

Statistica Sinica Preprint No: SS-2020-0149	
Title	Causal Inference from Possibly Unbalanced Split-Plot Designs: A Randomization-based Perspective
Manuscript ID	SS-2020-0149
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202020.0149
Complete List of Authors	Rahul Mukerjee and Tirthankar Dasgupta
Corresponding Author	Tirthankar Dasgupta
E-mail	td370@stat.rutgers.edu
Notice: Accepted version subject to English editing.	

CAUSAL INFERENCE FROM POSSIBLY UNBALANCED SPLIT-PLOT DESIGNS: A RANDOMIZATION-BASED PERSPECTIVE

Rahul Mukerjee and Tirthankar Dasgupta

Indian Institute of Management Calcutta and Rutgers University

Abstract:

Split-plot designs find wide applicability in multifactor experiments with randomization restrictions. Practical considerations often warrant the use of unbalanced designs. This paper investigates randomization based causal inference in split-plot designs that are possibly unbalanced. Extension of ideas from the recently studied balanced case yields an expression for the sampling variance of a treatment contrast estimator as well as a conservative estimator of the sampling variance. However, the bias of this variance estimator does not vanish even when the treatment effects are strictly additive. A careful and involved matrix analysis is employed to overcome this difficulty, resulting in a new variance estimator, which becomes unbiased under milder conditions. A construction procedure that generates such an estimator with minimax bias is proposed. Empirical studies suggest superiority of the latter estimator with respect to bias uniformly across different populations, and such superiority does not come at a heavy price of large inflation of the mean squared error.

Key words and phrases: Bias, Factorial experiment, Finite population, Minimax-
ity, Treatment-effect additivity.

1. Introduction

Factorial experiments were originally developed in the context of agricultural experiments (Fisher, 1925, 1935; Yates, 1935) and later extensively used in industrial and engineering applications (Wu and Hamada, 2009). Such experiments have currently been undergoing a third popularity surge among social, behavioral, and biomedical sciences. However, one of the key challenges of using standard principles of designing and analyzing factorial experiments in these fields arises from randomization restrictions. Consider a simplified version of the education experiment described in Dasgupta et al. (2015). Suppose the goal is to assess the causal effects of two interventions (referred to as factors in experimental design literature) – F_1 : a mid-year quality review by a team of experts, and F_2 : a bonus scheme for teachers – on the performances of 40 schools in the state of New York. Each factor has two levels denoted by 1 (application) and 0 (non-application). A completely randomized assignment of the 40 schools to the four treatment combinations 00, 01, 10, 11 is likely to disperse the schools assigned to level 1 of factor F_1 (i.e., schools to undergo review) all over the state. Such a

design may be prohibitive from the consideration of travel cost and time. A more practical alternative would be to divide these 40 schools by geographic proximity into four groups called whole-plots (WPs). Two of these WPs would then be randomly assigned to level 0 and the other two to level 1 of factor F_1 . The teacher bonus scheme can then be applied to half of the schools chosen randomly within each WP. Such a randomization scheme is an example of a classic split-plot design. See Kirk (1982), Cochran and Cox (1957), Box et al. (2005), and Wu and Hamada (2009) for formal definitions.

Randomization-based inference is a useful methodology to draw inference on causal effects of treatments from split-plot experiments in a finite population setting, as observed by Freedman (2006, 2008). The main advantage of randomization-based inference is the fact that it applies even if the experimental units are not randomly sampled from a larger population, which is the case in most social science experiments and clinical trials (Abadie et al., 2020; Olsen et al., 2013; Rosenberger et al., 2019). Recently, Zhao et al. (2018) developed a framework for randomization-based estimation procedure of finite-population causal effects for balanced split-plot designs, in which each WP consists of the same number of units or sub-plots (SPs), and any treatment combination of the SP factors occurs equally often in all WPs; vide (4) below. However, unbalanced split-plot

designs are quite common in the social sciences. Consider the school experiment described earlier. Suppose the 40 schools are spread over four counties with 8, 8, 12 and 12 schools in these counties. In this case, each county can be considered as a natural WP. Thus the design is unbalanced and the estimation methodology proposed by Zhao et al. (2018) is no longer applicable.

In this paper we investigate randomization based causal inference in split-plot designs that are possibly unbalanced, using the potential outcomes framework originally introduced by Neyman (1923), formalized in Rubin (1974) and extended in the subsequent works of Rubin. We start with a natural unbiased estimator of a typical treatment contrast and first examine how far the approach of Zhao et al. (2018) for the balanced case can be adapted to our more general setup. It is seen that this approach, aided by a variable transformation, yields an expression for the sampling variance of the treatment contrast estimator but runs into difficulty in variance estimation. Specifically, as in the balanced case and other situations in causal inference, the resulting variance estimator is conservative in the sense of having a nonnegative bias. However, unlike in most standard situations, the bias does not vanish even under strict additivity or homogeneity of treatment effects. To overcome this problem, a careful matrix analysis is

employed leading, under wide generality, to a new variance estimator. This estimator is also conservative, but enjoys the nice property of becoming unbiased under between-WP additivity, a condition even milder than strict additivity. We also discuss the issue of minimaxity, with a view to controlling the bias in variance estimation and explore the bias and mean squared error (MSE) of the two estimators under treatment effect heterogeneity via simulations. Proofs of all results appear in the supplementary material.

2. Notation and background

Consider a factorial experiment conducted to assess causal effects of m_1 WP factors F_{11}, \dots, F_{1m_1} and m_2 SP factors F_{21}, \dots, F_{2m_2} on a finite population of N units. Each factor has two or more levels. The treatment combinations are denoted by $z = z_1 z_2$, where $z_k \in Z_k$ and Z_k is the set of factor levels of F_{k1}, \dots, F_{km_k} ($k = 1, 2$). For $i = 1, \dots, N$, let $Y_i(z_1 z_2)$ denote the potential outcome of unit i when exposed to treatment combination $z_1 z_2$. This notation assumes the Stable Unit Treatment Value Assumption (Rubin, 1980), which means (1) that the potential outcome of unit i depends only on the treatment combination it is assigned to, and (2) that there are no hidden versions of treatments not represented by all level combinations of the $m_1 + m_2$ factors. A typical treatment contrast for unit i is of the

form

$$\tau_i = \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} g(z_1 z_2) Y_i(z_1 z_2), \quad (1)$$

where $g(z_1 z_2)$, $z_1 \in Z_1, z_2 \in Z_2$ are known, not all zeros, and sum to zero.

In the school example, unit i has four potential outcomes $Y_i(00)$, $Y_i(01)$, $Y_i(10)$, $Y_i(11)$. Following Dasgupta et al. (2015), the unit-level main effect of factor F_1 is defined as $\tau_i^{F_1} = \{-Y_i(00) - Y_i(01) + Y_i(10) + Y_i(11)\}/2$. The contrast coefficients $g(00), g(01), g(10), g(11)$ are $-1/2, -1/2, 1/2, 1/2$ respectively. The contrast coefficients for the main effect of factor F_2 and the interaction $F_1 F_2$ can be similarly defined. Let

$$\bar{Y}(z_1 z_2) = N^{-1} \sum_{i=1}^N Y_i(z_1 z_2), \quad (2)$$

denote the average potential outcome for treatment combination $z_1 z_2$, and let

$$\bar{\tau} = N^{-1} \sum_{i=1}^N \tau_i = \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} g(z_1 z_2) \bar{Y}(z_1 z_2), \quad (3)$$

denote a treatment contrast for the finite population of N units. We define $\bar{\tau}$ as the finite-population causal estimand of interest and consider the problem of drawing inference on $\bar{\tau}$ using the outcomes observed from the experiment.

