1 Introduction

This paper provides a framework for variance (bound) estimation in experimental analysis. Results are applicable to virtually any conceivable experimental design, any linear estimator (including but not limited to OLS, WLS, IPTW, AIPTW, doubly-robust, difference-in-means, Horvitz-Thompson, Hajek) and any identified variance bound. Variance bounds are defined in a new and useful way, opening the area to systematic study. A general estimator is proposed that subsumes Eicker-Huber-White (aka. "robust", "heteroskedastic consistent", "sandwich", "White", "Huber-White", "HC", etc.). In so doing the unified framework allows for systematic analysis of design-estimator-bound combinations in ways never before proposed.

Contributions:

- Unification: Provides a framework for the evaluation and comparative asymptotic analysis of a virtually unlimited set of design-estimator-bound triplets.
- Variance bounds: Opens the study of variance bounding to systematic study. Proposes a novel algorithm for variance bounding as a proof-of-concept, illustrating the utility of the framework and demonstrating the need for further study of an area that has been largely neglected to date. Contextualizes existing Neyman and Aronow-Samii bounds in light of this framework while providing elegant expressions and novel proofs.
- Variance bound estimation: A general variance bound estimator is proposed that can be applied to virtually any design, any linear estimator, and any identified bound. It subsumes the canonical Eicker-Huber-White (EHW, robust, sandwich, heteroskedastic consistent, HC, etc.) variance bound estimator, being algebraically equivalent to it for Bernoulli (independent assignment) designs with the OLS estimator. It likewise subsumes so-called "cluster robust" variance bounds, being algebraically equivalent for Bernoulli (independent) assignment of clusters, the OLS estimator and the Generalized Neyman bound (defined below). However, the proposed variance bound estimator is general.

1.1 The Neyman Causal Model

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 $\{y_{1i}, y_{2i}, ..., y_{ki}\}$, 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) distribution.

The only random element in the NCM is the treatment assignment indicators $\{R_{1i}, R_{2i}, ..., R_{ki}\}$, and they determine which potential outcome will be observed by the researcher. Since a unit can only be assigned to one arm of the experiment, only one of the indicators will realize a value of one, and the rest will be zero, such that $R_{1i} + R_{2i} + ... + R_{ki} = 1$ for all i.

A standard representation of the *observed* outcome for the i^{th} unit under the NCM would be,

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

which is itself random, due to the assignment indicators. For each unit, the observed data can then be represented as $\{Y_i^{obs}, R_{1i}, R_{2i}, ..., R_{ki}, x_i\}_{\forall i}$, where x_i is an additional vector

of k covariates. Like the potential outcomes, x_i is nonrandom, but unlike the potential outcomes the same value is observed irrespective of the assignment. In practice this might be ensured by collecting this data before the assignment.

Ideally, we would like to know, for a given individual, i, the difference between responses under various arms, called a *treatment effect*. It is clear from the definition of Y_i^{obs}, however, that individual treatment effects are not observable since only one of the potential outcomes can be observed for an individual, a problem known as *fundamental problem of causal inference* (Holland, 1986). As a result, researchers often try to estimate *averages* of across the units in study.

Example (*Treatment/Control Experiment*) : In an experiment with a control group (arm 0) and a treatment group (arm 1) the individual-level treatment effect, $y_{1i} - y_{0i}$, but this is not identified, so a researcher might estimate have the average treatment effect $n^{-1} \sum_i (y_{1i} - y_{0i})$.

Example (2×2 *Factorial Experiment*) : Consider a 2×2 factorial design with treatments A and B. Units in arm 1 are controls (no treatments), units in arm 2 are given treatment A only, units in arm 3 are given B only, and units in arm 4 are given both A and B. Similar to the treatment/control example, one could contrast the mean of an arm with a single treatment against the control mean, e.g., the average effect of A compared to no treatments, $n^{-1} \sum_i (y_{2i} - y_{1i})$. Another quantity of interest might be an *average marginal causal effect* (AMCE), e.g., the effect of A marginalizing over the levels of B, $n^{-1} \sum_i \frac{1}{2} (y_{2i} - y_{1i} + y_{4i} - y_{3i})$. Another example might be an omnibus test based on the contrast $n^{-1} \sum_i (y_{2i}/3 + y_{3i}/3 + y_{4i}/3 - y_{1i})$.

Target quantities such as *local average treatment effects* or *conditional average treatment effects* might also be considered in this framework, but the primary focus of this paper is *variance* estimation for linear estimators for virtually any design.

Suffice to say that developing variance estimators before considering point estimation is appealing, if somewhat counter-intuitive, for two reasons. On the one hand, asymptotic analysis for point estimators can be made easier by having first established general variance expressions (for all linear estimators and virtually any design). On the other hand, a general framework for variance (bound) estimation can be developed even while a particular estimation target has yet to be defined, and even if an "estimator" does not estimate anything of interest, it's variance can still be studied.

1.2 Modified Notation

To simplify notation, let y_1, y_2,...,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, stack these vectors to create

$$y := (y_1' \;\; y_2' \;\; \dots \;\; y_k')',$$

which is a column vector and has length kn containing all k potential outcomes for all n units.

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

defined as,

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

which, for example, allows us to express a k-length vector of the means of each arm as $(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}'y$. Next, define c as the *contrast vector*, of length k, such that $c'(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}'y$ gives contrasts between potential outcome means for the various arms.

Remark 1. *Note that $(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}'y$ has the form of an ordinary least squares (OLS) coefficient.*

Example (*Treatment/Control Experiment, continued*) : With two arms, control (arm 1) and treatment (arm 2), define $c = (-1 \;\; 1)'$. Then $c'(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}'y = n^{-1}\sum_i (y_{2i} - y_{1i})$.

Example (2×2 *Factorial Experiment, continued*) : If one were interested in the average difference in responses between the first two arms, then $c = (-1 \;\; 1 \;\; 0 \;\; 0)'$ and we have $c'(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}'y = n^{-1}\sum_i (y_{2i} - y_{1i})$. Likewise, if one chooses $c = (-\frac{1}{2} \;\; \frac{1}{2} \;\; -\frac{1}{2} \;\; \frac{1}{2})'$ then $c'(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}'y = n^{-1}\sum_i \frac{1}{2}(y_{2i} - y_{1i} + y_{4i} - y_{3i})$. 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 is not be strictly necessary.

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 $\mathbf{R}_2$, $\mathbf{R}_3$, ..., $\mathbf{R}_k$ analogously. Arange 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 $\boldsymbol{\pi} := \mathrm{E}[\mathbf{R}]$, 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 is convenient to define the $kn \times (k+l)$ matrix,

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

which augments the intercept vector, $\mathbb{1}$, with covariates.

Remark 2. *For some cases, such adjusting for covariates separately by arm, it might be useful to define $\mathbb{x}$ 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.*

2 Linear estimators

This paper focuses on the variance, bounding and variance bound estimation of the class of estimators that are linear in the observed outcome, y. This class includes everything from the difference-of-means, to the Horvitz-Thomposon estimator, to regression.

Note, however, that beyond presenting a general approach to variance bound estimation for the class of linear estimators, point estimation itself will be the focus of the third and fourth papers in the series. Questions such as consistency will be and causal identification will be considered then. For now, suffice it to be said that an estimator need not be consistent for any quantity of interest at all (causal or otherwise) in order to derive variance expressions for it.

2.1 Definition

Definition 2.1 (Linear Estimators)**.** *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 some common covariate adjusted estimators.

In common cases (e.g., OLS, WLS, Hajek), $\mathbf{W}$ will be a random matrix and is therefore represented as upper case (for random) and bold (for matrix). In others, it might be a constant matrix (e.g. Horvitz-Thompson). Still others it might be a random vector (e.g., generalized regression).

