

Covariate adjustment in stratified experiments

MAX CYTRYNBAUM

Department of Economics, Yale University

This paper studies covariate adjusted estimation of the average treatment effect in stratified experiments. We work in a general framework that includes matched tuples designs, coarse stratification, and complete randomization as special cases. Regression adjustment with treatment-covariate interactions is known to weakly improve efficiency for completely randomized designs. By contrast, we show that for stratified designs such regression estimators are generically inefficient, potentially even increasing estimator variance relative to the unadjusted benchmark. Motivated by this result, we derive the asymptotically optimal linear covariate adjustment for a given stratification. We construct several feasible estimators that implement this efficient adjustment in large samples. In the special case of matched pairs, for example, the regression including treatment, covariates, and pair fixed effects is asymptotically optimal. We also provide novel asymptotically exact inference methods that allow researchers to report smaller confidence intervals, fully reflecting the efficiency gains from both stratification and adjustment. Simulations and an empirical application demonstrate the value of our proposed methods.

KEYWORDS. Matched pairs, analysis of covariance, blocking, robust standard error, treatment effects.

JEL CLASSIFICATION. C10, C14, C90.

1. INTRODUCTION

This paper studies covariate adjusted estimation of the average treatment effect (ATE) in stratified experiments. Researchers often make use of both stratified treatment assignment and ex post covariate adjustment to improve the precision of experimental estimates. Indeed, out of a survey of over 50 experimental papers published in the AER and AEJ between 2018–2023, we found that 57% use stratified randomization, and 80% used some form of ex post covariate adjustment. An influential paper by [Lin \(2013\)](#) showed in a design-based setting that the regression estimator with full treatment-covariate interactions is always asymptotically weakly more efficient than difference of means estimation for completely randomized designs. [Negi and Wooldridge \(2021\)](#) extended these results to estimation of the ATE using data sampled from a superpopulation. However, questions remain about the interaction between stratification and regression adjustment and the implications of combining these methods for both estimator efficiency and the power and validity of inference methods. To study these questions, we work in

Max Cytrynbaum: max.cytrynbaum@yale.edu

I thank the anonymous referees for helpful suggestions during the revision process.

the stratified randomization framework of [Cytrynbaum \(2023\)](#), which includes matched tuples designs (e.g., matched pairs), coarse stratification, and complete randomization as special cases.

We show that the [Lin \(2013\)](#) interacted regression adjustment is generically inefficient in the family of linearly adjusted estimators, with asymptotic efficiency only in the limiting case of complete randomization. Motivated by this finding, we characterize the efficient linear covariate adjustment for a given stratified design, providing several new estimators that achieve the optimal variance.

Our first result derives the optimal linear adjustment coefficient for a given stratification. We show that asymptotically the interacted regression estimator uses the wrong objective function, minimizing a marginal variance objective that is totally insensitive to the stratification. By contrast, the optimal adjustment coefficient minimizes a *mean-conditional* variance objective, conditional on the covariates used to stratify. Intuitively, the efficient covariate adjustment is tailored to the stratification, ignoring fluctuations of the estimator that are predictable by the stratification covariates. Section 3.2 draws an interesting connection with partially linear regression ([Robinson \(1988\)](#)), showing that efficient linear adjustment of a stratified design is asymptotically equivalent to doubly-robust semiparametric adjustment of an i.i.d. design. Intuitively, stratification contributes the nonparametric component of the semiparametric adjustment function.

Our second set of results develops feasible versions of the optimal linear adjustment derived in Section 3.1. First, we show that if the conditional expectation of the adjustment covariates is linear in a known set of transformations of the stratification variables, then adding the latter to the interacted regression restores optimality. Next, we relax this assumption, providing four different regression estimators that are asymptotically efficient under weak conditions. For matched pairs experiments or in settings with limited treatment effect heterogeneity, the noninteracted regression with a full set of pair fixed effects is asymptotically efficient. More generally, we show asymptotic optimality of within-stratum (inconsistently) partialled versions of the Lin and tyranny-of-the-minority estimators ([Lin \(2013\)](#)). We also define a “group OLS” estimator, extending a proposal of [Imbens and Rubin \(2015\)](#) for matched pairs experiments to a larger class of designs. We show that this group OLS estimator is also asymptotically optimal.

Our final contribution is to develop novel asymptotically exact inference methods for covariate adjusted estimation under stratified designs. Confidence intervals based on the usual heteroskedasticity robust variance estimator are known to be conservative in stratified experiments ([Bai, Romano, and Shaikh \(2021\)](#)). By contrast, the coverage probabilities of our proposed confidence intervals converge to the specified nominal level, with no overcoverage. Our approach applies to a generic family of linear covariate adjustments and randomization schemes, including as special cases noninteracted regression adjustment, the [Lin \(2013\)](#) interacted regression, and all of the other estimators considered in this paper. Simulations and an empirical application to the experiment in [Baysan \(2022a\)](#) suggest that the usual robust confidence intervals can substantially overcover in stratified experiments, while our confidence intervals have close to nominal coverage.

We present several extensions of our main results in the Supplemental Appendix (Cytrynbaum (2024)). In the first, we consider estimation and inference in stratified experiments with noncompliance. As a simple corollary of our results on ATE estimation, we characterize the optimal linearly adjusted Wald estimator for the LATE (Imbens and Angrist (1994)), construct a feasible implementation of the efficient adjustment, and provide asymptotically exact inference methods. We also study efficient linear adjustment for finely stratified designs with nonconstant treatment proportions, as in Cytrynbaum (2023), and briefly consider the problem of efficient nonlinear adjustment.

There has been significant interest in treatment effect estimation under different experimental designs in the recent literature. Some papers studying covariate adjustment under stratified randomization include Bugni, Canay, and Shaikh (2018), Fogarty (2018), Liu and Yang (2020), Lu and Liu (2024), Wei, Tu, and Liu (2022), Reluga, Ye, and Zhao (2024), Wang, Wang, and Liu (2021), Ting, Shao, Yi, and Zhao (2022), Ke, Liu, and Yang (2024), and Chang (2023). These works differ from our paper in at least one of the following ways: (1) studying inference on the sample average treatment effect (SATE) rather than the ATE in a superpopulation, (2) restricting to coarse stratification (stratum size going to infinity), or (3) proving weak efficiency gains but not optimality. In a finite population setting, Ke, Liu, and Yang (2024) shows asymptotic efficiency of a projection-based estimator numerically equivalent to the “partialled Lin” approach considered in Section 3.4.2. In the same setting, Lu and Liu (2024) prove efficiency of a tyranny-of-the-minority style regression similar but not equivalent to one the considered in Section 3.4.4. Both papers give conservative inference on the SATE, while we provide asymptotically exact inference on the ATE using a generalized pairs-of-pairs (Abadie and Imbens (2008)) style approach. Remarks 3.19 and 3.22 in Section 3.4 below provide a detailed comparison.