The observed outcomes are generated through an assignment mechanism, which is the process of allocating treatment combinations to the N units. Here we consider a split-plot assignment mechanism which can be

described as follows. Suppose there is a partitioning of the N experimental units into $W(\geq 2)$ disjoint sets $\Omega_1, \dots, \Omega_W$, called WPs, such that Ω_w consists of $M_w(\geq 2)$ units, called SPs, $w = 1, \dots, W$, and $M_1 + \dots + M_W = N$. Consider now a two-stage randomization, which assigns $r_1(z_1)$ (≥ 2) WPs to level combination z_1 of F_{11}, \dots, F_{1m_1} and then, for each $w = 1, \dots, W$, assigns $r_{w2}(z_2)$ SPs within WP Ω_w to level combination z_2 of F_{21}, \dots, F_{2m_2} . Here at each stage all assignments are equiprobable, the $r_1(z_1)$ and $r_{w2}(z_2)$ are fixed positive integers, and $\sum_{z_1 \in Z_1} r_1(z_1) = W$, $\sum_{z_2 \in Z_2} r_{w2}(z_2) = M_w$ for $w = 1, \dots, W$.

Note that the above assignment mechanism yields a *balanced* split-plot design if

$$M_1 = \dots = M_W, \quad r_{12}(z_2) = \dots = r_{W2}(z_2), \text{ for all } z_2 \in Z_2. \quad (4)$$

In the school example described in Section 1, the WPs represent sets of schools within a county and we have $N = 40$, $W = 4$, $M_1 = M_2 = 8$, $M_3 = M_4 = 12$, $Z_1 = Z_2 = \{0, 1\}$. Finally, for all $z_2 \in Z_2$, $r_{w2}(z_2) = 4$ for $w = 1, 2$ and $r_{w2}(z_2) = 6$ for $w = 3, 4$. Thus, the design is unbalanced.

To define the observed outcomes of the experiment, we introduce two sets of random treatment assignment indices at the WP and the SP levels. Let $T_1(z_1)$ denote the set of indices w such that WP Ω_w is randomly assigned to level combination z_1 of F_{11}, \dots, F_{1m_1} . Similarly, for $z_2 \in Z_2$ and $w =$

$1, \dots, W$, let $T_{w2}(z_2)$ be the set of SPs in Ω_w randomly assigned to level combination z_2 of F_{21}, \dots, F_{2m_2} . For any treatment combination $z_1 z_2$, the observed outcomes from the WP Ω_w , $w \in T_1(z_1)$, are then $Y_i(z_1 z_2)$, $i \in T_{w2}(z_2)$. Let

$$\bar{Y}_w^{\text{obs}}(z_1 z_2) = \{r_{w2}(z_2)\}^{-1} \sum_{i \in T_{w2}(z_2)} Y_i(z_1 z_2), \quad (5)$$

denote the average observed outcome for treatment combination $z_1 z_2$ within WP Ω_w for $w \in T_1(z_1)$. In the spirit of the usual unbiased estimator of the population mean in two-stage sampling (Cochran, 1977), define

$$\bar{Y}^{\text{obs}}(z_1 z_2) = \frac{W}{Nr_1(z_1)} \sum_{w \in T_1(z_1)} M_w \bar{Y}_w^{\text{obs}}(z_1 z_2) = \frac{1}{r_1(z_1)} \sum_{w \in T_1(z_1)} \frac{M_w}{\bar{M}} \bar{Y}_w^{\text{obs}}(z_1 z_2), \quad (6)$$

where $\bar{M} = (M_1 + \dots + M_W)/W = N/W$ is the average WP size. From (5) and (6), it is straightforward to verify by conditioning on the randomization at the WP level that $E \left\{ \bar{Y}^{\text{obs}}(z_1 z_2) \right\} = \bar{Y}(z_1 z_2)$, where $\bar{Y}(z_1 z_2)$ is given by (2). Using (3), an immediate consequence of this fact is Proposition 1.

Proposition 1. *An unbiased estimator of the finite population treatment contrast $\bar{\tau}$ is given by*

$$\hat{\bar{\tau}} = \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} g(z_1 z_2) \bar{Y}^{\text{obs}}(z_1 z_2), \quad (7)$$

where $\bar{Y}^{\text{obs}}(z_1 z_2)$ is given by (6).

2.1 Why don't results from block designs work?⁹

Proposition 1 yields a point estimator of $\bar{\tau}$. However, to quantify the uncertainty associated with the point estimator, one needs to derive and estimate the sampling variance of $\hat{\bar{\tau}}$ with respect to its distribution induced by the randomization in the split-plot design. In the next two subsections we briefly describe two areas of research that are related to our problem setting. The first is the case of *unbalanced block designs*, that we discuss in Section 2.1 and argue why the results from this setting do not apply to ours. The second is the case of *balanced split-plot designs*, that motivates our work.

2.1 Why don't results from block designs work?

The problem of Neymanian variance estimation has been recently investigated for block designs that are unbalanced, i.e., have unequal block sizes, and the results (Pashley and Miratrix, 2019) have been applied to other settings (Schochet et al., 2020). A natural question that arises is, whether the ideas used to derive variance estimators for unbalanced block designs that have desirable properties can be applied to the case of split-plot designs. As seen from the subsequent discussion, the fundamentally different setting of the two designs does not permit such adaptation.

Consider a block design in W blocks $\Omega_1, \dots, \Omega_W$ of sizes M_1, \dots, M_W ,

2.1 Why don't results from block designs work?10

such that any treatment combination $z_1 z_2$ is randomly assigned to $r_w(z_1 z_2)$ units of Ω_w , $w = 1, \dots, W$. The fixed positive integers $r_w(z_1 z_2)$ sum to M_w for each w . Subject to this, all treatment assignments within each block are equally likely. Also, randomization is done independently for different blocks. Similar to (3) (See also (15) and (16) defined later), let $\bar{\tau}_1, \dots, \bar{\tau}_W$ be block level treatment contrasts, leading to the population level treatment contrast $\bar{\tau} = (1/W) \sum_{w=1}^W (M_w/\bar{M}) \bar{\tau}_w$, where $\bar{M} = N/W$ and $N = M_1 + \dots + M_W$. If $\hat{\tau}_w$ is an unbiased estimator of $\bar{\tau}_w$ on the basis of observed responses from Ω_w , then $\bar{\tau}$ is unbiasedly estimated by

$$\hat{\bar{\tau}} = (1/W) \sum_{w=1}^W (M_w/\bar{M}) \hat{\tau}_w. \quad (8)$$

The sampling variance of this estimator is $(1/W)^2 \sum_{w=1}^W (M_w/\bar{M})^2 \text{var}(\hat{\tau}_w)$, because $\bar{\tau}_1, \dots, \bar{\tau}_W$ are independent across blocks. In the special case of a single treatment factor with two levels, Pashley and Miratrix (2019) proposed several estimators of this variance depending on the composition of the blocks. For blocks containing at least two treated and two control units, this estimator is straightforward and is obtained by substituting a conservative estimator of $\text{var}(\hat{\tau}_w)$ that vanishes under strict additivity. For blocks that contain only one treated or control unit, Pashley and Miratrix (2019) proposed two estimators based on the weighted version of $\sum_{w=1}^W (\hat{\tau}_w - \hat{\bar{\tau}})^2$. Both estimators are conservative in general, but unbiased under the assump-

2.1 Why don't results from block designs work?¹¹

tion of block-level additivity, i.e., the average treatment effect is the same across all blocks. This approach is easily extendable to experiments with multiple treatments and factorial experiments.