Also, for convenience, define $\mathbf{w}$ to be $\mathbf{W}$ evaluated at $\mathbf{R} = \boldsymbol{\pi}$, i.e.,

$$\mathbf{w} := \{\mathbf{W}|_{\mathbf{R}=\boldsymbol{\pi}}\}. \tag{2}$$

Definition 2.2 (Horvitz-Thompson estimator)**.** *The Horvitz-Thompson estimator written as in Definition 2.1 with,*

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

noting that $\mathbf{W}^{\text{HT}} = \mathbf{w}^{\text{HT}}$ since $\mathbf{W}^{\text{HT}}$ is nonrandom.

Definition 2.3 (Contrast-of-means)**.** *Contrast-of-means (e.g., difference-of-means) can be written as in Definition 2.1 with,*

$$\mathbf{W} = \mathbf{W}^{\text{CM}} := \left(\mathbb{1}'\mathbf{R}\mathbb{1}\right)^{-1}\mathbb{1}'.$$

Definition 2.4 (Hajek estimator)**.** *The Hajek estimator can be written as Definition (2.1) with,*

$$\mathbf{W} = \mathbf{W}^{\text{HJ}} := \left(\mathbb{1}'\boldsymbol{\pi}^{-1}\mathbf{R}\mathbb{1}\right)^{-1}\mathbb{1}'\boldsymbol{\pi}^{-1}.$$

Definition 2.5 (WLS estimators)**.** *WLS estimators can be written as in Definition 2.1 with,*

$$\mathbf{W} = \mathbf{W}^{\text{WLS}} := \left(\mathbb{x}'\mathbf{m}\mathbf{R}\mathbb{x}\right)^{-1}\mathbb{x}'\mathbf{m}.$$

Remark 3. *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,* $\mathbb{x} = \mathbb{1}$ *(there are no covariates), WLS is equivalent to the difference-of-means. If* $\mathbf{m} = \boldsymbol{\pi}^{-1}$ *and* $\mathbb{x} = \mathbb{1}$*, 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. Likewise, deleting from* $\mathbb{x}$ *the columns associated with matrix of covariates,* $\mathbf{x}$*, then* $\mathbb{x} = \mathbb{1}$*, and* $\mathbf{W}^{\text{OLS}} = \mathbf{W}^{\text{CM}}$*, underscoring the observation that the difference of means is equivalent to OLS with no covariates (cf. Freedman 2008, Theorem 1).*

2.2 First-order Taylor approximation

In this section, a general approach to obtaining asymptotically valid variance expressions for linear estimators is given using a first-order approximation of a Tyalor series. The method is often used when an exact, closed-form variance expression is not tractable, as may be the case with any number of linear estimators. Examination of the W vectors defined above shows that, with the exception of Horvitz-Thompson, the estimators all had random denominators (i.e., inverted random matrices), making closed form variance expressions difficult.

The original estimator and its Taylor approximation are asymptotically equivalent (cite Pashley). As such, the original estimator "borrows" the closed-form variance expression given for the Taylor approximation, again justified given the asymptotic equivalence.

Lemma 2.6 (First-order Taylor approximation for linear estimators)**.** *First, assume a linear estimator as defined in Definition 2.1. Then, let* $\left\{ \left. . \right|_{\mathbf{R}=\boldsymbol{\pi}} \right\}$ *represent a function that evaluates the argument to the left of the vertical line at* $\mathbf{R} = \boldsymbol{\pi}$*. Similarly, let* $\left\{ \left. . \right|_{\mathbf{R}=\boldsymbol{\pi}} (\mathbf{R} - \boldsymbol{\pi}) \right\}$ *evaluate its argument at* $\mathbf{R} = \boldsymbol{\pi}$ *and then multiply by* $(\mathbf{R} - \boldsymbol{\pi})$*. Then from Taylor's theorem and the product rule, we have the first-order Taylor approximation,* $\widehat{\delta} \approx \widehat{\delta}^{\text{T}}$*, with*

$$\begin{aligned}
\widehat{\delta}_c^{\text{T}} :=& \left\{ \left. \widehat{\delta}_c \right|_{\mathbf{R}=\boldsymbol{\pi}} \right\} + \left\{ \left. \frac{d}{d\mathbf{R}} c'\mathbf{W} \right|_{\mathbf{R}=\boldsymbol{\pi}} (\mathbf{R} - \boldsymbol{\pi}) \right\} \left\{ \left. \mathbf{R} \right|_{\mathbf{R}=\boldsymbol{\pi}} \right\} y \\
&+ \left\{ \left. c'\mathbf{W} \right|_{\mathbf{R}=\boldsymbol{\pi}} \right\} \left\{ \left. \frac{d}{d\mathbf{R}} \mathbf{R} \right|_{\mathbf{R}=\boldsymbol{\pi}} (\mathbf{R} - \boldsymbol{\pi}) \right\} y \qquad (3) \\
=\ & a_c + \left\{ \left. \frac{d}{d\mathbf{R}} c'\mathbf{W} \right|_{\mathbf{R}=\boldsymbol{\pi}} \mathbf{R} \right\} \boldsymbol{\pi} y + c'\mathbf{w}\mathbf{R}y
\end{aligned}$$

where a_c *is a constant.*

Remark 4. *An expression for a_c is not given, since it does not contribute to the variance of the approximation. Recall that the purpose of deriving a first-order Taylor approximation, $\widehat{\delta}_c^{\mathrm{T}}$, is to identify a closed-form variance expression that might then be "borrowed" by the original linear estimator given in Definition 2.1, i.e., so that we might use the fact that $\mathrm{V}\left(\widehat{\delta}_c\right) \approx \mathrm{V}\left(\widehat{\delta}_c^{\mathrm{T}}\right)$ to justify using the variance estimator, $\widehat{\mathrm{V}}\left(\widehat{\delta}_c\right) := \widehat{\mathrm{V}}\left(\widehat{\delta}_c^{\mathrm{T}}\right)$.*

Theorem 2.7. *For a constant, a_c, and vector of constants, z_c, first-order Taylor approximations for linear estimators may be written as,*

$$\widehat{\delta}_c^{\mathrm{T}} = a_c + n1_k' \mathbf{w}^{\mathrm{HT}} \mathbf{R} z_c.$$

Hence, a first-order approximation of a Tyalor series using Taylor's theorem variance approximations will be expressed as the variance of a Horvitz-Thompson estimator of the ATE of z_c with contrast vector $n1_k$.

Table 1: Examples of linear estimators. $\mathbf{W}$ is as defined in Definition 2.1, z_c is as defined in Theorem 2.7.