Relative to the above papers, the superpopulation framework considered here creates some new technical challenges. For example, as pointed out in Bai, Romano, and Shaikh (2021), matching units into data-dependent strata post-sampling produces a complicated dependence structure between the treatment assignments and random covariates. We deal with this using a tight-matching condition (Equation (2.1)) and martingale CLT analysis similar to Cytrynbaum (2022). This setting also has analytical advantages, which allow us to establish new conceptual results. For example, the population level characterization of the optimal adjustment coefficient in Section 3.1 allows us to give explicit necessary and sufficient conditions for the efficiency of several commonly used regression estimators. The efficiency of interacted regression under a “rich covariates” condition, as well as the equivalence between optimal linear adjustment of stratified designs and doubly-robust semiparametric adjustment appear to be new observations in this literature. To the best of our knowledge, we give the first asymptotically exact inference on the ATE for general covariate adjusted estimators under finely stratified randomization.

Independently, Bai, Jiang, Romano, Shaikh, and Zhang (2024b) study covariate adjustment under matched pairs randomization in a superpopulation framework. They also find that regression adjustment without pair fixed effects may be inefficient, while adding pair fixed effects restores efficiency. Relative to our work, they additionally study

regularized regression adjustment under high-dimensional asymptotics, which we do not consider. By contrast, we study more general forms of stratification, allowing coarse and fine stratification with arbitrary treatment proportions $p \neq 1/2$. For such designs, the strata fixed-effects estimator may still be inefficient. To fix this, we introduce novel forms of linear adjustment that are efficient under general stratified designs.

The rest of the paper is organized as follows. In Section 2, we define notation and introduce the family of stratified designs that we will consider throughout the paper. Section 3 gives our main results, characterizing optimal covariate adjustment, and constructing efficient estimators. Section 4 provides asymptotically exact inference on the ATE for generic linearly adjusted estimators. In Sections 5 and 6, we study the finite-sample properties of our method, including both simulations and an empirical application to the experiment in Baysan (2022a). Section 7 concludes with some recommendations for practitioners.

2. FRAMEWORK AND STRATIFIED DESIGNS

For a binary treatment $d \in \{0, 1\}$, let $Y_i(1), Y_i(0)$ denote the treated and control potential outcomes, respectively. For treatment assignment D_i , let $Y_i = Y_i(D_i) = D_i Y_i(1) + (1 - D_i) Y_i(0)$ be the observed outcome. Let X_i denote covariates. Consider data $(X_i, Y_i(1), Y_i(0))_{i=1}^n$ sampled i.i.d. from a superpopulation of interest. We are interested in estimating the average treatment effect in this population, $ATE = E[Y(1) - Y(0)]$. After sampling units $i = 1, \dots, n$, treatments $D_{1:n}$ are assigned by stratified randomization. In particular, we use the “local randomization” framework introduced in Cytrynbaum (2022).

DEFINITION 2.1 (Local randomization). Let treatment proportions $p = a/k$ with $\gcd(a, k) = 1$.¹ Suppose that n is divisible by k for notational simplicity. Partition the experimental units into n/k groups g with $\{1, \dots, n\} = \bigcup_g g$ disjointly and $|g| = k$. Let $\psi(X) \in \mathbb{R}^{d_\psi}$ denote a vector of stratification variables. Suppose that the groups that satisfy a homogeneity condition with respect to $\psi(X)$ such that

$$\frac{1}{n} \sum_g \sum_{i, j \in g} |\psi(X_i) - \psi(X_j)|_2^2 = o_p(1). \tag{2.1}$$

Require that the groups only depend on the stratification variables $\psi_{1:n}$ and data-independent randomness π_n , so that $g = g(\psi_{1:n}, \pi_n)$ for each g . Independently, for each $|g| = k$, draw treatment variables $(D_i)_{i \in g}$ by setting $D_i = 1$ for exactly a out of k units, completely at random. For a stratification satisfying these conditions, we denote $D_{1:n} \sim \text{Loc}(\psi, p)$.

EXAMPLE 2.2 (Matched tuples). Equation (2.1) requires units in a group to have similar $\psi(X_i)$ values and can be thought of as a tight-matching condition. Cytrynbaum (2023) provides an iterative pairing algorithm to match units into groups that provably satisfy

¹ $\gcd(a, k)$ stands for greatest common divisor.

this condition for any k . Drawing treatments $D_{1:n} \sim \text{Loc}(\psi, p)$ produces a “matched tuples” design for $p = a/k$. Matched pairs corresponds to the case $p = 1/2$.

EXAMPLE 2.3 (Complete randomization). We say variables $D_{1:n}$ are completely randomized with treatment probability p if $D_{1:n}$ is drawn uniformly from all vectors $d_{1:n}$ with $d_i = 1$ for exactly proportion p of the units. Formally, $P(D_{1:n} = d_{1:n}) = 1/\binom{n}{pn}$ for all such vectors. We denote complete randomization by $D_{1:n} \sim \text{CR}(p)$. Complete randomization may be obtained in our framework by setting $\psi = 1$ and forming groups $|g| = k$ at random, which automatically satisfies Equation (2.1). For example, assigning 2 out of 3 units in each group to treatment gives a “random matched triples” representation of complete randomization with $p = 2/3$.

REMARK 2.4 (Coarse stratification). Similarly, coarse stratification with large fixed strata $S(X) \in \{1, \dots, m\}$ can also be obtained in our framework by setting $\psi(X) = S(X)$ and matching units with identical $S(X)$ values into groups at random. Because of this, our framework enables a unified asymptotic analysis for a wide range of stratifications.

Experiment timing Suppose that the experimenter does the following:

- (1) Samples units and observes their baseline covariates.
- (2) Partitions the units into data-dependent groups $g = g(\psi_{1:n}, \pi_n)$ that satisfy Equation (2.1) for some stratification variables $\psi(X)$.
- (3) Draws treatment assignments $D_{1:n} \sim \text{Loc}(\psi, p)$, observes outcomes $Y_i(D_i)$, and forms an estimate of the ATE, potentially adjusting for covariates $h(X)$.

We are agnostic about the exact time at which the covariates are observed, subject to the constraints above. For example, it could be that only $\psi(X)$ is observed at the design stage, while the full vector X is collected later with the outcomes, and the experimenter chooses to adjust for $h(X) \subseteq X$. Alternatively, the full vector X could be observed at the design stage, but the experimenter chooses to only stratify on $\psi(X)$, and adjusts for $h(X) \subseteq X$ at step (3). We may or may not have $\psi(X) \subseteq h(X)$.²

Consider the unadjusted estimator given by the coefficient $\hat{\theta}$ on D in the regression $Y \sim 1 + D$. Before discussing covariate adjustment, we first state a helpful variance decomposition for $\hat{\theta}$ that will be used extensively below. Let $c(X) = E[Y(1) - Y(0)|X]$ denote the conditional average treatment effect (CATE) and $\sigma_d^2(X) = \text{Var}(Y(d)|X)$ the heteroskedasticity function. Define the *balance function*:

$$b(X; p) = E[Y(1)|X] \left(\frac{1-p}{p}\right)^{1/2} + E[Y(0)|X] \left(\frac{p}{1-p}\right)^{1/2}. \tag{2.2}$$

We often denote $b = b(X; p)$ in what follows. [Cytrynbaum \(2022\)](#) shows that if $D_{1:n} \sim \text{Loc}(\psi, p)$ then $\sqrt{n}(\hat{\theta} - \text{ATE}) \Rightarrow \mathcal{N}(0, V)$ with

$$V = \text{Var}(c(X)) + E[\text{Var}(b|\psi)] + E\left[\frac{\sigma_1^2(X)}{p} + \frac{\sigma_0^2(X)}{1-p}\right]. \tag{2.3}$$

²Our asymptotic framework lets $h(X), \psi(X)$ be fixed as $n \rightarrow \infty$.