To see why *neither of the above estimators can be defined* for split-plot designs, note that in order for them to be defined, the estimator $\widehat{\tau}_w$ of the WP-level contrast $\bar{\tau}_w$ needs to be defined for each $w = 1, \dots, W$. However, in a split-plot design, $\widehat{\tau}_w$ cannot be defined for *any* w because of the WP-level randomization. To understand this further, consider a toy example with twelve units and two factors (each at two levels 0 and 1) in two settings shown in Table 2.1: (i) in $W = 2$ blocks of sizes $M_1 = 8$ and $M_2 = 4$ (left half) and (ii) in $W = 2$ WPs of sizes $M_1 = 8$ and $M_2 = 4$ (right half), where the first factor is the WP-factor. Assume that the contrast τ of interest is the main effect of the second factor. The unit-level contrasts $\tau_i = \{-Y_i(00) + Y_i(01) - Y_i(10) + Y_i(11)\}/2$, $i = 1, \dots, 12$ and the population level contrast $\bar{\tau}$ are the same in both situations. To make the two settings comparable, assume that for the block design $r_1(z_1 z_2) = 2$ and $r_2(z_1 z_2) = 1$ for $z_1, z_2 \in \{0, 1\}$ and for the split-plot design $r_1(z_1) = 1$, $r_{12}(z_2) = 4$, $r_{22}(z_2) = 2$. One possible realization of the treatment assignment for each design is shown in Table 2.1. In the split-plot design, the first WP receives level 0 of the first factor, whereas the second WP receives level 1.

2.1 Why don't results from block designs work?12

For the block design, $\widehat{\tau}_w$ is readily defined for each block w . The block-level variance $\text{var}(\widehat{\tau}_w)$ can be estimated for $w = 1$ but not for $w = 2$, and hence the estimators based on $\sum_{w=1}^W \left(\widehat{\tau}_w - \widehat{\bar{\tau}}\right)^2$ can be used. However, for the split-plot design, neither $\widehat{\tau}_w$ nor $\text{var}(\widehat{\tau}_w)$ can be defined for any of the WPs because we observe potential outcomes for only two of the four treatment combinations within each WP.

Table 1: Block design versus Split-plot design

Block	Unit	Factor levels				WP	Unit	Factor levels			
		00	01	10	11			00	01	10	11
1	1	X				1	1	X			
	2	X					2	X			
	3		X				3		X		
	4		X				4		X		
	5			X			5	X			
	6			X			6	X			
	7				X		7		X		
	8				X		8		X		
2	9	X				2	9			X	
	10		X				10			X	
	11			X			11				X
	12				X		12				X

2.2 Sampling variance and its estimation: Prior work on balanced designs13

2.2 Sampling variance and its estimation: Prior work on balanced designs

Zhao et al. (2018) derived an expression for the sampling variance of $\hat{\tau}$ for a balanced split-plot design, that is, when conditions (4) are satisfied. Under these conditions, we can write $r_{w2}(z_2) = r_2(z_2)$ for $w = 1, \dots, W$ and $z_2 \in Z_2$. Equations (5) and (6) respectively reduce to:

$$\begin{aligned}\bar{Y}_w^{\text{obs}}(z_1 z_2) &= \{r_2(z_2)\}^{-1} \sum_{i \in T_{w2}(z_2)} Y_i(z_1 z_2), \\ \bar{Y}^{\text{obs}}(z_1 z_2) &= \{r_1(z_1)\}^{-1} \sum_{w \in T_1(z_1)} \bar{Y}_w^{\text{obs}}(z_1 z_2).\end{aligned}$$

Next, Zhao et al. (2018) defined the following quantities for any $z_1, z_1^* \in Z_1$ and $z_2, z_2^* \in Z_2$:

$$S_{\text{bt}}(z_1 z_2, z_1^* z_2^*) = \frac{M}{W-1} \sum_{w=1}^W \{\bar{Y}_w(z_1 z_2) - \bar{Y}(z_1 z_2)\} \{\bar{Y}_w(z_1^* z_2^*) - \bar{Y}(z_1^* z_2^*)\}, \quad (9)$$

$$S_{\text{in}}(z_1 z_2, z_1^* z_2^*) = \frac{\sum_{w=1}^W \sum_{i \in \Omega_w} \{Y_i(z_1 z_2) - \bar{Y}_w(z_1 z_2)\} \{Y_i(z_1^* z_2^*) - \bar{Y}_w(z_1^* z_2^*)\}}{W(M-1)}. \quad (10)$$

Here M is the common size of the WPs in the balanced case and $\bar{Y}_w(z_1 z_2) = M^{-1} \sum_{i \in \Omega_w} Y_i(z_1 z_2)$, for each w and $z_1 z_2$. The quantities $S_{\text{bt}}(z_1 z_2, z_1^* z_2^*)$ and $S_{\text{in}}(z_1 z_2, z_1^* z_2^*)$ represent, respectively, the between and within WP mean squares or products in an analysis of variance/covariance

2.2 Sampling variance and its estimation: Prior work on balanced designs14

decomposition of the potential outcomes, i.e., of the quantity

$$\sum_{w=1}^W \sum_{i \in \Omega_w} \{Y_i(z_1 z_2) - \bar{Y}(z_1 z_2)\} \{Y_i(z_1^* z_2^*) - \bar{Y}(z_1^* z_2^*)\}.$$

Also, they defined the following quantity, which is a function of the observed outcomes and can be computed from experimental data:

$$\hat{v}(\hat{\tau}) = \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} \sum_{z_2^* \in Z_2} \{r_1(z_1)\}^{-1} g(z_1 z_2) g(z_1 z_2^*) s(z_1 z_2, z_1 z_2^*), \quad (11)$$

where

$$s(z_1 z_2, z_1 z_2^*) = \sum_{w \in T_1(z_1)} \frac{\{\bar{Y}_w^{\text{obs}}(z_1 z_2) - \bar{Y}^{\text{obs}}(z_1 z_2)\} \{\bar{Y}_w^{\text{obs}}(z_1 z_2^*) - \bar{Y}^{\text{obs}}(z_1 z_2^*)\}}{r_1(z_1) - 1}. \quad (12)$$

We now summarize the main results of Zhao et al. (2018) in the following Theorem on the sampling variance of $\hat{\tau}$ and its estimation.

Theorem 1 (Zhao). (a) *The sampling variance of $\hat{\tau}$ is given by*

$$\begin{aligned} \text{var}_{\text{S-P}}(\hat{\tau}) &= \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} \sum_{z_2^* \in Z_2} \frac{g(z_1 z_2) g(z_1 z_2^*) \{S_{\text{bt}}(z_1 z_2, z_1 z_2^*) - S_{\text{in}}(z_1 z_2, z_1 z_2^*)\}}{M r_1(z_1)} \\ &+ \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} \frac{\{g(z_1 z_2)\}^2 S_{\text{in}}(z_1 z_2, z_1 z_2)}{r_1(z_1) r_2(z_2)} - \frac{\sum_{w=1}^W (\bar{\tau}_w - \bar{\tau})^2}{W(W-1)}, \end{aligned}$$

where

$$\bar{\tau}_w = M^{-1} \sum_{i \in \Omega_w} \tau_i = \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} g(z_1 z_2) \bar{Y}_w(z_1 z_2), \quad (w = 1, \dots, W).$$

(b) $E\{\hat{v}(\hat{\tau})\} \geq \text{var}_{\text{S-P}}(\hat{\tau})$, with equality holding for every treatment contrast τ if and only if between-WP additivity holds, which means $\bar{Y}_w(z_1 z_2) -$

$\bar{Y}_w(z_1^* z_2^*)$ is the same over $w = 1, \dots, W$, for every pair of treatment combinations $z_1 z_2$ and $z_1^* z_2^*$.

Thus, to summarize, Zhao et al. (2018) obtained an estimator of the sampling variance that, like most variance estimators in finite population causal inference (Mukerjee et al., 2018), has a nonnegative bias. Further, they noted that this bias vanishes for every treatment contrast $\bar{\tau}$ if and only if between-WP additivity holds, which means

$$\bar{Y}_1(z_1 z_2) - \bar{Y}_1(z_1^* z_2^*) = \dots = \bar{Y}_W(z_1 z_2) - \bar{Y}_W(z_1^* z_2^*), \quad (13)$$

for every pair of treatment combinations $z_1 z_2$ and $z_1^* z_2^*$.

3. Sampling variance and its estimation generalizing the balanced case

In this section, we derive an expression for the sampling variance and find a variance estimator generalizing the results in Section 2.2 to the unbalanced case, and examine the properties of the estimator. It is worthwhile to note that these results revolve around the quantities $S_{\text{bt}}(z_1 z_2, z_1^* z_2^*)$ and $S_{\text{in}}(z_1 z_2, z_1^* z_2^*)$ defined in equations (9) and (10) respectively, and both of them depend on the common WP size M . Because the M_w 's are different across whole plots in an unbalanced split-plot design, it is difficult to guess

what the counterparts of these two quantities will be in the unbalanced case.