Estimator	$\mathbf{W}$	z_c
Horvitz-Thompson	$(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}'\boldsymbol{\pi}^{-1}$	$\mathrm{diag}(y)\mathbb{1}(\mathbb{1}'\mathbb{1})^{-1}c$
Contr.-of-means	$(\mathbb{1}'\mathbf{R}\mathbb{1})^{-1}\mathbb{1}'$	$\boldsymbol{\pi}\mathrm{diag}(y-\mathbb{1}(\mathbb{1}'\boldsymbol{\pi}\mathbb{1})^{-1}\mathbb{1}'\boldsymbol{\pi}y)\mathbb{1}(\mathbb{1}'\boldsymbol{\pi}\mathbb{1})^{-1}c$
Hajek	$\left(\mathbb{1}'\boldsymbol{\pi}^{-1}\mathbf{R}\mathbb{1}\right)^{-1}\mathbb{1}'\boldsymbol{\pi}^{-1}$	$\mathrm{diag}(y-\mathbb{1}(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}'y)\mathbb{1}(\mathbb{1}'\mathbb{1})^{-1}c$
OLS	$(\mathbb{x}'\mathbf{R}\mathbb{x})^{-1}\mathbb{x}'$	$\boldsymbol{\pi}\mathrm{diag}\left(y-\mathbb{x}b^{\mathrm{OLS}}\right)\mathbb{x}(\mathbb{x}'\boldsymbol{\pi}\mathbb{x})^{-1}c$
WLS	$(\mathbb{x}'\mathbf{m}\mathbf{R}\mathbb{x})^{-1}\mathbb{x}'\mathbf{m}$	$\boldsymbol{\pi}\mathrm{diag}\left(y-\mathbb{x}b^{\mathrm{WLS}}\right)\mathbf{m}\mathbb{x}(\mathbb{x}'\mathbf{m}\boldsymbol{\pi}\mathbb{x})^{-1}c$
GR ($b=b^{\mathrm{WLS}}$)	$\mathbf{w}^{\mathrm{HT}}\left(\mathbf{i}_{kn}-(\mathbf{R}-\boldsymbol{\pi})\mathbf{W}^{\mathrm{WLS}}\right)$	$\mathrm{diag}\left(y-\mathbb{x}b^{\mathrm{WLS}}\right)\mathbb{1}(\mathbb{1}'\mathbb{1})^{-1}c$
IV	$\left(\widehat{\tilde{\mathbb{x}}}'\mathbf{R}\widehat{\tilde{\mathbb{x}}}\right)^{-1}\widehat{\tilde{\mathbb{x}}}'$ with: $\widehat{\tilde{\mathbb{x}}} := \mathbb{z}(\mathbb{z}'\mathbf{R}\mathbb{z})^{-1}\mathbb{z}'\mathbf{R}\mathbb{x}$	$\boldsymbol{\pi}\mathrm{diag}\left(y-\mathbb{x}b^{\mathrm{IV}}\right)\tilde{\mathbb{x}}(\tilde{\mathbb{x}}'\boldsymbol{\pi}\tilde{\mathbb{x}})^{-1}c$ with: $\tilde{\mathbb{x}} := \mathbb{z}(\mathbb{z}'\boldsymbol{\pi}\mathbb{z})^{-1}\mathbb{z}'\boldsymbol{\pi}\mathbb{x}$, $b^{\mathrm{IV}} := (\tilde{\mathbb{x}}'\boldsymbol{\pi}\tilde{\mathbb{x}})^{-1}\tilde{\mathbb{x}}'\boldsymbol{\pi}y$

Example (*Weighted least squares*) : Weighted least squares gives as special cases OLS ($\mathbf{m}=\mathbf{i}_{kn}$), contrast-of-means (e.g., difference of means, with $\mathbf{m}=\mathbf{i}_{kn}$ and $\mathbb{x}=\mathbb{1}$) and the Hajek estimator ($\mathbf{m}=\boldsymbol{\pi}^{-1}$ and $\mathbb{x}=\mathbb{1}$). To derive its Taylor approximation, first let $\mathbf{w}^{\mathrm{WLS}} = \mathbf{W}^{\mathrm{WLS}}|_{\mathbf{R}=\boldsymbol{\pi}} = (\mathbb{x}'\mathbf{m}\boldsymbol{\pi}\mathbb{x})^{-1}\mathbb{x}'\mathbf{m}$, and define outcome vector, z^{WLS}, as,

$$z^{\mathrm{WLS}} = \boldsymbol{\pi}\mathrm{diag}\left(y-\mathbb{x}b^{\mathrm{WLS}}\right)\mathbf{m}\mathbf{w}^{\mathrm{WLS}\prime}c \tag{4}$$

where the vector $b^{\mathrm{WLS}} := \mathbf{w}^{\mathrm{WLS}}\mathbf{m}y$ is, loosely speaking, the WLS coefficient "estimand". Then, using equation (3) we have,

$$\begin{aligned}
\widehat{\delta_c}^{\mathrm{WLS}} &\approx a_c^{\mathrm{WLS}} - c'\mathbf{w}^{\mathrm{WLS}}\mathbf{R}\mathbb{x}\mathbf{w}^{\mathrm{WLS}}y + c'\mathbf{w}^{\mathrm{WLS}}\mathbf{R}y \\
&= a_c^{\mathrm{WLS}} + c'\mathbf{w}^{\mathrm{WLS}}\mathbf{R}\left(y-\mathbb{x}b^{\mathrm{WLS}}\right) \\
&= a_c^{\mathrm{WLS}} + 1_{kn}'\mathbf{R}(\boldsymbol{\pi}^{-1}\boldsymbol{\pi})\mathrm{diag}(y-\mathbb{x}b^{\mathrm{WLS}})\mathbf{w}^{\mathrm{WLS}\prime}c \\
&= a_c^{\mathrm{WLS}} + n1_k'\mathbf{w}^{\mathrm{HT}}\mathbf{R}z_c^{\mathrm{WLS}},
\end{aligned}$$

Therefore the first-order Taylor approximation is equivalent to a constant, a_c^{WLS} (the exact form of which is unimportant) plus a Horvitz-Thompson estimator with outcome vector, z^{WLS}, and contrast vector, $n1_k$.

Remark 5. *Expressions given in the examples illustrate that first order Taylor approximations for linear estimators are Horvitz-Thompson estimators. This highlights the importance of studying Horvitz-Thompson variance in order to develop asymptotic variance expressions for linear estimators in general.*

Remark 6. *Taylor approximations of this sort include terms such as b^{WLS} that are not known. The next section will show that the plug-in principle provides a basis for asymptotically valid variance expressions nonetheless.*

3 Variance

Now that the importance of Horvitz-Thompson estimators for asymptotic variance expressions for the entire class of linear estimators (which includes, for example, OLS, WLS, Hajek, and difference-of-means) has been established, this section will give the variance of HT estimators and first-order approximates of linear estimators.

Throughout, we will make use of the $kn \times kn$ "design matrix",

$$\mathbf{d} := \mathrm{V}\left(1'_{kn}\boldsymbol{\pi}^{-1}\mathbf{R}\right), \tag{5}$$

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

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

$$\mathrm{V}\left(\widehat{\delta}_c^{\text{T}}\right) = z_c'\mathbf{d}z_c, \tag{6}$$

with examples of z_c given in Table 1.

Remark 7. *Equation (6) is an exact variance expression, however, they are not identified. The next subsection introduces the necessary concept of variance bounding.*

Remark 8. *The design matrix, $\mathbf{d}$, will provide useful device in the study the best designs for outcomes with different characteristics as the next example will show.*