The variance V is in fact the [Hahn \(1998\)](#) semiparametric variance bound³ for the ATE (with covariates $\psi(X)$), providing a formal sense in which stratification does nonparametric regression adjustment “by design.” The middle term is the most important for our analysis below. For example, in this notation the difference in asymptotic efficiency between stratifications ψ_1 and ψ_2 (for fixed p) is simply $E[\text{Var}(b|\psi_1)] - E[\text{Var}(b|\psi_2)]$. Note also that $E[\text{Var}(b|\psi)] \leq \text{Var}(b)$ for any ψ , showing how stratification removes the variance due to fluctuations that are predictable by $\psi(X)$.

Moving beyond the difference of means estimator $\hat{\theta}$, suppose that at the analysis stage, the experimenter has access to covariates $h(X)$, which may strictly contain $\psi(X)$. One may try to further improve the efficiency of ATE estimation by regression adjustment using these covariates, either using the standard regression $Y \sim 1 + D + h$ or the regression $Y \sim 1 + D + h + Dh$ (with demeaned covariates) studied in [Lin \(2013\)](#). We study the interaction between covariate adjustment and stratification in Section 3.1 below, characterizing the optimal linear adjustment.

3. MAIN RESULTS

3.1 *Efficient linear adjustment in stratified experiments*

In this section, we begin by studying the efficiency of commonly used covariate-adjusted estimators of the ATE under stratified randomization. [Lin \(2013\)](#) showed that in a completely randomized experiment, equivalent to $D_{1:n} \sim \text{Loc}(\psi, p)$ with $\psi = 1$, regression adjustment with full treatment-covariate interactions is asymptotically weakly more efficient than difference of means estimation. [Negi and Wooldridge \(2021\)](#) extended this result to ATE estimation in the superpopulation framework that we use in this paper. Interestingly, we show that this result is atypical. For a general stratified experiment with $\psi \neq 1$, [Lin \(2013\)](#) style regression adjustment may be strictly inefficient relative to difference of means. The issue is that the interacted regression solves the wrong optimization problem, minimizing a marginal variance objective when, due to the stratification, it should instead minimize a mean-conditional variance objective, conditional on the stratification variables ψ . In fact, the Lin estimator is totally insensitive to the stratification, estimating the same adjustment coefficient for any stratified design $D_{1:n} \sim \text{Loc}(\psi, p)$. Because of this, interacted regression is generically suboptimal and in some cases can even be strictly less efficient than difference of means. Before proceeding, we state our main assumption.

ASSUMPTION 3.1 (Smoothness and moment conditions). *Assume the following:*

- (i) *The conditional expectations $E[h(X)|\psi]$ and $E[Y(d)|\psi]$ for $d \in \{0, 1\}$ are Lipschitz continuous in the stratification variables ψ .*
- (ii) *The moments $E[Y(d)^4] < \infty$ for $d \in \{0, 1\}$ and $E[|h_t(X)|^4] < \infty$ for all $1 \leq t \leq \dim(h)$, $|\psi(X)|_2 < K < \infty$ a.s. and $\text{Var}(h) > 0$.*

³[Armstrong \(2022\)](#) shows that this variance bound also holds for stratified designs.

Now we are ready to define the Lin estimator and state our first result. Denote $h_i = h(X_i)$ and demeaned covariates $\tilde{h}_i = h_i - E_n[h_i]$, with $E_n[h_i] \equiv n^{-1} \sum_{i=1}^n h_i$. The Lin estimator $\hat{\theta}_L$ is the coefficient on D_i in the interacted regression

$$Y_i \sim 1 + D_i + \tilde{h}_i + D_i \tilde{h}_i. \tag{3.1}$$

Define the within treatment arm covariate means $\bar{h}_1 = E_n[h_i D_i] / E_n[D_i]$ and $\bar{h}_0 = E_n[h_i(1 - D_i)] / E_n[1 - D_i]$. The Lin estimator $\hat{\theta}_L$ can be related to the difference of means estimator $\hat{\theta}$ as

$$\hat{\theta}_L = \hat{\theta} - \hat{\gamma}'_L (\bar{h}_1 - \bar{h}_0). \tag{3.2}$$

Here, the *adjustment coefficient* $\hat{\gamma}_L$ is $\hat{\gamma}_L = (1 - p)(\hat{a}_1 + \hat{a}_0) + p\hat{a}_0$, where \hat{a}_0 and \hat{a}_1 are the coefficients on \tilde{h}_i and $D_i \tilde{h}_i$ in Equation (3.1). The following theorem characterizes the asymptotic properties of this estimator under stratified designs.

THEOREM 3.2. *Let Assumption 3.1 hold. If $D_{1:n} \sim \text{Loc}(\psi, p)$, then the Lin estimator $\sqrt{n}(\hat{\theta}_L - \text{ATE}) \Rightarrow \mathcal{N}(0, V)$ with*

$$V = \text{Var}(c(X)) + E[\text{Var}(b - \gamma'_L h | \psi)] + E\left[\frac{\sigma_1^2(X)}{p} + \frac{\sigma_0^2(X)}{1 - p}\right].$$

The adjustment coefficient satisfies $\hat{\gamma}_L \xrightarrow{p} \gamma_L$ with $\gamma_L = \text{argmin}_{\gamma \in \mathbb{R}^{d_h}} \text{Var}(b - \gamma' h)$.

The variance V differs from the variance of the unadjusted estimator only in the middle term, which changes from $E[\text{Var}(b | \psi)]$ in the unadjusted case to $E[\text{Var}(b - \gamma'_L h | \psi)]$ for the interacted regression. Crucially, the second statement of Theorem 3.2 shows that the adjustment coefficient γ_L attempts to minimize a marginal variance, instead of the mean-conditional variance that shows up in V above. Because of this, the estimator may be inefficient for stratifications $\psi \neq 1$, since in general

$$\gamma_L = \text{argmin}_{\gamma \in \mathbb{R}^{d_h}} \text{Var}(b - \gamma' h) \neq \text{argmin}_{\gamma \in \mathbb{R}^{d_h}} E[\text{Var}(b - \gamma' h | \psi)] \equiv \gamma^*.$$

Observe that the Lin estimator is completely insensitive to the experimental design, estimating the same adjustment coefficient $\gamma_L = \text{argmin}_{\gamma} \text{Var}(b - \gamma' h)$ for any stratification variables $\psi(X)$. The following example shows that this can lead to strict inefficiency relative to difference of means estimation.