Note that by (2), $\bar{Y}(z_1 z_2) = N^{-1} \sum_{w=1}^W M_w \bar{Y}_w(z_1 z_2)$, where $\bar{Y}_w(z_1 z_2) = M_w^{-1} \sum_{i \in \Omega_w} Y_i(z_1 z_2)$ is the average potential outcome of all units in WP Ω_w for treatment combination $z_1 z_2$. A helpful feature of the balanced case is that $\bar{Y}(z_1 z_2)$ is then the simple average of $\bar{Y}_w(z_1 z_2)$ for $w = 1, \dots, W$, and similarly, $\bar{Y}^{\text{obs}}(z_1 z_2)$ in (6) is also the simple average of the $\bar{Y}_w^{\text{obs}}(z_1 z_2)$, $w \in T_1(z_1)$. We first translate this “simple average” feature to the unbalanced case, with a view to facilitating the derivation there. To that end, we convert the “raw” potential outcomes $Y_i(z_1 z_2)$ to “adjusted” potential outcomes

$$U_i(z_1 z_2) = (M_w / \bar{M}) Y_i(z_1 z_2), \quad (14)$$

for each $z_1 \in Z_1$, $z_2 \in Z_2$, $i \in \Omega_w$ and $w = 1, \dots, W$. For each $z_1 z_2$, define

$$\bar{U}_w(z_1 z_2) = M_w^{-1} \sum_{i \in \Omega_w} U_i(z_1 z_2), \quad w = 1, \dots, W, \quad \text{and} \quad \bar{U}(z_1 z_2) = W^{-1} \sum_{w=1}^W \bar{U}_w(z_1 z_2).$$

By (14), $\bar{U}(z_1 z_2) = \bar{Y}(z_1 z_2)$, i.e., $\bar{Y}(z_1 z_2)$ is the simple average of $\bar{U}_w(z_1 z_2)$, $w = 1, \dots, W$, irrespective of whether M_1, \dots, M_W are equal or not. As seen later in equation (18), a similar simple average relationship holds also between their observed counterparts. The points just noted simplify the derivation to some extent, but additional complications remain to be ad-

dressed, for example, because the $r_{w2}(z_2)$ are not constant over w in the unbalanced case.

Next, for $z_1, z_1^* \in Z_1$ and $z_2, z_2^* \in Z_2$, define

$$S_{\text{bt}}(z_1 z_2, z_1^* z_2^*) = \frac{\overline{M}}{W-1} \sum_{w=1}^W \{ \overline{U}_w(z_1 z_2) - \overline{U}(z_1 z_2) \} \{ \overline{U}_w(z_1^* z_2^*) - \overline{U}(z_1^* z_2^*) \},$$

$$S_{\text{in},w}(z_1 z_2, z_1^* z_2^*) = \frac{1}{M_w-1} \sum_{i \in \Omega_w} \{ U_i(z_1 z_2) - \overline{U}_w(z_1 z_2) \} \{ U_i(z_1^* z_2^*) - \overline{U}_w(z_1^* z_2^*) \}.$$

Some more notation will help. First, Let

$$\bar{\tau}_w = (1/M_w) \sum_{i \in \Omega_w} \tau_i = \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} g(z_1 z_2) \bar{Y}_w(z_1 z_2), \quad w = 1, \dots, W, \quad (15)$$

denote the WP level treatment contrasts, where $\bar{Y}_w(z_1 z_2)$ defined earlier is the average potential outcome of all units in WP Ω_w for treatment combination $z_1 z_2$. The second equality in (15) follows from (1). Also, from (3) and (15), it follows that

$$\bar{\tau} = (1/W) \sum_{w=1}^W (M_w/\overline{M}) \bar{\tau}_w. \quad (16)$$

Now define

$$\Delta = \frac{1}{W(W-1)} \sum_{w=1}^W \{ (M_w/\overline{M}) \bar{\tau}_w - \bar{\tau} \}^2, \quad (17)$$

where $\bar{\tau}_w$ is given by (15). Then, extending the ideas of Zhao et al. (2018), after considerable algebra, we obtain the following result on the sampling variance of $\hat{\bar{\tau}}$, the unbiased estimator of $\bar{\tau}$.

Theorem 2. *The sampling variance of $\widehat{\tau}$ is*

$$\begin{aligned} \text{var}(\widehat{\tau}) &= \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} \sum_{z_2^* \in Z_2} \frac{g(z_1 z_2) g(z_1 z_2^*)}{r_1(z_1)} \left(\frac{S_{\text{bt}}(z_1 z_2, z_1 z_2^*)}{\overline{M}} - \sum_{w=1}^W \frac{S_{\text{in},w}(z_1 z_2, z_1 z_2^*)}{W M_w} \right) \\ &+ \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} \frac{\{g(z_1 z_2)\}^2}{W r_1(z_1)} \sum_{w=1}^W \frac{S_{\text{in},w}(z_1 z_2, z_1 z_2)}{r_{w2}(z_2)} - \Delta. \end{aligned}$$

Next, to obtain an estimator of the sampling variance, we first define the counterparts of $\overline{Y}_w^{\text{obs}}(z_1 z_2)$ and $\overline{Y}^{\text{obs}}(z_1 z_2)$ in (5) and (6) in terms of the adjusted potential outcomes:

$$\begin{aligned} \overline{U}_w^{\text{obs}}(z_1 z_2) &= \frac{1}{r_{w2}(z_2)} \sum_{i \in T_{w2}(z_2)} U_i(z_1 z_2), \quad w \in T_1(z_1) \quad \text{and} \\ \overline{U}^{\text{obs}}(z_1 z_2) &= \frac{1}{r_1(z_1)} \sum_{w \in T_1(z_1)} \overline{U}_w^{\text{obs}}(z_1 z_2). \end{aligned}$$

Then it is easy to see from (5), (6) and (14) that

$$\overline{Y}^{\text{obs}}(z_1 z_2) = \overline{U}^{\text{obs}}(z_1 z_2). \quad (18)$$

Note that $\overline{U}^{\text{obs}}(z_1 z_2)$ is the simple average of $\overline{U}_w^{\text{obs}}(z_1 z_2)$, $w \in T_1(z_1)$, irrespective of whether M_1, \dots, M_W are equal or not. This is precisely what the relationship between $\overline{Y}^{\text{obs}}(z_1 z_2)$ and $\overline{Y}_w^{\text{obs}}(z_1 z_2)$ in (6) reduces to when $M_1 = \dots = M_W$, providing us with the intuition to generalize the results of Zhao et al. (2018) via the use of the adjusted potential outcomes, in particular, replacing $\overline{Y}^{\text{obs}}(z_1 z_2)$ in (7) by $\overline{U}^{\text{obs}}(z_1 z_2)$ because of (18). We

now define the following estimator of the sampling variance in Theorem 2:

$$\hat{V}(\hat{\tau}) = \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} \sum_{z_2^* \in Z_2} \frac{g(z_1 z_2) g(z_1 z_2^*)}{r_1(z_1)} \hat{S}(z_1 z_2, z_1 z_2^*), \quad (19)$$

where

$$\hat{S}(z_1 z_2, z_1 z_2^*) = \frac{1}{r_1(z_1) - 1} \sum_{w \in T_1(z_1)} \left\{ \bar{U}_w^{\text{obs}}(z_1 z_2) - \bar{U}^{\text{obs}}(z_1 z_2) \right\} \left\{ \bar{U}_w^{\text{obs}}(z_1 z_2^*) - \bar{U}^{\text{obs}}(z_1 z_2^*) \right\}.$$

Again, considerable algebra yields the following result:

Theorem 3. *The variance estimator $\hat{V}(\hat{\tau})$ given by (19) estimates the sampling variance of $\hat{\tau}$ with a nonnegative bias Δ defined by (17), that is, $E \left\{ \hat{V}(\hat{\tau}) \right\} = \text{var}(\hat{\tau}) + \Delta$.*

Remark 1. Theorem 3 shows that $\hat{V}(\hat{\tau})$ is a conservative estimator of $\text{var}(\hat{\tau})$ with a non-negative bias Δ . This property is in line with variance estimators in other situations of randomization based causal inference. Moreover, in the balanced case, by (17), the bias Δ vanishes when $\bar{\tau}_1 = \dots = \bar{\tau}_W = \bar{\tau}$, which happens for every treatment contrast $\bar{\tau}$ if and only if (13), i.e., between-WP additivity holds. This shows how the results of Zhao et al. (2018) presented in Section 2.2 follow from Theorem 3. A disturbing feature of the variance estimator $\hat{V}(\hat{\tau})$, however, emerges in the unbalanced case which is the main focus of this paper. Then $\hat{V}(\hat{\tau})$

remains biased even if between-WP additivity holds, because by (15) and (16), condition (13) implies $\bar{\tau}_1 = \dots = \bar{\tau}_W = \bar{\tau}$ and hence

$$\Delta = \frac{\bar{\tau}^2}{W(W-1)\bar{M}^2} \sum_{w=1}^W (M_w - \bar{M})^2,$$

which is positive when M_1, \dots, M_W are not all equal unless $\bar{\tau} = 0$. The situation remains unchanged even under the stronger assumption of strict additivity or homogeneity of treatment effects (Neyman, 1923), which enforces the constancy of $Y_i(z_1 z_2) - Y_i(z_1^* z_2^*)$ over $i = 1, \dots, N$ for every pair of treatment combinations $z_1 z_2$ and $z_1^* z_2^*$.

This property of $\hat{V}(\hat{\tau})$ described in Remark 1 is a matter of concern because a requirement typically imposed on a variance estimator in causal inference is that it should become unbiased at least under Neymannian strict additivity, if not under milder versions thereof such as between-WP additivity in the present context. The estimator $\hat{V}(\hat{\tau})$, obtained by generalizing the arguments in the balanced case fails to meet this requirement when M_1, \dots, M_W are not all equal. In the rest of the paper, we investigate the existence of a variance estimator that overcomes this difficulty and show how, under wide generality, such an estimator can be obtained by appropriately modifying $\hat{V}(\hat{\tau})$ as given by (19).

4. A new variance estimator

We begin our search for an improved variance estimator by expanding the bias term Δ defined in (17) as follows:

$$\Delta = (1/N)^2 \left[\sum_{w=1}^W M_w^2 \bar{\tau}_w^2 - \sum_{w=1}^W \sum_{w^* (\neq w)=1}^W \{M_w M_{w^*} / (W-1)\} \bar{\tau}_w \bar{\tau}_{w^*} \right]. \quad (20)$$

Note that in (20), the term $\bar{\tau}_w^2$ is not unbiasedly estimable, but for $w \neq w^*$, $\bar{\tau}_w \bar{\tau}_{w^*}$ allows unbiased estimation. This is because, by (15),

$$\bar{\tau}_w^2 = (1/M_w)^2 \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} \sum_{z_1^* \in Z_1} \sum_{z_2^* \in Z_2} \sum_{i \in \Omega_w} \sum_{i^* \in \Omega_w} g(z_1 z_2) g(z_1^* z_2^*) Y_i(z_1 z_2) Y_{i^*}(z_1^* z_2^*). \quad (21)$$

The sums over i and i^* in (21) include the case $i = i^*$. There is at least one pair of distinct treatment combinations $z_1 z_2$ and $z_1^* z_2^*$ such that $g(z_1 z_2) g(z_1^* z_2^*) \neq 0$ and $Y_i(z_1 z_2) Y_{i^*}(z_1^* z_2^*)$ is never observable as unit i cannot be assigned simultaneously to both $z_1 z_2$ and $z_1^* z_2^*$. Hence, $\bar{\tau}_w^2$ does not allow unbiased estimation. On the other hand, for $w \neq w^*$, $\bar{\tau}_w \bar{\tau}_{w^*}$ does not involve terms like $Y_i(z_1 z_2) Y_{i^*}(z_1^* z_2^*)$, and is unbiasedly estimable. For each w , let z_{1w} denote the level combination of the WP factors assigned to WP Ω_w . Now define

$$G_w^{\text{obs}} = \sum_{z_2 \in Z_2} g(z_{1w} z_2) \bar{Y}_w^{\text{obs}}(z_{1w} z_2).$$

Proposition 2. For $w, w^* = 1, \dots, W$, $w \neq w^*$, an unbiased estimator of $\bar{\tau}_w \bar{\tau}_{w^*}$ is given by

$$H_{ww^*} = \frac{W(W-1)G_w^{\text{obs}}G_{w^*}^{\text{obs}}}{r_1(z_{1w})\{r_1(z_{1w^*}) - \delta(z_{1w}, z_{1w^*})\}},$$

where $\delta(z_{1w}, z_{1w^*})$ is an indicator that equals one if $z_{1w} = z_{1w^*}$ and zero otherwise.

We can now use Proposition 2 to construct a new estimator of $\text{var}(\hat{\bar{\tau}})$.

Consider any symmetric matrix $B = ((b_{ww^*}))$ of order W such that $b_{ww} = M_w^2$ for $w = 1, \dots, W$. Now define the variance estimator

$$\tilde{V}(\hat{\bar{\tau}}) = \hat{V}(\hat{\bar{\tau}}) + (1/N^2) \sum_{w=1}^W \sum_{w^* (\neq w)=1}^W [b_{ww^*} + \{M_w M_{w^*}/(W-1)\}] H_{ww^*}, \quad (22)$$

where $\hat{V}(\hat{\bar{\tau}})$ is the variance estimator defined in Section 3, and H_{ww^*} is as defined in Proposition 2. Then, from (20), (22), Theorem 3 and Proposition 2 it is easy to see that

$$\begin{aligned} E \left\{ \tilde{V}(\hat{\bar{\tau}}) \right\} &= \text{var}(\hat{\bar{\tau}}) + \Delta + (1/N^2) \sum_{w=1}^W \sum_{w^* (\neq w)=1}^W [b_{ww^*} + \{M_w M_{w^*}/(W-1)\}] \bar{\tau}_w \bar{\tau}_{w^*} \\ &= \text{var}(\hat{\bar{\tau}}) + \tilde{\Delta}, \end{aligned}$$

where

$$\tilde{\Delta} = (1/N^2) \sum_{w=1}^W \sum_{w^*=1}^W b_{ww^*} \bar{\tau}_w \bar{\tau}_{w^*}. \quad (23)$$

Clearly, if the matrix B is nonnegative definite, then the bias $\tilde{\Delta}$ is nonnegative, making $\tilde{V}(\hat{\tau})$ a conservative estimator of $\text{var}(\hat{\tau})$. Furthermore, by (23), this bias vanishes if and only if $\bar{\tau}_1 = \cdots = \bar{\tau}_W$, when B has each row sum zero, and is a positive semidefinite (psd) matrix of rank $W - 1$. These facts are summarized in Theorem 4, which is the main result of this section.

Theorem 4. *Let there exist a psd matrix $B = ((b_{ww*}))$ of order W and satisfying the conditions: (c1) $b_{ww} = M_w^2$, $w = 1, \dots, W$, (c2) $\sum_{w*=1}^W b_{ww*} = 0$, $w = 1, \dots, W$, and (c3) $\text{rank}(B) = W - 1$. Then the variance estimator $\tilde{V}(\hat{\tau})$ defined in (22) estimates $\text{var}(\hat{\tau})$ with a nonnegative bias $\tilde{\Delta}$ given by (23), which vanishes if and only if $\bar{\tau}_1 = \cdots = \bar{\tau}_W$.*

Remark 2. Recall that the between-WP additivity condition (13) is equivalent to $\bar{\tau}_1 = \cdots = \bar{\tau}_W$ for every treatment contrast. Thus, even when the WP sizes M_1, \dots, M_W are not all equal, by Theorem 4, the bias $\tilde{\Delta}$ vanishes for every treatment contrast if and only if between-WP additivity holds. Thus, if a psd matrix B satisfying conditions (c1)-(c3) is available, then Theorem 4 provides us with a variance estimator that possesses properties similar to the ones derived by Zhao et al. (2018) for the balanced case. However, the issue of existence of such a matrix turns out to be quite

challenging, and will be explored in the next section.