Example (*Comparing complete randomization and paired randomization*) : Consider a treatment/control (two-arm) experiment that is pair-randomized. A pair-randomized design is a special case of a block-randomized (i.e., stratified) design where blocks have size 2. In each pair/block, one unit is assigned to treatment and the other in control with equal (.5) probability. When $n = 4$ (and assuming w.l.o.g. that the data are sorted by pair), the design matrix is

$$\mathbf{d}^{pr} = \begin{bmatrix} 1 & -1 & & & -1 & 1 & & \\ -1 & 1 & & & 1 & -1 & & \\ & & 1 & -1 & & & -1 & 1 \\ & & -1 & 1 & & & 1 & -1 \\ -1 & 1 & & & 1 & -1 & & \\ 1 & -1 & & & -1 & 1 & & \\ & & -1 & 1 & & & 1 & -1 \\ & & 1 & -1 & & & -1 & 1 \end{bmatrix},$$

noting that empty cells represent 0. The design matrix for complete randomization (where 2 of 4 are randomly assigned to treatment) is

$$\mathbf{d}^{cr} = \begin{bmatrix} 1 & -1/3 & -1/3 & -1/3 & -1 & 1/3 & 1/3 & 1/3 \\ -1/3 & 1 & -1/3 & -1/3 & 1/3 & -1 & 1/3 & 1/3 \\ -1/3 & -1/3 & 1 & -1/3 & 1/3 & 1/3 & -1 & 1/3 \\ -1/3 & -1/3 & -1/3 & 1 & 1/3 & 1/3 & 1/3 & -1 \\ -1 & 1/3 & 1/3 & 1/3 & 1 & -1/3 & -1/3 & -1/3 \\ 1/3 & -1 & 1/3 & 1/3 & -1/3 & 1 & -1/3 & -1/3 \\ 1/3 & 1/3 & -1 & 1/3 & -1/3 & -1/3 & 1 & -1/3 \\ 1/3 & 1/3 & 1/3 & -1 & -1/3 & -1/3 & -1/3 & 1 \end{bmatrix}.$$

Eigendecomposition of $\mathbf{d}^{cr} - \mathbf{d}^{pr}$ gives eigenvalues $2.67, 0, 0, 0, 0, 0, -1.33$, and -1.33 corresponding eigenvectors in Table 2. The eigenvectors associated with nonzero eigenvalues provide insight into the subspace in $\mathbb{R}^{2n}$ where one design may be preferable to another, for example, when the estimator is difference-in-means (which is equivalent to both Horvitz-Thopson and Hajek for these designs).

Table 2: Eigenvectors of $\mathbf{d}^{cr} - \mathbf{d}^{pr}$

e1	e2	e3	e4	e5	e6	e7	e8
-0.354	0.791	0.000	0.000	0.000	0.000	0.500	0.000
-0.354	0.158	-0.573	-0.178	-0.250	-0.421	-0.500	0.000
0.354	0.158	-0.180	-0.450	0.585	-0.149	0.000	0.500
0.354	0.158	0.319	-0.600	-0.260	-0.264	0.000	-0.500
0.354	0.474	0.282	0.524	0.100	-0.190	-0.500	0.000
0.354	-0.158	-0.291	0.346	-0.150	-0.611	0.500	0.000
-0.354	-0.158	0.102	0.073	0.685	-0.339	-0.000	-0.500
-0.354	-0.158	0.602	-0.077	-0.161	-0.454	-0.000	0.500

In this example, assuming the contrast matrix is $c = (-1, 1)'$, and examining the first eigenvector with eigenvalue 2.67, one can conclude that if the outcomes for the four units given in Table 3, then the difference-of-means would be much less precise under the completely randomized design. So, the eigenvector in a sense represents a "best-case" (normed) potential outcome vector for paired randomization. Inspection of the outcomes themselves confirms the intuition that pair randomization is better than complete randomization when units are homogenous within pairs.

Table 3: Pair randomization better than complete randomization

unit id	pair id	y_0	y_1
1	1	.3536	.3536
2	1	.3536	.3536
3	2	-.3536	-.3536
4	2	-.3536	-.3536

Next, considering the two eigenvectors associated with the eigenvalue -1.33, we see the implied potential outcomes in Table 4 give potential outcomes for which complete randomization is preferable. Note that either of the two sets is a "worst-case" scenario for paired randomization, as is any set of potential outcomes that can be generated by linear combinations of the two eigenvectors. Inspection of these outcomes is consistent with the observation

that complete randomization can be better than paired randomization when paired units are maximally heterogeneous.

Table 4: Complete randomization better than pair randomization

unit id	pair id	y_0	y_1	y_0	y_1
1	1	-.5	-.5	0	0
2	1	.5	.5	0	0
3	2	0	0	-.5	-.5
4	2	0	0	.5	.5

Remark 9. *The observation that pair randomization can hurt precision when units are not homogeneous within pairs is not new. However, this approach to comparing designs is perfectly general and can be applied to any design. Moreover, the example illustrates a relatively effortless method of identifying key insights about arbitrary designs.*

4 Variance bounds

In spite of an exact expression for first-order Taylor approximations in Equation (9), the quantity is unidentified because not all terms in the quadratic can be observed. Even if the elements of $\mathbf{R}z_c$ was observed directly, some pairs of potential outcomes can never be jointly observed. For example, for a given unit, only one of two (or more) potential outcomes can be observed, a problem is referred to as the "fundamental problem of causal inference" (Holland, 1986). Other design features, such as clustering or pair randomization, can also render various combinations of potential outcomes unobservable.

Starting with Neyman (1923) one proposed solution to unidentified variance has been to estimate a *variance bound*, i.e., a quantity that is provably greater than the variance, but which is identified. It should be understood that while the term *variance estimation* is often used as a shorthand in the literature, it is not, in general, an accurate phrase. *Variance bound estimation* is a more precise so it will be used here.

Definition 4.1 (Variance bound matrix)**.** *Let $\tilde{\mathbf{d}}$ be an arbitrary $kn \times kn$ matrix and let be z an arbitrary vector with length kn. Then $\tilde{\mathbf{d}}$ is a* variance bound matrix *(or* bounding matrix*) for $\mathbf{d}$ if, for all $z \in \mathbb{R}^{kn}$, $z'\mathbf{d}z \leq z'\tilde{\mathbf{d}}z$.*

Lemma 4.2. *$\tilde{\mathbf{d}}$ is a bounding matrix $\mathbf{d}$ if and only if matrix $\tilde{\mathbf{d}} - \mathbf{d}$ is positive semi-definite.*

Proof. By the definition of a bound, $z'\tilde{\mathbf{d}}z - z'\mathbf{d}z \geq 0$ for all $z \in \mathbb{R}^{kn}$. This implies that $z'(\tilde{\mathbf{d}} - \mathbf{d})z \geq 0$, i.e., that $\tilde{\mathbf{d}} - \mathbf{d}$ is positive semi-definite. □

Definition 4.3 (Identified variance bound)**.** *Let $\tilde{\mathbf{d}}$ be bounding matrix for $\mathbf{d}$. It gives an* identified *variance bound if*

$$\mathrm{I}(\mathbf{d} = -1) \circ \mathrm{I}(\tilde{\mathbf{d}} = 0) = \mathrm{I}(\mathbf{d} = -1)$$

where $\circ$ is element-wise multiplication, $\mathrm{I}(\mathbf{d} = -1)$ is an indicator function returning an $kn \times kn$ matrix of ones and zeros indicating whether each element of $\mathbf{d}$ is equal to -1 (an indication that the associated term in the variance quadratic is impossible to observe), and $\mathrm{I}(\tilde{\mathbf{d}} = 0)$ is, similarly, an indicator function returning an $kn \times kn$ matrix of ones and zeros indicating the location of zeros in $\tilde{\mathbf{d}}$.

4.1 Generalizing Neyman's variance bound

This section proposes a generalization of Neyman's (1923) variance bound. Let matrix $\mathbf{d}$ be partitioned into k^2 partitions of size $n \times n$. Then for $i, j \in \{1, 2, ..., k\}$, let the $\mathbf{d}_{ij}$ be the $(i, j)^{th}$ partition, having dimension $n \times n$. Also, let c_i be the i^{th} element of the length-k contrast vector, c. Then the following bounding method produces an identified bound for experiments in which $c_i \mathrm{I}(\mathbf{d}_{ii} = -1) = 0_{n \times n}$ for all i (no negative ones in block diagonals associated with nonzero contrast values) and $\mathbf{d}_{ij} = \mathbf{d}_{ji}$ for all $i, j \in \{1, 2, ..., k\}$ (off-diagonal blocks are symmetric). This covers a variety of common designs.