EXAMPLE 3.3 (Random assignment to class size). Suppose $Y(d)$ are student test scores under random assignment to one of two class sizes $d \in \{0, 1\}$. Let $h(X)$ be parent's wealth and $\psi(X)$ previous year (baseline) test scores. Suppose parent's wealth is predictive of future test scores marginally so that $\text{Cov}(h, Y(d)) > 0$. Then $\text{Cov}(h, b) > 0$ and the Lin coefficient is $\gamma_L = \text{Var}(h)^{-1} \text{Cov}(h, b) > 0$. However, if on average parent's wealth has no predictive power for test scores *conditional* on a student's baseline scores (a proxy for ability), then $E[\text{Cov}(h, Y(d) | \psi)] = 0$. In this case, regression adjustment for parent's

wealth $h(X)$ in an experiment stratified on the earlier scores $\psi(X)$ will be strictly less efficient than unadjusted estimation since

$$\begin{aligned} V_{\text{lin}} - V_{\text{unadj}} &= E[\text{Var}(b - \gamma'_L h|\psi)] - E[\text{Var}(b|\psi)] \\ &= -2\gamma_L E[\text{Cov}(h, b|\psi)] + \gamma_L^2 E[\text{Var}(h|\psi)] = \gamma_L^2 E[\text{Var}(h|\psi)] > 0. \end{aligned}$$

An important special case occurs when the design is completely randomized ($\psi = 1$) or if the covariates and stratification variables are independent $h(X) \perp\!\!\!\perp \psi(X)$. In this case, the Lin estimator is weakly more efficient than difference of means since we have

$$E[\text{Var}(b - \gamma'_L h|\psi)] = \text{Var}(b - \gamma'_L h) = \min_{\gamma} \text{Var}(b - \gamma'h) \leq \text{Var}(b).$$

An analogue of Theorem 3.2 also holds for the noninteracted regression estimator $Y_i \sim 1 + D_i + h_i$ under stratified designs $D_{1:n} \sim \text{Loc}(\psi, p)$. The noninteracted estimator is known to be inefficient relative to difference of means even for completely randomized experiments unless $p = 1/2$ or treatment effects are homogeneous. For completeness, we give asymptotic theory and optimality conditions for this estimator under stratified randomization in Section A.3 in the Supplemental Appendix.

We noted above that the Lin estimator $\hat{\theta}_L$ can be written in the canonical form $\hat{\theta}_L = \hat{\theta} - \hat{\gamma}'_L(\bar{h}_1 - \bar{h}_0)$. In fact, most commonly used adjusted estimators can be written in the standard form $\hat{\theta}_{\text{adj}} = \hat{\theta} - \hat{\gamma}'(\bar{h}_1 - \bar{h}_0)$ for some $\hat{\gamma}$, up to order $O_p(n^{-1})$ factors. The following theorem describes the asymptotic properties of general covariate-adjusted estimators $\hat{\theta}_{\text{adj}}$ of this form. To avoid carrying around factors of p in our variance expressions, in what follows, we scale adjusted estimators by the normalization constant $c_p = \sqrt{p(1-p)}$.

THEOREM 3.4. *Let Assumption 3.1 hold. Suppose $\hat{\gamma} \xrightarrow{p} \gamma$ and consider the adjusted estimator*

$$\hat{\theta}_{\text{adj}} = \hat{\theta} - \hat{\gamma}'(\bar{h}_1 - \bar{h}_0)c_p.$$

If $D_{1:n} \sim \text{Loc}(\psi, p)$ then $\sqrt{n}(\hat{\theta}_{\text{adj}} - \text{ATE}) \Rightarrow \mathcal{N}(0, V(\gamma))$ with

$$V(\gamma) = \text{Var}(c(X)) + E[\text{Var}(b - \gamma'h|\psi)] + E\left[\frac{\sigma_1^2(X)}{p} + \frac{\sigma_0^2(X)}{1-p}\right]. \tag{3.3}$$

We define a linearly-adjusted estimator to be asymptotically efficient if it globally minimizes the asymptotic variance $V(\gamma)$ in the previous theorem.

DEFINITION 3.5 (Optimal linear adjustment). The estimator $\hat{\theta}_{\text{adj}} = \hat{\theta} - \hat{\gamma}'(\bar{h}_1 - \bar{h}_0)c_p$ is *efficient* for the design $D_{1:n} \sim \text{Loc}(\psi, p)$ and covariates $h(X)$ if $\hat{\gamma} \xrightarrow{p} \gamma^*$ for an optimal adjustment coefficient

$$\gamma^* \in \underset{\gamma \in \mathbb{R}^{d_h}}{\text{argmin}} E[\text{Var}(b - \gamma'h|\psi)].$$

In particular, $V(\gamma^*) = \min_{\gamma \in \mathbb{R}^{d_h}} V(\gamma)$.

Note that efficiency is defined *relative* to a design $D_{1:n} \sim \text{Loc}(\psi, p)$ and covariates $h(X)$. Setting $\gamma = 0$ recovers unadjusted estimation, so any optimal estimator is in particular weakly more efficient than difference of means.

Optimal adjustment coefficient If $E[\text{Var}(h|\psi)] > 0$, then the optimization in Definition 3.5 is solved uniquely by a mean-conditional OLS coefficient

$$\gamma^* = E[\text{Var}(h|\psi)]^{-1} E[\text{Cov}(h, b|\psi)]. \tag{3.4}$$

Intuitively, fine stratification makes treatment-control imbalances in the covariates $h(X)$ and the potential outcomes $Y(d)$ that are predictable by ψ small enough that they do not contribute to first-order asymptotic variance. Because of this, the optimal covariate-adjusted estimator $\widehat{\theta} - \gamma^*(\bar{h}_1 - \bar{h}_0)c_p$ ignores such fluctuations, minimizing the *mean-conditional* variance objective $E[\text{Var}(b - \gamma'h|\psi)]$, instead of the marginal variance $\text{Var}(b - \gamma'h)$ targeted by the Lin estimator.

Optimal covariates The form of the variance in Equation (3.3) suggests adjusting for variables h that contain predictive information not already contained in ψ . The (unknown) optimal covariates are $h^* = b$. In this case, $\gamma^* = 1$ makes the middle variance term identically zero, and $\widehat{\theta}_{\text{adj}}$ achieves the [Armstrong \(2022\)](#) semiparametric variance bound.

Sample average treatment effect Theorem 3.4 may be extended to covariate-adjusted estimation of the sample average treatment effect $\text{SATE} = E_n[Y_i(1) - Y_i(0)]$. Defining the conditional treatment effect variance $\sigma_\tau^2(X) = \text{Var}(Y(1) - Y(0)|X)$, one can show that $\sqrt{n}(\widehat{\theta} - \text{SATE}) \Rightarrow \mathcal{N}(0, V_S(\gamma))$ with

$$V_S(\gamma) = E[\text{Var}(b - \gamma'h|\psi)] + E\left[\frac{\sigma_1^2(X)}{p} + \frac{\sigma_0^2(X)}{1-p} - \sigma_\tau^2(X)\right]. \tag{3.5}$$

In particular, the optimal adjustment for estimating the ATE and the SATE are the same, with $\gamma_{\text{SATE}}^* = \gamma^*$.