Remark 3. In Theorem 4, conditions (c1) and (c2) ensure the “if” part, while (c3) accounts for the “only if” part. To see the role of (c3) in some detail, from (23) note that $\tilde{\Delta}$ vanishes, and hence $\tilde{V}(\hat{\tau})$ becomes unbiased for $\text{var}(\hat{\tau})$, whenever $(\bar{\tau}_1, \dots, \bar{\tau}_W) \in \mathcal{R}^{\text{orth}}(B)$, where $\mathcal{R}^{\text{orth}}(B)$ is the orthocomplement of the row space of B . If $\text{rank}(B)$ is allowed to be less than $W-1$ by dropping (c3), then $\mathcal{R}^{\text{orth}}(B)$ is broader than the space of vectors $(\bar{\tau}_1, \dots, \bar{\tau}_W)$ that have all elements equal, and hence $\tilde{V}(\hat{\tau})$ becomes unbiased under a wider variety of situations than $\bar{\tau}_1 = \dots = \bar{\tau}_W$ alone. However, this gain comes with a serious price. Along the lines of Proposition 3 in Section 6, given (c1) and (c2), the largest eigenvalue of B is at least as large as $\sum_{w=1}^W M_w^2 / \text{rank}(B)$, and hence it can get considerably inflated if $\text{rank}(B) < W-1$. As discussed later in Section 6 and illustrated in Example 1, this, in turn, can significantly increase the bias of $\tilde{V}(\hat{\tau})$ when $(\bar{\tau}_1, \dots, \bar{\tau}_W) \notin \mathcal{R}^{\text{orth}}(B)$. This is why we retain condition (c3) which not only ensures the “only if” part of Theorem 4, but also helps in controlling the bias when the condition $\bar{\tau}_1 = \dots = \bar{\tau}_W$ does not hold.

5. Existence and construction

We will now study the existence of a psd matrix B satisfying conditions (c1)-(c3) stated in Theorem 4 as a purely mathematical problem. Without loss of generality, we assume hereafter that

$$M_1 \leq M_2 \leq \cdots \leq M_W. \quad (24)$$

To motivate the ideas, consider first the case $W = 3$, where conditions (c1) and (c2) determine B uniquely as

$$B = \begin{bmatrix} M_1^2 & (M_3^2 - M_1^2 - M_2^2)/2 & (M_2^2 - M_1^2 - M_3^2)/2 \\ (M_3^2 - M_2^2 - M_1^2)/2 & M_2^2 & (M_1^2 - M_2^2 - M_3^2)/2 \\ (M_2^2 - M_3^2 - M_1^2)/2 & (M_1^2 - M_3^2 - M_2^2)/2 & M_3^2 \end{bmatrix}. \quad (25)$$

This matrix is also psd and satisfies (c3) if and only if its principal minor, given by the first two rows and columns, is positive. Simplification of this condition and application of (24) yields $M_3 < M_1 + M_2$ as the necessary and sufficient condition for B to satisfy (c1)-(c3). This construction of B for $W = 3$ raises the following questions with respect to the general case $W \geq 3$:

(a) Is the condition

$$M_W < M_1 + \cdots + M_{W-1}, \quad (26)$$

necessary and sufficient for existence of a psd matrix B satisfying (c1)-(c3)?

(b) If so, then under (26), can one construct such a matrix B by an extension of the form in (25) to the general case?

Later in this section, Theorem 5 answers (a) in the affirmative. On the other hand, the question in (b) does not allow a conclusive answer. To see why, observe that the most obvious extension of (25) to general $W \geq 3$ is given by $B = ((b_{ww^*}))$, with

$$\begin{aligned} b_{ww} &= M_w^2, w = 1, \dots, W, \\ b_{ww^*} &= \frac{M_1^2 + \dots + M_W^2}{(W-1)(W-2)} - \frac{M_w^2 + M_{w^*}^2}{W-2}, w, w^* = 1, \dots, W, w \neq w^*. \end{aligned} \quad (27)$$

The divisors in (27) ensure condition (c2) about zero row sums and make it consistent with (25) when $W = 3$. The form (27) is also natural because, in keeping with M_1^2, \dots, M_W^2 as the diagonal elements of B , it takes the off-diagonal elements as linear combinations of M_1^2, \dots, M_W^2 in a systematic manner. However, unlike the case of $W = 3$, the matrix B given by (27) may not be psd for $W \geq 4$, even when the condition (26) holds. For instance, if $W = 4$, then this condition holds for both the configurations $(M_1, \dots, M_4) = (8, 8, 12, 12)$ and $(6, 6, 14, 14)$. The matrix B in (27) is psd of rank 3 ($= W - 1$) for the first configuration, but has a negative eigenvalue

for the second.

The above discussion makes it clear that, in general, the task of obtaining a psd matrix B satisfying (c1)-(c3) under condition (26) can be far more complex than what the form (25) arising for $W = 3$ suggests. Theorem 5 establishes condition (26) as a necessary and sufficient condition for existence of such a matrix.

Theorem 5. *Let $W \geq 3$. Then condition (26), that is, $M_W < M_1 + \dots + M_{W-1}$, is necessary and sufficient for the existence of a psd matrix $B = ((b_{ww^*}))$ of order W and satisfying the conditions (c1) $b_{ww} = M_w^2$, $w = 1, \dots, W$, (c2) $\sum_{w^*=1}^W b_{ww^*} = 0$, $w = 1, \dots, W$, and (c3) $\text{rank}(B) = W - 1$.*

The sufficiency part of the proof of Theorem 5 leads to a construction procedure of the matrix B satisfying conditions (c1)-(c3). If $M_1 = \dots = M_W (= M, \text{ say})$, then one can simply take M^2 at each diagonal position of B and $-M^2/(W - 1)$ at each off-diagonal position. Turning next to the case of unequal $M_1 \leq \dots \leq M_W$, suppose condition (26) holds. Let $\mu = (M_1, \dots, M_{W-1})'$, where the prime denotes transposition, and let e denote the $(W - 1) \times 1$ vector of ones. Then the steps involved in the construction of the matrix B are:

Step 1: Find a vector x with elements ± 1 satisfying the condition

$$|\mu'x| < M_W. \quad (28)$$

Step 2: Find nonnegative constants a_1 and a_2 , satisfying $a_1 + a_2 < 1$ such that

$$a_1 \left\{ (\mu'x)^2 - \mu'\mu \right\} + a_2 \left\{ (\mu'e)^2 - \mu'\mu \right\} = M_W^2 - \mu'\mu. \quad (29)$$

Step 3: Construct the matrix $A = D \{a_1 xx' + a_2 ee' + (1 - a_1 - a_2)I\} D$,

where x , a_1 and a_2 are obtained from steps 1 and 2 above, I is the identity matrix of order $W - 1$ and $D = \text{diag}(M_1, \dots, M_{W-1})$.

Step 4: Construct matrix B as follows:

$$B = \begin{bmatrix} A & -Ae \\ -e'A & e'Ae \end{bmatrix}$$

Then B is psd of order W and satisfies (c1)-(c3) by the proof of the sufficiency part of Theorem 5. A lemma, crucial in this proof, appears in the supplementary materials and guarantees the existence of vector x in step 1 and constants a_1 and a_2 in step 2 under condition (26).