Definition 4.4 (Generalized Neyman variance bound)**.** *The "Generalized Neyman bound" is the is the bound corresponding to the block-diagonal bounding matrix,* $\tilde{\mathbf{d}}^{\mathrm{N}}$*, with block* (i, i) *given by,*

$$\tilde{\mathbf{d}}^{\mathrm{N}}_{ii} := \sum_{j=1}^{k} \frac{c_j}{c_i} \mathbf{d}_{ij}$$

where c_i *and* c_j *are, respectively, elements* i *and* j *from from the contrast vector,* c*.*

4.2 A novel proof of the Aronow-Samii bound

Consider an identified bound proposed by Aronow and Samii (2017) that has the a unusual virtue of being perfectly general, i.e., applicable to arbitrary (identified) designs.

Definition 4.5 (Aronow-Samii variance bound)**.** *The "Aronow-Samii variance bound" is the bound corresponding to the bounding matrix,*

$$\tilde{\mathbf{d}}^{\mathrm{AS}} := \mathbf{d} + \mathrm{I}\left(\mathbf{d} = -1\right) + \mathrm{diag}(\mathrm{I}\left(\mathbf{d} = -1\right) 1_{kn})$$

and diag(.) *creates a diagonal matrix from a vector.*

Theorem 4.6. *The Aronow-Samii variance bound,* $n^{-2} y' \tilde{\mathbf{d}}^{\mathrm{AS}} y$*, is an identified bound for* $n^{-2} y' \mathbf{d} y$*.*

Proof. By definition of $\tilde{\mathbf{d}}^{\mathrm{AS}}$,

$$\tilde{\mathbf{d}}^{\mathrm{AS}} - \mathbf{d} = \mathrm{I}\left(\mathbf{d} = -1\right) + \mathrm{diag}\left(\mathrm{I}\left(\mathbf{d} = -1\right) 1_{kn}\right).$$

Note that by construction $(\tilde{\mathbf{d}}^{\mathrm{AS}} - \mathbf{d})$ has diagonal elements set equal to the sum of the off-diagonal elements in its row (which by construction are either 0 or 1). The Gershgorin circle theorem implies that a real matrix is positive semi-definite if, for all i, the i^{th} diagonal element is greater or equal to the sum of the absolute values of the other elements in the i^{th} row. So, by the Gershgorin circle theorem $\tilde{\mathbf{d}}^{\mathrm{AS}} - \mathbf{d}$ is positive semidefinite. Therefore, by Lemma (4.2), $n^{-2} y' \tilde{\mathbf{d}}^{\mathrm{AS}} y$ is a variance bound. Moreover, as long as the design is an identified design (i.e., $0 < \pi_{1i} < 1$ for all i), it is an identified bound because $\mathrm{I}\left(\mathbf{d} = -1\right)$ ensures that the elements of $\mathbf{d}$ equal to -1 correspond to 0's in $\tilde{\mathbf{d}}^{\mathrm{AS}}$. □

Aronow and Samii (2017) derive their bound using Young's inequality. The above-theorem and proof using the Gershgorin circle theorem tie their insight to the current framework.

4.3 A proposed algorithm for finding a variance bound for any design

The following is an algorithm which that can obtain an identified variance bound. Like the AS bound it has the virtue of being applicable to virtually any design. The algorithm is a proof of concept, demonstrating the utility of the notation scheme which allows for the application of matrix theory for the creation of alternative bounds. The subject of comparing bounds will be considered further in Section 4.4.

Algorithm 4.7.

1. Initialize $kn \times kn$ matrix $\mathbf{t}$. Examples could be $\mathrm{I}(\mathbf{d} = -1)$ or, if the conditions for the Neyman bound not be applicable, start with $\tilde{\mathbf{d}}^{\mathrm{N}} - \mathbf{d}$ which may approximate a bound
2. Obtain the eigen decomposition of matrix $\mathbf{t}$. If all eigenvalues are non-negative (within tolerance), goto Step 6, otherwise continue
3. Update $\mathbf{t} = \mathbf{v}(\mathbf{e} \circ \mathrm{I}(\mathbf{e} > 0))\mathbf{v}'$ where $\mathbf{v}$ is the matrix of eigenvectors and $\mathbf{e}$ is a diagonal matrix of eigenvalues
4. Update $\mathbf{t} = \mathrm{I}(\mathbf{d} = -1) + \mathrm{I}(\mathbf{d} \neq -1) \circ \mathbf{t}$
5. Return to Step 2
6. Set $\tilde{\mathbf{d}}^{\mathrm{M}} = \mathbf{d} + \mathbf{t}$

As above, $\circ$ is elementwise multiplication and, for example, $\mathrm{I}(\mathbf{e} > 0)$ is an indicator function returning a matrix of ones and zeros indicating which elements of $\mathbf{e}$ are greater than zero.

Conceptually, the goal of the algorithm is to create a matrix $\mathbf{t}$ that can be added to $\mathbf{d}$ yielding a $\tilde{\mathbf{d}}$ matrix that corresponds to an identified variance bound. By Lemma 4.2 and Definition 4.3, there are two requirements for $\mathbf{t}$. First it must be positive semi-definite, and, second, elements corresponding to -1's in the matrix $\mathbf{d}$ must equal one. In step 1, $\mathbf{t}$ meets the second criterion, but not the first. In step 3, the algorithm creates an approximation to the initial $\mathbf{t}$ matrix by way of the eigen decomposition that ensures positive semi-definiteness, thus meeting the first criterion. However, due to the approximation, $\mathbf{t}$ no longer meets the second criterion. Therefore, in step 4 the algorithm forces $\mathbf{t}$ to have 1's wherever $\mathbf{d}$ has -1's in order to again meet the second criteria. But doing so means that $\mathbf{t}$ will no longer meet the first criteria. So, the algorithm iterates through steps 2-4 until convergence is achieved (i.e., until all eigenvalues are non-negative in step 2) at which point $\mathbf{t}$ meets both criteria and, thus, $\tilde{\mathbf{d}}^{\mathrm{M}}$ corresponds to an identified bound.

4.4 Comparing bounds

Definition 4.8 (Tighter bound)**.** *Let $\tilde{\mathbf{d}}^a$ and $\tilde{\mathbf{d}}^b$ correspond to two identified bounds. Matrix $\tilde{\mathbf{d}}^a$ is corresponds to a tighter bound than $\tilde{\mathbf{d}}^b$ if $\tilde{\mathbf{d}}^b - \tilde{\mathbf{d}}^a$ is positive semidefinite.*

Definition 4.9 (Invariant bounding matrix)**.** *A matrix $\tilde{\mathbf{d}}$ is an invariant bounding matrix if it is an bounding matrix and if all $n \times n$ partitions, $\tilde{\mathbf{d}}_{ij} 1_{kn} = 0_{kn}$, i.e., all rows of the partition (or, equivalently, all columns) sum to zero.*

Example (*Paired randomization*) : Consider a pair-randomized design, whereby units are "blocked" (i.e., stratified) into groups of two, and then, in each block, one of the two units

is randomly assigned to treatment while the other is assigned to control. When $n = 4$ (and assuming w.l.o.g. that the data are sorted by pair), the matrix $\mathbf{d}$ is

$$\mathbf{d} = \begin{bmatrix} 1 & -1 & & & -1 & 1 & & \\ -1 & 1 & & & 1 & -1 & & \\ & & 1 & -1 & & & -1 & 1 \\ & & -1 & 1 & & & 1 & -1 \\ -1 & 1 & & & 1 & -1 & & \\ 1 & -1 & & & -1 & 1 & & \\ & & -1 & 1 & & & 1 & -1 \\ & & 1 & -1 & & & -1 & 1 \end{bmatrix},$$

noting that empty cells represent 0.