REMARK 3.6 (Nonuniqueness). In general, the optimal adjustment coefficient γ^* may not be unique. For example, if $h(x) = (z(\psi), w(x))$ with $z(\psi)$ a Lipschitz function of the stratification variables, then the variance objective is constant in the coefficient on $z(\psi)$,

$$E[\text{Var}(b - \gamma'_z z - \gamma'_w w|\psi)] = E[\text{Var}(b - \gamma'_w w|\psi)] \quad \forall \gamma_z \in \mathbb{R}^{d_z}.$$

In fact, our analysis shows that the adjustment term $\gamma'_z(\bar{z}_1 - \bar{z}_0) = o_p(n^{-1/2})$ for any coefficient γ_z in this case. Intuitively, since the covariate $z(\psi)$ is already finely balanced by stratifying on $\psi(X)$, *ex post* adjustment by $z(\psi)$ cannot improve first-order efficiency. However, there may still be finite sample efficiency gains from such adjustments, if the covariates $z(\psi)$ are not completely balanced by the stratification. Section 3.5 below provides methods to further adjust for covariates that are functions of the stratification variables.

3.1.1 *Extensions* Before continuing, we briefly mention some extensions to the framework above that are studied in detail in Appendices A.1–A.4.

Experiments with noncompliance In settings with noncompliance, we may instead consider estimation and inference on the local average treatment effect (LATE) of [Imbens and Angrist \(1994\)](#). As a simple application of our main results, Section A.1 characterizes the optimal linear adjustment for estimating the LATE, constructs feasible efficient estimators, and provides asymptotically exact inference on the LATE under stratified randomization with ex post covariate adjustment.

Varying treatment proportions [Cytrynbaum \(2023\)](#) extends Definition 2.1 to allow fine stratification with nonconstant assignment propensity $p(\psi)$. Section A.2 in the Supplemental Appendix characterizes the optimal adjustment coefficient for such designs and derives a feasible efficient estimator.

Nonlinear adjustment In some settings, it may be more natural to use nonlinear or nonparametric covariate adjustment to improve efficiency, for example, in experiments with binary outcomes. Section A.4 in the Supplemental Appendix characterizes the optimal adjustment over a general function space \mathcal{H} for finely stratified designs with varying propensity $p(\psi)$. Feasible estimation of the optimal nonlinear adjustment is an interesting problem that we leave for future work.

3.2 Equivalence with partially linear regression adjustment

This section shows that optimal linear adjustment of a stratified design is as efficient as semiparametric *partially linear* regression adjustment in an experiment with i.i.d. treatments, with adjustment function that is linear in $h(X)$ and nonparametric in $\psi(X)$. This suggests that experimenters stratify on a small set of covariates expected to be most predictive of outcomes at design time, and (efficiently) adjust for the remaining covariates ex post. See below for a more detailed discussion of stratification versus adjustment.

The main result of this section shows first-order asymptotic equivalence of the following (design, estimator) pairs:

$$(D_{1:n} \sim \text{Loc}(\psi, p), \text{optimal linear}) \\ \iff (D_i \overset{\text{iid}}{\sim} \text{Bernoulli}(p), \text{optimal semiparametric}).$$

To define the latter, consider the within-arm partially linear regression models

$$(g_d^*, \gamma_d^*) = \underset{g \in L_2(\psi), \gamma \in \mathbb{R}^{d_h}}{\text{argmin}} E[(Y(d) - g(\psi) - \gamma'h)^2] \tag{3.6}$$

for $d \in \{0, 1\}$. Define the partially linear adjustment function $F_d(x) = g_d^*(\psi(x)) + h(x)' \gamma_d^*$ and consider a [Robins and Rotnitzky \(1995\)](#) style augmented inverse propensity weighting (AIPW) estimator:

$$\widehat{\theta}_{\text{AIPW}} = E_n[F_1(X_i) - F_0(X_i)] + E_n \left[\frac{D_i(Y_i - F_1(X_i))}{p} \right] - E_n \left[\frac{(1 - D_i)(Y_i - F_0(X_i))}{1 - p} \right].$$

The next theorem shows that optimal linear adjustment of the design $D_{1:n} \sim \text{Loc}(\psi, p)$ is asymptotically equivalent to optimal semiparametric adjustment with nonparametric $\psi(X)$ and linear $h(X)$ components.

THEOREM 3.7. *Require Assumption 3.1 and suppose $D_i \stackrel{\text{iid}}{\sim} \text{Bernoulli}(p)$. Then $\sqrt{n}(\widehat{\theta}_{\text{AIPW}} - \text{ATE}) \Rightarrow \mathcal{N}(0, V^*)$ with*

$$V^* = \text{Var}(c(X)) + \min_{\gamma \in \mathbb{R}^{d_h}} E[\text{Var}(b - \gamma'h|\psi)] + E\left[\frac{\sigma_1^2(X)}{p} + \frac{\sigma_0^2(X)}{1-p}\right].$$

The limiting variance V^* is the same as the optimal linearly adjusted variance $V(\gamma^*)$ from Definition 3.5. Intuitively, stratification contributes the nonlinear component of the optimal model $F_d(x)$ above, while optimal adjustment contributes the linear component. The optimal adjustment coefficient $\gamma^* = \sqrt{\frac{1-p}{p}}\gamma_1 + \sqrt{\frac{p}{1-p}}\gamma_0$, for partially linear coefficients γ_1, γ_0 defined in Equation (3.6) above.

Stratification versus regression adjustment Theorem 3.7 shows that stratification provides nonparametric control over the fluctuations of the outcomes predictable by $\psi(X)$, while (linear) adjustment only provides linear control. In first-order asymptotics, this suggests that we stratify on all available covariates, since the variance V^* above is minimized by setting $\psi(X) = X$. However, this may perform poorly in finite samples due to a curse of dimensionality for stratification as $\dim(\psi)$ increases. For example, [Cytrynbaum \(2023\)](#) shows the variance convergence rate $n \text{Var}(\widehat{\theta}) = V + O_p(n^{-2/(\dim(\psi)+1)})$ for the variance V in Equation (3.3), which may be slow even for moderate $\dim(\psi)$. Intuitively, this suggests stratifying on a small set⁴ of covariates $\psi(X)$ expected to be most predictive of outcomes at design time, and planning to optimally adjust for less predictive covariates $h(X)$ ex post.

The next two sections show how to construct linearly-adjusted estimators for the design $D_{1:n} \sim \text{Loc}(\psi, p)$ that achieve the optimal variance V^* .