Remark 4. It is satisfying that the condition (26) holds under wide generality. It only requires the largest WP to be not too large compared to the

others and holds, in particular, when there is a tie about the largest WP. Moreover, in some situations it may be possible to adjust the composition of the WPs so as to meet (26) without significantly increasing the cost of the experiment. As an illustration, consider a variant of the school example of Section 1, where the 40 schools are spread over four counties with 8, 6, 6 and 20 schools, so that (26) does not hold. In this case, if the third and fourth counties are contagious and two schools of the fourth county are close to their border, then these two may be clubbed with the third country, leading to WPs of sizes 8, 6, 8 and 18, ensuring (26) is met.

Remark 5. For $W = 3$, one can check that the construction stated above yields the unique B in (25). For $W \geq 4$, however, a psd matrix B meeting (c1)-(c3) is non-unique. Indeed, then the above construction itself can yield a wide class of such matrices B considering all vectors x which satisfy (28), and for each such x , all nonnegative a_1, a_2 satisfying $a_1 + a_2 < 1$ and (29). Thus, the issue of discriminating among rival choices of B becomes important. Such a discriminating strategy is discussed in Section 6.

6. Minimax estimators unbiased under between-WP additivity

As seen in Section 5, while condition (26) guarantees the existence of matrix B and consequently a variance estimator that is unbiased under between-WP additivity, such a matrix is non-unique. Thus, it is important to define a criterion that can discriminate among possible choices of B . Clearly, a good choice should control the bias $\tilde{\Delta} = (1/N^2) \sum_{w=1}^W \sum_{w^*=1}^W b_{ww^*} \bar{\tau}_w \bar{\tau}_{w^*}$ given by (23) that is associated with the estimation of $\text{var}(\widehat{\bar{\tau}})$. The hurdle here is that $\bar{\tau}_1, \dots, \bar{\tau}_W$ are unknown. Even the idea of minimaxity does not work without further refinement, because B is psd, and hence $\tilde{\Delta}$ is unbounded with respect to variation of $\bar{\tau}_1, \dots, \bar{\tau}_W$ in the W -dimensional real space. On the other hand, by (16), multiplication of $\bar{\tau}_1, \dots, \bar{\tau}_W$ by any nonzero constant only rescales the treatment contrast $\bar{\tau}$, without essentially altering it. We, therefore, consider minimization of $\tilde{\Delta}$ subject to $\sum_{w=1}^W \bar{\tau}_w^2 = 1$. This is motivated by Mukerjee et al. (2018) who touched upon split-plot designs only in the balanced case. It is easy to see that the above formulation calls for obtaining B , subject to (c1)-(c3), so as to minimize $\lambda_{\max}(B)$, the largest eigenvalue of B . The following proposition provides us with a lower bound for $\lambda_{\max}(B)$.

Proposition 3. *For any psd matrix B satisfying (c1)-(c3), a lower bound*

for $\lambda_{\max}(B)$ is given by $\lambda_0 = \sum_{w=1}^W M_w^2 / (W-1)$, but this bound is not sharp whenever M_1, \dots, M_W are not all equal.

Given Proposition 3, an analytical solution to the minimaxity problem above seems to be intractable in the unbalanced case. This is anticipated, because a complete characterization of matrices B satisfying (c1)-(c3) is hard, even though in Section 5, we were able to outline a general method for constructing such matrices when condition (26) holds. As a practical strategy, therefore, it makes sense to concentrate on matrices B that can be obtained via this method, with a view to minimizing $\lambda_{\max}(B)$ among these matrices. It is reassuring that even then the class of competing matrices B is quite large, as noted in Remark 5. In practice, these competing matrices B can be generated quite fast by (i) enumerating vectors x with elements ± 1 and satisfying (28), and (ii) for each such x , performing a grid search to find nonnegative (a_1, a_2) satisfying $a_1 + a_2 < 1$ and (29). To implement (ii), we vary a_1 from 0 to 0.9999 in steps of 0.0001, such that the corresponding a_2 , found from (29), meets $0 \leq a_2 < 1 - a_1$. Finding the matrix B that has the smallest λ_{\max} among all candidate matrices identified through (i) and (ii) above is fast and easy. This enumeration becomes even faster noting that, without loss of generality, the first element of x in (i) can be taken as

+1, because replacing x by $-x$ does not affect Steps 1-4 of Section 5.

Example 1. Returning to the school example in Sections 1 and 2, where we have $N = 40$, $W = 4$ and $(M_1, M_2, M_3, M_4) = (8, 8, 12, 12)$, the smallest $\lambda_{\max}(B)$ obtainable via steps 1 through 4 described in Section 5 is 192, which corresponds to

$$B = \begin{bmatrix} 64 & 32 & -48 & -48 \\ 32 & 64 & -48 & -48 \\ -48 & -48 & 144 & -48 \\ -48 & -48 & -48 & 144 \end{bmatrix},$$

as given by $x = (1, 1, -1)'$, $a_1 = 0.5$ and $a_2 = 0$. Clearly, this B meets (c1)-(c3) of Theorem 4, and quite reassuringly, it has the smallest possible $\lambda_{\max}(B)$ also over all psd matrices B satisfying (c1) and (c2), as one can verify numerically. On the other hand, if B satisfies only (c1) and (c2) but not (c3), i.e., $\text{rank}(B) < W-1 (= 3)$, then following Remark 3, $\lambda_{\max}(B)$ gets inflated. For example, if $\text{rank}(B) = 2$, then $\lambda_{\max}(B) \geq \sum_{w=1}^W M_w^2 / 2 = 208$. Further, if $\text{rank}(B) = 1$, then by (c1) and (c2), $B = qq'$ where q equals $(8, -8, -12, 12)'$ or $(8, -8, 12, -12)'$ and $\lambda_{\max}(B)$ becomes as large as 416.

7. Simulation results and performance comparisons

Whereas Theorem 4 establishes unbiasedness of $\tilde{V}(\hat{\tau})$ under (26) and between-WP additivity, and consideration of minimaxity is expected to provide protection under extreme departures from additivity, it is also important to understand how the bias of $\tilde{V}(\hat{\tau})$ would compare to that of $\hat{V}(\hat{\tau})$ under different levels of treatment effect heterogeneity. Also, the theoretical results do not provide any clue regarding whether the bias adjustment comes at the price of an undesirable inflation of the MSE. We now conduct some simulations to study these two aspects. We consider the estimation of the main effect of factor F_2 in the setting of Example 1. The unit-level treatment contrast τ_i equals $\{-Y_i(00) + Y_i(01) - Y_i(10) + Y_i(11)\}/2$ for $i = 1, \dots, 40$ (Dasgupta et al., 2015). The finite population contrast of interest is $\bar{\tau} = \sum_{i=1}^{40} \tau_i/40$. The vector of potential outcomes for unit i , denoted by $Y_i = (Y_i(00), Y_i(01), Y_i(10), Y_i(11))$, is generated using the multivariate normal model $Y_i \sim N_4(\theta_w, \Sigma_w)$, $i \in \Omega_w$, $w = 1, \dots, 4$, where $\Sigma_w = \sigma_w^2 \{(1 - \rho_w)I_4 + \rho_w J_4\}$ is the covariance matrix for WP Ω_w that depends on two parameters: the variance σ_w^2 and correlation ρ_w . Matrices I_n and J_n respectively denote the n th order identity matrix and the matrix of ones. Seven possible scenarios (listed in Table 2) for generating the

potential outcomes are considered.