For the pair-randomized design, the Neyman bound cannot be applied because $\mathbf{d}_{00}$ and $\mathbf{d}_{11}$ have negative entries. The Aronow-Samii bound and Algorithm 4.7 have bounding matrices

$$\tilde{\mathbf{d}}^{\text{AS}} = \begin{bmatrix} 3 & & & & 1 & & & \\ & 3 & & 1 & & & & 1 \\ & & 3 & & & 1 & & \\ & 1 & & 3 & & & 1 & \\ 1 & & & & 3 & & & \\ & & 1 & & & 3 & & \\ & & & 1 & & & 3 & \\ & & & & 1 & & & 3 \end{bmatrix}, \quad \text{and} \quad \tilde{\mathbf{d}}^{\text{M}} = \begin{bmatrix} 2 & & & & 2 & & & \\ & 2 & & 2 & & & & 2 \\ & & 2 & & & 2 & & \\ & 2 & & 2 & & & 2 & \\ 2 & & & & 2 & & & \\ & & 2 & & & 2 & & \\ & & & 2 & & & 2 & \\ & & & & 2 & & & 2 \end{bmatrix},$$

respectively. By the Gershgorian circle theorem the difference,

$$\tilde{\mathbf{d}}^{\text{AS}} - \tilde{\mathbf{d}}^{\text{M}} = \begin{bmatrix} 1 & & & & -1 & & & \\ & 1 & & -1 & & & & \\ & & 1 & & & & & -1 \\ & & & 1 & & & -1 & \\ & -1 & & & 1 & & & \\ -1 & & & & & 1 & & \\ & & & -1 & & & 1 & \\ & & -1 & & & & & 1 \end{bmatrix},$$

is positive semi-definite, proving that $\tilde{\mathbf{d}}^{\text{M}}$ corresponds to a tighter variance bound. Confirmation also comes from eigendecomposition of the difference, $\tilde{\mathbf{d}}^{\text{AS}} - \tilde{\mathbf{d}}^{\text{M}}$, which yields all non-negative eigenvalues: 2, 2, 2, 2, 0, 0, 0, and 0.

One might alternatively choose the invariant bounding matrix,

$$\tilde{\mathbf{d}}^{\text{INVAR}} = \begin{bmatrix} 2 & & -1 & -1 & & 2 & -1 & -1 \\ & 2 & -1 & -1 & 2 & & -1 & -1 \\ -1 & -1 & 2 & & -1 & -1 & & 2 \\ -1 & -1 & & 2 & -1 & -1 & 2 & \\ & 2 & -1 & -1 & 2 & & -1 & -1 \\ 2 & & -1 & -1 & & 2 & -1 & -1 \\ -1 & -1 & & 2 & -1 & -1 & 2 & \\ -1 & -1 & 2 & & -1 & -1 & & 2 \end{bmatrix}.$$

It is verifiably a bound because the eigenvalues of $\tilde{\mathbf{d}}^{INV} - \mathbf{d}$ are 8, 0, 0, 0, 0, 0, 0, and 0. However, eigendecomposition of $\tilde{\mathbf{d}}^{INV} - \tilde{\mathbf{d}}^{\text{M}}$ gives eigenvalues 4, 0, 0, 0, 0, 0, 0, and -4, indicating that the better bound may depend on outcome vector, y, and perhaps the estimator as well.

To further adjudicate between these bounds, consider the first order approximation for the Hajek estimator, where the difference in bounds can be written,

$$\begin{aligned} \tilde{\text{V}}^{\text{INVAR}}\left(\widehat{\delta}^{\text{T(HJ)}}\right) - \tilde{\text{V}}^{\text{M}}\left(\widehat{\delta}^{\text{T(HJ)}}\right) =& n^{-2} z^{\text{HJ}\prime}(\tilde{\mathbf{d}}^{\text{INVAR}} - \tilde{\mathbf{d}}^{\text{M}}) z^{\text{HJ}} \\ =& n^{-2} y' \mathbf{m} y, \end{aligned}$$

where $\mathbf{m} = (\mathbf{i}_{2n} - \mathbb{1}(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}')(\tilde{\mathbf{d}}^{\text{INVAR}} - \tilde{\mathbf{d}}^{\text{M}})(\mathbf{i}_{2n} - \mathbb{1}(\mathbb{1}'\mathbb{1})^{-1}\mathbb{1}')$. Eigendecomposition of $\mathbf{m}$ gives eigenvalues 4, 0, 0, 0, 0, 0, 0 and 0, which shows that the bounding matrix $\tilde{\mathbf{d}}^{\text{M}}$ is preferable with respect to the first-order Hajek approximation. Now, Hajek, Horvitz-Thompson and difference-in-means are algebraically equivalent for this design, so all have the same variance in truth. That said, in light of this result (and perhaps if the number of pairs were larger in an actual application), one might be inclined to use general variance bound estimator proposed in (9) in the next section with bounding matrix, $\tilde{\mathbf{d}}^{\text{M}}$, and approximate z-vector, $\mathbf{R}\widehat{z}^{\text{HJ}} = \mathbf{R}\left(\mathbf{i}_{2n} - \mathbb{1}(\mathbb{1}'\mathbf{R}\mathbb{1})^{-1}\mathbb{1}'\right)\mathbf{R}y$.

5 Variance bound estimation

With an identified variance bounds defined and several methods of obtaining matrices, $\tilde{\mathbf{d}}$, this section turns to the subject of variance bound *estimation*.

First define the $kn \times kn$ matrix of probabilities and joint probabilities of assignment,

$$\mathbf{p} := \mathrm{E}\left[\mathbf{R}1_{kn}1_{kn}'\mathbf{R}\right].$$

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

$$\tilde{\mathbf{d}}_{/\mathbf{p}} := \tilde{\mathbf{d}}/\mathbf{p} \tag{7}$$

with / denoting element-wise division defined such that division by zero equals zero. Then an estimator of a variance bound for the Horvitz-Thompson estimator can be written,

$$\widehat{\widetilde{\mathrm{V}}}\left(\widehat{\delta}^{\mathrm{HT}}\right) := n^{-2}y'\mathbf{R}\tilde{\mathbf{d}}_{/\mathbf{p}}\mathbf{R}y. \tag{8}$$

It is unbiased for the variance bound $n^{-2}y'\tilde{\mathbf{d}}y$ because $\mathrm{E}\left[\mathbf{R}\tilde{\mathbf{d}}_{/\mathbf{p}}\mathbf{R}\right] = \tilde{\mathbf{d}}$ by construction. Being inverse-probability weighted, the variance bound estimator in (8) is, itself, a Horvitz-Thompson estimator.

For linear estimators, examples of which are given in Table 1, the bound $n^{-2}z'\tilde{\mathbf{d}}z$ cannot be estimated unbiasedly because the definition of z will often include quantities that, themselves, must be estimated. However, an appeal to the plug-in principle suggests the use of

$$\widehat{\widetilde{\mathrm{V}}}\left(\widehat{\delta}^{\mathrm{T}}\right) := n^{-2}\widehat{z}'\mathbf{R}\tilde{\mathbf{d}}_{/\mathbf{p}}\mathbf{R}\widehat{z} \tag{9}$$

with $\widehat{z}$ containing sample analogues for elements of z where necessary. The usual asymptotic arguments can provide further justification.