3.3 Efficiency by rich strata controls

This section provides a “rich covariates” style condition on the relationship between adjustment covariates and stratification variables under which a simple parametric correction of the Lin estimator is fully efficient. The basic idea is to include rich functions $z(\psi)$ of the stratification variables in the adjustment set alongside the additional covariates we would like to adjust for ex post. The main result of this section shows that including $z(\psi)$ as covariates forces the Lin estimator to solve the mean-conditional variance minimization problem of Definition 3.5, restoring asymptotic optimality. An analogous result holds for the noninteracted regression estimator $Y \sim 1 + D + h$ if $p = 1/2$. As a simple application of this section’s results, Example 3.12 shows that for coarsely stratified designs the Lin estimator with leave-one-out strata indicators is efficient.

⁴It is difficult to give concrete guidance for choosing $\dim(\psi)$, since the relevant quantities such as $E[\text{Var}(b|\psi)]$ are not estimable at design time, before we have outcome data. The rate above suggests $\dim(\psi) = o(\log n)$ to achieve the variance V in Equation (3.3).

Consider adjusting for covariates $h(X) = (w(X), z(\psi))$. The main assumption of this section requires that the conditional mean $E[w|\psi]$ is well approximated by known transformations $z(\psi)$ of the stratification variables.

ASSUMPTION 3.8. *There exist $c \in \mathbb{R}^{d_w}$ and $\Lambda \in \mathbb{R}^{d_w \times d_z}$ such that $E[w|\psi] = c + \Lambda z(\psi)$.*

Our next theorem shows that adding such transformations $z(\psi)$ to the adjustment set recovers full efficiency for the Lin estimator.

THEOREM 3.9. *Suppose Assumptions 3.1 and 3.8 hold. Fix adjustment set $h(x) = (w(x), z(\psi))$. Then the Lin estimator $\hat{\theta}_L$ is fully efficient for the design $D_{1:n} \sim \text{Loc}(\psi, p)$. In particular, $\sqrt{n}(\hat{\theta}_L - \text{ATE}) \Rightarrow \mathcal{N}(0, V^*)$ with*

$$V^* = \text{Var}(c(X)) + \min_{\gamma \in \mathbb{R}^{d_h}} E[\text{Var}(b - \gamma' h|\psi)] + E\left[\frac{\sigma_1^2(X)}{p} + \frac{\sigma_0^2(X)}{1-p}\right].$$

Moreover, the asymptotic variance has

$$\min_{\gamma \in \mathbb{R}^{d_h}} E[\text{Var}(b - \gamma' h|\psi)] = \min_{\alpha \in \mathbb{R}^{d_w}} E[\text{Var}(b - \alpha' w|\psi)].$$

In practice, Theorem 3.9 suggests including flexible functions $z(\psi)$ of the stratification variables in the adjustment set. The proof is given in Section A.6 of the Supplemental Appendix. The following corollary follows shows that if $p = 1/2$ (matched pairs) or if treatment effect heterogeneity is limited then the noninteracted regression $Y \sim 1 + D + w + z$ with rich strata controls $z(\psi)$ is also asymptotically efficient.

COROLLARY 3.10. *Suppose additionally that $p = 1/2$ or $E[\text{Cov}(Y(1) - Y(0), w|\psi)] = 0$. Then the coefficient θ_N on D_i in the regression $Y \sim 1 + D + w + z$ is asymptotically efficient.*

The condition $E[\text{Cov}(Y(1) - Y(0), w|\psi)] = 0$ limits the explanatory power of covariates w for treatment effect heterogeneity, conditional on the stratification variables.

REMARK 3.11 (Indirect efficiency gain). The second statement of the theorem shows that optimal adjustment for $h(X)$ is as efficient as optimal adjustment for the subvector $w(X) \subseteq h(X) = (w(X), z(\psi))$. In this sense, the efficiency improvement due to including $z(\psi)$ is indirect. Indeed, our analysis shows that $\hat{\theta} - \gamma'_z(\bar{z}_1 - \bar{z}_0) = \hat{\theta} + o_p(n^{-1/2})$ for any $\gamma_z \in \mathbb{R}^{d_z}$, so adjustment for $z(\psi)$ alone cannot affect the first-order asymptotic variance. Intuitively, we are just using the inclusion of $z(\psi)$ as a device to “tilt” the coefficient on $w(X)$, forcing the Lin estimator to solve the correct mean-conditional variance optimization problem.

The next example uses Theorem 3.9 to show that including leave-one-out strata indicators as covariates in the Lin estimator restores asymptotic efficiency for coarsely stratified designs.

EXAMPLE 3.12 (Coarse stratification). Consider stratified randomization $D_{1:n} \sim \text{Loc}(S, p)$ with fixed strata $S(x) \in \{1, \dots, m\}$. Let the adjustment covariates be $h(x) = (w(x), z(x))$ with leave-one-out strata indicators $z(S_i) = (\mathbb{1}(S_i = k))_{k=1}^{m-1}$. In this case, Assumption 3.8 is automatically satisfied since we can write $E[w|S] = c + \Lambda z$ with $c = E[w|S = m]$ and $\Lambda_{jk} = (E[w_j|S = k] - E[w_j|S = m])_{jk}$. Then by Theorem 3.9, the Lin estimator $\widehat{\theta}_L$ with covariates $h_i = (w_i, z_i)$ is efficient. In particular, we have $\sqrt{n}(\widehat{\theta}_L - \text{ATE}) \Rightarrow \mathcal{N}(0, V^*)$ with optimal variance

$$V^* = \text{Var}(c(X)) + \min_{\gamma} E[\text{Var}(b - \gamma'w|S)] + E\left[\frac{\sigma_1^2(X)}{p} + \frac{\sigma_0^2(X)}{1-p}\right].$$

Similarly, by Corollary 3.10 if $p = 1/2$, then including leave-one-out strata fixed effects in the noninteracted regression restores efficiency.

REMARK 3.13 (Fine stratification). Note that the argument in Example 3.12 only applies to coarse stratification, where the strata $S(x) \in \{1, \dots, m\}$ are data-independent and fixed as $n \rightarrow \infty$. For fine stratification $D_{1:n} \sim \text{Loc}(\psi, p)$ with continuous covariates $\psi(x)$, the strata are data-dependent and number of strata $m \asymp n$, so Theorem 3.23 does not apply. Indeed, for matched pairs the Lin regression in Example 3.12 would have $n + 2 \dim(h) > n$ covariates, producing collinearity. This collinearity problem occurs more generally; see Remark 3.18 below for further discussion.

Leaving behind the rich covariates Assumption 3.8, the next section provides new adjusted estimators that are fully efficient for any design in the class $D_{1:n} \sim \text{Loc}(\psi, p)$ under weak conditions.

3.4 Generic efficient adjustment

In this section, we study several adjusted estimators that are asymptotically efficient under weak conditions for any stratified design $D_{1:n} \sim \text{Loc}(\psi, p)$. For matched pairs designs, or in settings with limited treatment effect heterogeneity, the noninteracted regression including treatment, covariates, and pair fixed effects is efficient. More generally, we show that the following estimators are efficient under weak assumptions:

- (1) **PL** – A partialled Lin estimator with within-stratum (inconsistently) partialled covariates.
- (2) **GO** – A “Group OLS” estimator, generalizing a proposal of Imbens and Rubin (2015) for matched pairs designs.
- (3) **TM** – A tyranny-of-the-minority (ToM) estimator for stratified designs.