Table 2: Simulation settings

Population	θ_1	θ_2	θ_3	θ_4	σ_1^2	σ_2^2	σ_3^2	σ_4^2	ρ_1	ρ_2	ρ_3	ρ_4
I	(5,8,7,8)	(7,9,4,6)	(8,11,7,8)	(7,8,6,9)	2.5	2	2	3	.5	.5	.5	.5
II	(10,5,9,8)	(5,9,10,8)	(10,9,8,5)	(10,5,8,9)	2.5	2	2	3	1	1	1	1
III	(10,5,9,8)	(5,9,10,8)	(10,9,8,5)	(10,5,8,9)	2.5	2	2	3	.5	.5	.5	.5
IV	(10,5,9,8)	(5,9,10,8)	(10,9,8,5)	(10,5,8,9)	2.5	2	2	3	.2	.4	.6	.8
V	(10,5,9,8)	(5,9,10,8)	(10,9,8,5)	(10,5,8,9)	2.5	2	2	3	0	0	0	0
VI	(10,5,9,8)	(5,9,10,8)	(10,9,8,5)	(10,5,8,9)	2.5	2	2	3	-.3	-.3	-.3	-.3
VII	(10,5,9,8)	(5,9,10,8)	(10,9,8,5)	(10,5,8,9)	2.5	2	2	3	-.3	.3	-.3	.3

The potential outcomes for population I are forced to ensure, via an appropriate command in R, that the WP means $\bar{\tau}_1, \dots, \bar{\tau}_4$ are always two. Population II generates different $\bar{\tau}_1, \dots, \bar{\tau}_4$ but guarantees the same τ_i within each WP. Populations III through VII differ only with respect to the correlation parameters that lead to different types of treatment effect heterogeneity. These include all zero correlations in population V, all negative correlations in population VI, and a mix of positive and negative correlations in population VII.

From each population, 25 sets of potential outcomes are generated, and the biases of variance estimators $\hat{V}(\hat{\tau})$ and $\tilde{V}(\hat{\tau})$ are compared. Note that these biases are Δ given by (17) and $\tilde{\Delta}$ given by (23). Because we do not have any theoretical expressions for the MSEs of $\hat{V}(\hat{\tau})$ and $\tilde{V}(\hat{\tau})$, we estimate them empirically. For each set of potential outcomes, we generate

10000 treatment assignments and thus 10000 sets of observed outcomes, calculate $\hat{V}(\hat{\tau})$ and $\tilde{V}(\hat{\tau})$ for each set, and estimate their MSEs from these 10000 values. The results are summarized in Table 3.

Table 3: Median Bias and Median MSE of the variance estimators from 25 sets of potential outcomes under seven settings

Population	$\hat{V}(\hat{\tau})$		$\tilde{V}(\hat{\tau})$	
	Bias	MSE	Bias	MSE
I	0.0533	0.2317	0.0000	0.2613
II	0.6400	5.9011	0.2800	5.8343
III	0.6403	5.7001	0.2740	5.5989
IV	0.7096	5.5829	0.3306	5.8072
V	0.7284	6.1943	0.3285	5.8122
VI	0.7667	5.2141	0.3719	5.3680
VII	0.5921	5.2011	0.2915	5.5709

The results suggest that with respect to bias, the variance estimator $\tilde{V}(\hat{\tau})$ outperforms the estimator $\hat{V}(\hat{\tau})$ uniformly across all the seven populations. As expected, the bias of $\tilde{V}(\hat{\tau})$ is zero for population I. However, with respect to MSE, there appears to be no clear winner. The estimator $\tilde{V}(\hat{\tau})$ performs slightly better than $\hat{V}(\hat{\tau})$ in populations II, III and V and slightly worse in the other populations with respect to MSE. However, it is

encouraging to note that the bias reduction associated the estimator $\tilde{V}(\hat{\tau})$ does not appear to come with a significant sacrifice of MSE.

8. Discussion

We have seen that the attempt to generalize Neymanian variance estimation from balanced to unbalanced split-plot designs is a highly non-trivial problem. Results from unbalanced block designs cannot be used in such a setting. At the expense of an intricate chain of arguments, it is possible to derive a variance estimator that is unbiased under assumptions of treatment homogeneity similar to that in the balanced case, and is also robust under departures from such an assumption. Empirical evidences suggest that this bias reduction does not come at a price of undesirably high inflation of MSE. It will be interesting to examine the performance of the two estimators considered in this paper with respect to coverage of asymptotic confidence intervals for the true contrast τ . Such a comparison will, however, entail considerable theoretical work to establish the asymptotic distribution of $\hat{\tau}$ because of the complex setting of the unbalanced split-plot assignment mechanism. We plan to pursue this as our future research.

Supplementary Materials

The supplementary materials contain the proofs of all results stated in the main manuscript.

Acknowledgements

This work was supported by the J.C. Bose National Fellowship, Government of India, and grants from Indian Institute of Management Calcutta and National Science Foundation, USA. We are grateful to two reviewers whose insightful comments resulted in significant improvement of the paper.

References

- Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. M. (2020). Sampling-based versus design-based uncertainty in regression analysis. *Econometrica*, 88:265–296.
- Box, G. E. P., Hunter, J. S., and Hunter, W. G. (2005). *Statistics for Experimenters: Design, Innovation, and Discovery*. John Wiley & Sons, Hoboken, New Jersey, 2nd edition.
- Cochran, W. G. (1977). *Sampling Techniques*. John Wiley & Sons: New York.

REFERENCES

- Cochran, W. G. and Cox, G. M. (1957). *Experimental Designs*. John Wiley & Sons, Hoboken, New Jersey, 2nd edition.
- Dasgupta, T., Pillai, N. S., and Rubin, D. B. (2015). Causal inference for 2^K factorial designs by using potential outcomes. *Journal of the Royal Statistical Society, Series B*, 77:727–753.
- Fisher, R. A. (1925). *Statistical Methods for Research Workers*. Oliver & Boyd, Edinburgh, Scotland.
- Fisher, R. A. (1935). *The Design of Experiments*. Oliver & Boyd, Oxford, England, 1st edition.
- Freedman, D. A. (2006). Statistical models for causation: What inferential leverage do they provide? *Evaluation Review*, 30:691–713.
- Freedman, D. A. (2008). On regression adjustments to experimental data. *Advances in Applied Mathematics*, 40:180–193.
- Kirk, R. E. (1982). *Experimental Design: Procedures for the Behavioral Sciences*. Brooks/Cole, Monterey, CA.
- Mukerjee, R., Dasgupta, T., and Rubin, D. B. (2018). Using standard tools from finite population sampling to improve causal inference for complex

REFERENCES

- experiments. *Journal of the American Statistical Association*, 113:868–881.
- Neyman, J. (1923). On the application of probability theory to agricultural experiments. Essay on principles. Section 9. *Statistical Science*, 5:465–472.
- Olsen, R., Orr, L., Bell, S., and Stuart, E. (2013). External validity in policy evaluations that choose sites purposively. *Journal of Policy Analysis and Management*, 32:107–121.
- Pashley, N. and Miratrix, L. (2019). Insights on variance estimation for blocked and matched pairs designs. *arXiv:1710.10342v5 [stat.ME]*.
- Rosenberger, W. F., Uschner, D., and Wang, Y. (2019). Randomization: The forgotten component of the randomized clinical trial. *Statistics in medicine*, 38:1–12.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66:688–701.
- Rubin, D. B. (1980). Comment on “Randomization analysis of experimen-

REFERENCES

- tal data: The Fisher randomization test by D. Basu". *Journal of the American Statistical Association*, 75:591–593.
- Schochet, P. S., Pashley, N., Miratrix, L., and Kautz, T. (2020). Design-based ratio estimators for clustered, blocked rcts. *arXiv:2002.01146v1 [stat.ME]*.
- Wu, C. F. J. and Hamada, M. S. (2009). *Experiments: Planning, Analysis, and Optimization*. John Wiley & Sons, Hoboken, New Jersey, 2nd edition.
- Yates, F. (1935). Complex experiments. *Supplement to the Journal of the Royal Statistical Society*, 2:181–247.
- Zhao, A., Ding, P., Mukerjee, R., and Dasgupta, T. (2018). Randomization-based causal inference from split-plot designs. *Annals of Statistics*, 46:1876–1903.
- Indian Institute of Management Calcutta, Joka, Diamond Harbour Road, Kolkata 700104, India
- E-mail: rmuk0902@gmail.com
- Department of Statistics, Rutgers University, 110 Frelinghuysen Road, Piscataway, New Jersey 08901, U.S.A.
- E-mail: tirthankar.dasgupta@rutgers.edu