Example (*Subsuming standard errors that go by names such as Eicker-Huber-White, "heteroskedastic consistent", "sandwich", and "robust"*) : For the OLS estimator, z^{OLS} is defined in Table (1). The plug-in principle motivates the use of

$$\widehat{z}^{\mathrm{OLS}} = \mathrm{diag}\left(\widehat{u}\right)\boldsymbol{\pi}\mathbb{x}\left(\mathbb{x}'\mathbf{R}\mathbb{x}\right)^{-1}c,$$

where $\widehat{u} = y - \mathbb{x}\widehat{b}^{\mathrm{OLS}}$ and $\widehat{b}^{\mathrm{OLS}} = \left(\mathbb{x}'\mathbf{R}\mathbb{x}\right)^{-1}\mathbb{x}'\mathbf{R}y$ is the OLS coefficient. Then from equation (9) we have,

$$\begin{aligned}\widehat{\widetilde{\mathrm{V}}}\left(\widehat{\delta}^{\mathrm{T(OLS)}}\right) =& c'\left(\mathbb{x}'\mathbf{R}\mathbb{x}\right)^{-1}\mathbb{x}'\boldsymbol{\pi}\mathrm{diag}\left(\widehat{u}\right)\mathbf{R}\tilde{\mathbf{d}}_{/\mathbf{p}}\mathbf{R}\mathrm{diag}\left(\widehat{u}\right)\boldsymbol{\pi}\mathbb{x}\left(\mathbb{x}'\mathbf{R}\mathbb{x}\right)^{-1}c \\ =& c'\left(\mathbb{x}'\mathbf{R}\mathbb{x}\right)^{-1}\mathbb{x}'\mathrm{diag}\left(\mathbf{R}\widehat{u}\right)\left(\boldsymbol{\pi}\tilde{\mathbf{d}}_{/\mathbf{p}}\boldsymbol{\pi}\right)\mathrm{diag}\left(\mathbf{R}\widehat{u}\right)\mathbb{x}\left(\mathbb{x}'\mathbf{R}\mathbb{x}\right)^{-1}c.\end{aligned}$$

This is a the variance bound estimator in (9) made specific to OLS. So far it is applicable to virtually any design and any variance bound.

Next, specify a Bernoulli design, in which units are assigned independently to treatment. (Probabilities of assignment may be equal across units, but they need not be in this example.) In this design, the diagonal elements of $\mathbf{d}$ are equal to the diagonal of $\boldsymbol{\pi}^{-1} - \mathbf{i}_{kn}$, where $\mathbf{i}_{kn}$ is an identity matrix. Further, any of the above bounding methods yields $\tilde{\mathbf{d}} = \boldsymbol{\pi}^{-1} - \mathbf{i} + \mathbf{i} = \boldsymbol{\pi}^{-1}$.

Thus $\tilde{\mathbf{d}}_{/\mathbf{p}} = \boldsymbol{\pi}^{-2}$ so that $\left(\boldsymbol{\pi}\tilde{\mathbf{d}}_{/\mathbf{p}}\boldsymbol{\pi}\right) = \mathbf{i}_{kn}$ is the identity matrix. So the OLS variance bound estimator for Bernoulli designs simplifies to,

$$\widehat{\tilde{\mathrm{V}}}^{\mathrm{B}}\left(\widehat{\delta}^{\mathrm{T(OLS)}}\right) = c'\left(\mathbb{x}'\mathbf{R}\mathbb{x}\right)^{-1}\mathbb{x}'\mathrm{diag}\left(\mathbf{R}\widehat{u}^2\right)\mathbb{x}\left(\mathbb{x}'\mathbf{R}\mathbb{x}\right)^{-1}c.$$

This is White's (1980) canonical "sandwich" variance estimator, sometimes referred to as HC0. The example shows that Eicker-Huber-White standard errors are a special case of (9) for OLS in a Bernoulli design. Note, however, that (9) is much more general. It applies to any linear estimator, virtually any design and any (identified) variance bound.

Remark 10. *Adjustments for degrees of freedom (e.g., HC1) or leverage (e.g., HC2, HC3, etc.) can be applied as well.*

Example (*Subsuming "cluster robust" standard errors*) : Also consider this variance bound estimator for OLS in designs in which clusters are assigned independently to treatment. Then, if we choose the Neyman bound $\tilde{\mathbf{d}}^{\mathrm{N}}_{/\mathbf{p}}$ (or $\tilde{\mathbf{d}}^{\mathrm{M}}_{/\mathbf{p}}$, which is equivalent in the case of for Bernoulli assignment of clusters), and assuming w.l.o.g. that units are sorted by cluster, then $\left(\boldsymbol{\pi}\tilde{\mathbf{d}}^{\mathrm{N}}_{/\mathbf{p}}\boldsymbol{\pi}\right)$ resolves to a block diagonal matrix of 1's with the blocks corresponding to clusters. Hence, (9) also reproduces the "cluster-robust" standard errors sometimes referred to as CR0 as a special case.

5.1 Asymptoodles

Now consider the variance of (n times) the variance estimator:

$$\begin{aligned}
\mathrm{V}\left(n\widehat{\tilde{\mathrm{V}}}\left(\widehat{\delta}^{\mathrm{HT}}\right)\right) =& \mathrm{E}\left[\left(n^{-1}y'\mathbf{R}\tilde{\mathbf{d}}_{/\mathbf{p}}\mathbf{R}y - n^{-1}y'\tilde{\mathbf{d}}y\right)^2\right] \\
\leq& n^{-2}\max\left(|y|\right)^4 1'_{4n^2}\left|\left(\tilde{\mathbf{d}}\otimes\tilde{\mathbf{d}}\right)\circ\left(\mathrm{E}\left[\left(\mathbf{R}1_n1'_n\mathbf{R}\right)\otimes\left(\mathbf{R}1_n1'_n\mathbf{R}\right)\right] - \mathbf{p}\otimes\mathbf{p}\right)/\left(\mathbf{p}\otimes\mathbf{p}\right)\right|1_{4n^2} \\
=& n^{-2}\max\left(|y|\right)^4\sum_i\sum_j\sum_k\sum_l\left|\left(\frac{\pi_{ij}-\pi_i\pi_j}{\pi_i\pi_j}+t_{ij}\right)\left(\frac{\pi_{kl}-\pi_k\pi_l}{\pi_k\pi_l}+t_{kl}\right)\left(\frac{\pi_{ijkl}-\pi_{ij}\pi_{kl}}{\pi_{ij}\pi_{kl}}\right)\right|
\end{aligned}$$

where i, j, k, l index from 1 to kn such that, for example, for $i \leq n$ define $\pi_i := \pi_{0i}$ and for $i > n$ we define $\pi_i := \pi_{1(i-n)}$. And, for example, t_{ij} is the i, j element of the matrix $\mathbf{t} := \left(\tilde{\mathbf{d}} - \mathbf{d}\right)$. That this quantity converges can be checked numerically for a hypothetical sequence of designs and populations.

References

Athey, S., and Imbens, G.W. 2017. The Econometrics of Randomized Experiments. *Handbook of Economic Field Experiments*, **1**: 73-140.

Arceneaux, Kevin, and David Nickerson. 2009. Modeling uncertainty with clustered data: A comparison of methods, Political Analysis, **17**: 177–90.

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. Forthcoming at *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.

Basse, G. and A. Feller. 2017. Analyzing two-stage experiments in the presence of interference, *Journal of the American Statistical Association*, DOI: 10.1080/01621459.2017.1323641

Bloniarz, A., Liu, H., Zhang, C.H., Sekhon, J.S., and Yu, B. 2016. Lasso adjustments of treatment effect estimates in randomized experiments, *Proceedings of the National Academy of Sciences*, **113**(27): 7383-90.