The main new condition we impose in this section is that the adjustment covariates are not collinear, conditionally on the stratification variables. This guarantees that the optimal adjustment coefficient γ^* is unique with $\gamma^* = E[\text{Var}(h|\psi)]^{-1} E[\text{Cov}(h, b|\psi)]$, as discussed in Section 3.1.

ASSUMPTION 3.14. *The conditional variance satisfies $E[\text{Var}(h|\psi)] > 0$.*

Note that this assumption rules out adjustment for functions $h(\psi)$ of the stratification variables. To see why it is necessary, consider that, for example, in a regression with full strata fixed effects $Y \sim D + h + z^n$, covariates $h_i = h(\psi_i)$ would be asymptotically collinear with the strata fixed effects $z^n = (\mathbb{1}(i \in g_j))_{j=1}^{n/k}$. More intuitively, the problem is that $h(\psi)$ has too little residual variation within local regions of $\psi(X)$ space defining the fine strata. We noted earlier that $\hat{\theta} - \alpha'(\bar{h}_1 - \bar{h}_0) = \hat{\theta} - o_p(n^{-1/2})$ for any $\alpha \in \mathbb{R}^{d_h}$, so such adjustment cannot improve first-order efficiency. Nevertheless, one may still wish to adjust for $h(\psi)$ to correct finite sample imbalances not controlled by the design. Adjustment for such variables needs to be handled slightly differently, and we construct modified efficient estimators for this purpose in Section 3.5 below.

3.4.1 *Strata fixed-effects estimator* Recall that for $p = a/k$, a finely stratified design $D_{1:n} \sim \text{Loc}(\psi, p)$ partitions the experimental units $\{1, \dots, n\}$ into n/k disjoint groups g . Define the fixed-effects estimator $\hat{\theta}_{\text{FE}}$ by the least squares equation

$$Y_i = \hat{\theta}_{\text{FE}}D_i + \hat{\gamma}'_{\text{FE}}h_i + \sum_{j=1}^{n/k} \hat{a}_j \mathbb{1}(i \in g_j) + e_i. \tag{3.7}$$

The next theorem shows that $\hat{\theta}_{\text{FE}}$ is fully efficient in the case of matched pairs or if treatment effect heterogeneity is limited, but may be inefficient in general.

THEOREM 3.15. *Suppose Assumptions 3.1 and 3.14 hold. The estimator has representation $\hat{\theta}_{\text{FE}} = \hat{\theta} - \hat{\gamma}'_{\text{FE}}(\bar{h}_1 - \bar{h}_0) + O_p(n^{-1})$. If $D_{1:n} \sim \text{Loc}(\psi, p)$, then $\sqrt{n}(\hat{\theta}_{\text{FE}} - \text{ATE}) \Rightarrow \mathcal{N}(0, V)$ with variance*

$$V = \text{Var}(c(X)) + E[\text{Var}(b - \gamma'_{\text{FE}}h|\psi)] + E\left[\frac{\sigma_1^2(X)}{p} + \frac{\sigma_0^2(X)}{1-p}\right].$$

and coefficient $\gamma_{\text{FE}} = \text{argmin}_{\gamma \in \mathbb{R}^{d_h}} E[\text{Var}(f - \gamma'h|\psi)]$ for target function

$$f(x) = m_1(x)\sqrt{\frac{p}{1-p}} + m_0(x)\sqrt{\frac{1-p}{p}}.$$

The function $f \neq b$ in general. If $p = 1/2$, then $f = b$ and the fixed-effects estimator is efficient. If $p \neq 1/2$, it is efficient if and only if $E[\text{Cov}(h, Y(1) - Y(0)|\psi)] = 0$.

See Section A.7 for the proof. Asymptotically exact inference for the ATE using $\hat{\theta}_{\text{FE}}$ is available using the tools in Section 4.

REMARK 3.16 (Conditions for efficiency). If $p = 1/2$, then $f = b$ and $\hat{\theta}_{\text{FE}}$ is efficient. More generally, $f(x) \neq b(x)$ and $\hat{\theta}_{\text{FE}}$ solves the wrong variance minimization problem, effectively targeting the wrong linear combination of outcomes. The necessary and sufficient condition $E[\text{Cov}(h, Y(1) - Y(0)|\psi)] = 0$ requires that treatment effect heterogeneity is not explained by the covariates $h(X)$, conditional on the stratification variables.

In the rest of this section, we develop estimators that are fully efficient for any finely stratified design, without imposing any assumptions on treatment effect heterogeneity or treatment proportions.

3.4.2 Partialled Lin estimator First, we define a partialled version of the Lin estimator. Let $g(i)$ denote the group that unit i belongs to and define the within-group partialled covariates

$$\check{h}_i = h_i - \frac{1}{k} \sum_{j \in g(i)} h_j.$$

For example, if $k = 2$ this is just the within-pair covariate difference $\check{h}_i = (1/2)(h_i - h_{m(i)})$, where i is matched to $m(i)$. We can think of \check{h}_i as an inconsistent but approximately unbiased signal for the nonparametrically residualized covariate $h_i - E[h_i|\psi_i]$. Next, we use these partialled covariates in the Lin regression

$$Y \sim 1 + D_i + \check{h}_i + D_i \check{h}_i. \tag{3.8}$$

Define the *partialled* Lin estimator $\hat{\theta}_{PL}$ to be the coefficient on D_i in this regression. For reference, similar to the Lin regression, we may write this in the standard form $\hat{\theta}_{PL} = \hat{\theta} - \hat{\gamma}_{PL}(\bar{h}_1 - \bar{h}_0)c_p$ with adjustment coefficient $\hat{\gamma}_{PL} = (\hat{a}_1 + \hat{a}_0)\sqrt{\frac{1-p}{p}} + \hat{a}_0\sqrt{\frac{p}{1-p}}$, where \hat{a}_0 and \hat{a}_1 are coefficients on \check{h}_i and $D_i\check{h}_i$.

Our main result in Theorem 3.23 below shows that the partialled Lin estimator $\hat{\theta}_{PL}$ is asymptotically efficient in the sense of Definition 3.5, with $\hat{\gamma}_{PL} \xrightarrow{p} \gamma^*$ for the optimal adjustment coefficient γ^* .

REMARK 3.17 (Intuition for optimality). Theorem 3.4 showed that an estimator $\hat{\theta} - \hat{\gamma}(\bar{h}_1 - \bar{h}_0)c_p$ is efficient if $\hat{\gamma} \xrightarrow{p} \gamma^*$ and γ^* solves the conditional-mean variance problem $\gamma^* \in \operatorname{argmin}_{\gamma} E[\operatorname{Var}(b - \gamma'h|\psi)]$. By using within-stratum partialled regressors \check{h}_i , we force the Lin estimator to only use covariate signal $h_i - E[h_i|\psi_i]$ that is mean-independent of the stratification variables.