Campbell, Stephen L., and Carl D. Meyer. 2009. Generalized Inverses of Linear Transformations. *https://doi.org/10.1137/1.9780898719048*

Fuller, W.A. 2009. *Sampling Statistics.* New Jersey: Wiley.

Fuller, W.A. and C.T. Isaki. 1981. Survey Design Under Superpopulation Models In: *Current Topics in Survey Sampling* Eds: Krewski, D. , J.N.K. Rao, R. Platek. New York, Academic Press.

Freedman, D.A. 2008a. On regression adjustments to experimental data. *Adv. in Appl. Math.* **40** 180–193.

Freedman, D.A. 2008b. On regression adjustments in experiments with several treatments. *Ann. Appl. Stat.* **2** 176–196.

Hansen, B. and Bowers, J. (2009). Attributing effects to a cluster-randomized get-out-the-vote campaign. *J. Amer. Statist. Assoc.* **104** 873–885.

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.

Isaki, C.T., and W.A. Fuller. 1982. Survey Design Under the Regression Superpopulation Model. *Journal of the American Statistical Association* **77**(377): 89-96

Li, Xinran and Ding, Peng. 2017. General Forms of Finite Population Central Limit Theorems with Applications to Causal Inference. *Journal of the American Statistical Association* **112**(520): 1759-1769

Li, Xinran, Peng Ding, and Donald B. Rubin. 2017. Asymptotic Theory of Rerandomization in Treatment-Control Experiments. *arXiv:1604.00698*

Lin, Winston. 2013. Agnostic Notes on Regression Adjustments to Experimental Data: Reexamining Freedman's Critique. *Annals of Applied Statistics* **7**(1): 295-318

Lu, J. 2016. Covariate adjustment in randomization-based causal inference for 2k factorial designs. *Statistics & Probability Letters*, **119**:1120.

Middleton, J.A. 2008. Bias of the regression estimator for experiments using clustered random assignment. *Stat. Probability Lett.* **78** 2654–2659.

Middleton, Joel A. 2018. A unified theory of regression adjustment for design-based inference. *arXiv:1803.06011*

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.

Raj, D. 1965. On a method of using multi-auxiliary information in sample surveys. *J. Amer. Statist. Assoc.* **60** 270–277.

Rohde, Charles. 1965. Generalized Inverses of Partitioned Matrices. *Journal of the Society for Industrial and Applied Mathematics* **13**(4): 1033-1035.

Rubin, Donald. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* **66**: 688–701.

Sarndal, C.-E., B. Swensson, and J. Wretman. 1992. Model Assisted Survey Sampling. New York: Springer.

Samii, Cyrus and Peter M. Aronow. 2012. On Equivalencies Between Design-Based and Regression-Based Variance Estimators for Randomized Experiments. *Statistics and Probability Letters*. **82**: 365–370.

Schochet, Peter Z. 2010. Is regression adjustment supported by the Neyman model for causal inference? *Journal of Statistical Planning and Inference* **140**: 246-259.

Sinclair, B., McConnell, M. and Green, D.P. 2012. Detecting Spillover Effects: Design and Analysis of Multilevel Experiments. *American Journal of Political Science* **56**(4): 1055-1069.

Wood, John. 2008. On the Covariance Between Related Horvitz-Thompson Estimators. *Journal of Official Statistics*. **24** 53–78.

Zhao, A., Ding, P., Mukerjee, R., and Dasgupta, T. 2017+. Randomization-Based Causal Inference From Unbalanced 2^2 Split-Plot Designs. *Annals of Statistics*, in press.

A Notation Index

n	Number of units in the finite population in the experiment
1_{kn}	Length-kn column vector of 1's. In matrix notation, serves as a replacement for the more common summation symbol, Σ
y_{0i}, y_{1i}	The control and treatment potential outcomes for the i^{th} unit, respectively
y_0, y_1	Length-n vectors of control and treatment potential outcomes, respetively
y	Length-kn vector of all potential outcomes. The first n elements are control potential outcomes multiplied by -1, followed by the treatment potential outcomes. Multiplication of control potential outcomes by -1 allows for the compact representation of the ATE as the sum of the elements of this vector divided by n
δ	Average treatment effect (ATE), the parameter of interest
R_{0i}, R_{1i}	Random indicators of the i^{th} unit's assignment to control and treatment, respectively
R_0, R_1	Length-n vectors of assignment indicators for control and treatment, respectively
$\mathbf{R}$	$kn \times kn$ diagonal matrix of assignment indicators. The first n diagonal elements represent the control indicators, followed by n treatment indicators
π_{0i}, π_{1i}	For the i^{th} unit, the probability of assignment to control and treatment, respectively
π_0, π_1	Length-n vectors of probabilities of assignment to control and treatment, respectively
$\boldsymbol{\pi}$	$kn \times kn$ diagonal matrix of assignment probabilities. The first n diagonal elements give the control probabilities, followed by the treatment probabilities
π_{0i0j}, π_{0i1j}, π_{1i0j}, π_{1i1j}	Joint assignment probabilities for units i and j. For example, π_{1i0j} is the probability that i is in treatment and j is in control
$\mathbf{d}$	$kn \times kn$ "design" matrix that gives the variance-covariance matrix of the vector $1'_{kn}\boldsymbol{\pi}^{-1}\mathbf{R}$. Allows for compact representation of variance of HT estimators as a quadratic in matrix form
$\mathbf{d}_{00}$, $\mathbf{d}_{01}$, $\mathbf{d}_{10}$, $\mathbf{d}_{11}$	The four $n \times n$ partitions of the matrix $\mathbf{d}$. For example, the top-right partition, $\mathbf{d}_{01}$, has i,j element $\frac{\pi_{0i1j}-\pi_{0i}\pi_{1j}}{\pi_{0i}\pi_{1j}}$
$\tilde{\mathbf{d}}$	A modified version of $\mathbf{d}$ that allows for compact representation of a variance *bound* for HT estimators as a quadratic in matrix form. While the variance of the HT estimator is not identified, a variance bound may be

$\mathbf{p}$	$kn \times kn$ "probability" matrix that gives the joint assignment probabilities
$\mathbf{p}_{00}$, $\mathbf{p}_{01}$, $\mathbf{p}_{10}$, $\mathbf{p}_{11}$	The four $n \times n$ quadrants of the matrix $\mathbf{p}$. For example, $\mathbf{p}_{01}$ has ij element π_{0i1j}
$\tilde{\mathbf{p}}$	A modified version of $\mathbf{p}$ that replaces zeros with ones. Allows for division by $\tilde{\mathbf{p}}$ without division-by-zero error
x_i	Length-k vector of covariates associated with the i^{th} unit
$\mathbf{x}$	An $n \times k$ matrix of covariates
$\tilde{\mathbf{x}}$	An $n \times (k+1)$ matrix representing the concatenation of an intercept vector, 1_n, and $\mathbf{x}$
$\mathbb{x}$	A $kn \times l$ matrix of covariates. The first n rows are multiplied by -1 to mirror the vector y. Represents an arbitrary specification
$\mathbb{x}_{\text{I}}$	A $kn \times (k+2)$ matrix of covariates. The "common slopes" specification. Elements in the first n rows are multiplied by -1 to mirror the vector y
$\mathbb{x}_{\text{II}}$	A $kn \times (2k+2)$ matrix of covariates. The "separate slopes" specification. Elements in the first n rows are multiplied by -1 to mirror the vector y

B Supplementary Proofs