REMARK 3.18 (Treatment-strata interactions). As an alternative to $\hat{\theta}_{PL}$, one may attempt to use the Lin regression $Y_i \sim (1, h_i, g^n(i)) + D_i(1, h_i, g^n(i))$ with leave-one-out strata fixed effects $g^n(i) = (\mathbb{1}(i \in g_j))_{j=1}^{n/k-1}$. Unfortunately, this produces collinear regressors for $p = a/k$ if either $a = 1$ or $a = k - 1$, which includes the case of matched pairs. To see the issue, one can show by Frisch–Waugh that in contrast to Equation (3.8) above, this estimator partials covariates h_i separately in each treatment arm, using $\check{h}_{i1} = h_i - a^{-1} \sum_{j \in g(i)} h_j D_j$ if $D_i = 1$ and $\check{h}_{i0} = h_i - (k - a)^{-1} \sum_{j \in g(i)} h_j (1 - D_j)$ if $D_i = 0$. For instance, if $a = 1$ then this is $\check{h}_i = h_i - h_i = 0$ for all i , showing collinearity. In the case $1 < a < k - 1$ where this estimator is feasible, it is asymptotically equivalent to the partialled Lin estimator. However, finite-sample properties will be worse due to noisier within-arm partialling.

REMARK 3.19. A calculation shows that our estimator $\hat{\theta}_{PL}$ is numerically equivalent to a regression estimator proposed in Ke, Liu, and Yang (2024), which the authors derive alternately through an optimal projection argument. They study estimation of the SATE under stratified randomization in a finite population framework, providing conservative inference. They do not derive the exact form of the asymptotic variance, instead leaving it as an infinite sum, which they assume converges to some limit. By contrast, we derive the exact form of the asymptotic variance under the data-adaptive stratifications in Definition 2.1, enabling asymptotically exact inference on the ATE using $\hat{\theta}_{PL}$.

3.4.3 *Group OLS estimator* Next, we generalize an estimator proposed by Imbens and Rubin (2015) for covariate adjustment in matched pairs experiments to more general stratified designs. For each group of units $g = 1, \dots, n/k$ in the design $D_{1:n} \sim \text{Loc}(\psi, p)$, define the within-group difference of means of outcomes and covariates

$$y_g = \frac{1}{k} \sum_{i \in g} \frac{Y_i D_i}{p} - \frac{1}{k} \sum_{i \in g} \frac{Y_i (1 - D_i)}{1 - p} \quad \text{and} \quad h_g = \frac{1}{k} \sum_{i \in g} \frac{h_i D_i}{p} - \frac{1}{k} \sum_{i \in g} \frac{h_i (1 - D_i)}{1 - p}.$$

For any group-indexed array $(x_g)_g$, denote $E_g[x_g] = \frac{k}{n} \sum_g x_g$. Define the *group OLS* estimator $\hat{\theta}_G$ by the regression

$$y_g = \hat{\theta}_G + \hat{\gamma}'_G h_g + e_g \tag{3.9}$$

with $E_g[(1, h_g) e_g] = 0$. For motivation, note that if $h = 0$ then this becomes $y_g = \hat{\theta}_G + e_g$ and $\hat{\theta}_G$ is just the unadjusted estimator $\hat{\theta}_G = \bar{Y}_1 - \bar{Y}_0$. More generally, the adjusted version can be written $\hat{\theta}_G = E_g[y_g] - \hat{\gamma}'_G E_g[h_g] = \hat{\theta} - \hat{\gamma}'_G (\bar{h}_1 - \bar{h}_0)$ with adjustment coefficient $\hat{\gamma}_G = \text{Var}_g(h_g)^{-1} \text{Cov}_g(h_g, y_g)$. The estimators $\hat{\theta}_G$ and $\hat{\theta}_{PL}$ are numerically identical for the case of matched pairs, but not for $p \neq 1/2$. The main result of this section shows that $\hat{\theta}_G$ is asymptotically equivalent to the partialled Lin estimator $\hat{\theta}_{PL}$, and both are asymptotically optimal.

REMARK 3.20 (Intuition for efficiency). The estimator $\hat{\theta}_G$ uses within-group differences of covariates $\bar{h}_{g1} - \bar{h}_{g0}$ to predict within-group outcome differences $\bar{Y}_{1g} - \bar{Y}_{0g}$. Similar to the partialled Lin strategy, by doing this we only measure the variation in covariates and potential outcomes orthogonal to the stratification variables. This forces least squares to compute a conditional variance-covariance trade-off, solving the optimal adjustment problem in Definition 3.5. In particular, the proof of Theorem 3.23 shows that if $D_{1:n} \sim \text{Loc}(\psi, p)$, then the adjustment coefficient

$$\hat{\gamma}_G = \text{Var}_g(h_g)^{-1} \text{Cov}_g(h_g, y_g) \xrightarrow{p} c_p \underset{\gamma}{\text{argmin}} E[\text{Var}(b - \gamma' h | \psi)].$$

REMARK 3.21. Imbens and Rubin (2015) propose $\hat{\theta}_G$ in the case of matched pairs $p = 1/2$. Their analysis uses a toy sampling model where the pairs themselves are drawn “pre-matched” from a superpopulation. By contrast, we model the experimental units as being sampled from a superpopulation, with units matched into data-dependent strata post-sampling. This more realistic model complicates the analysis, producing different

limiting variances and requiring different inference procedures. In a design-based setting, Fogarty (2018) shows that the Imbens and Rubin (2015) estimator is weakly more efficient than difference of means for matched pairs designs. By contrast, we extend this estimator to a larger family of fine stratifications strictly containing matched pairs, and show that it is asymptotically optimal among linearly adjusted estimators.

3.4.4 *Tyranny-of-the-minority estimator* Finally, we define tyranny-of-the-minority (ToM) adjustment, extending Lin (2013). To do so, define the adjustment coefficient:

$$\hat{\gamma}_{\text{TM}} = \text{Var}_n(\check{h}_i)^{-1} \left(\text{Cov}_n(\check{h}_i, Y_i | D_i = 1) \sqrt{\frac{1-p}{p}} + \text{Cov}_n(\check{h}_i, Y_i | D_i = 0) \sqrt{\frac{p}{1-p}} \right). \quad (3.10)$$

Define the ToM estimator $\hat{\theta}_{\text{TM}} = \hat{\theta} - \hat{\gamma}_{\text{TM}}(\bar{h}_1 - \bar{h}_0)c_p$. The main difference between the ToM and partialled Lin adjustment coefficients is that $\hat{\gamma}_{\text{TM}}$ estimates the conditional variance $E[\text{Var}(h|\psi)]$ only once, using the sample variance $\text{Var}_n(\check{h}_i)$ for the full experimental sample. By contrast, partialled Lin estimates this term separately in each treatment arm, using $\text{Var}_n(\check{h}_i | D_i = 1)$ and $\text{Var}_n(\check{h}_i | D_i = 0)$. Because of this, we expect $\hat{\theta}_{\text{TM}}$ to be more stable than $\hat{\theta}_{\text{PL}}$ in small experiments.

REMARK 3.22. Lu and Liu (2024) propose an alternate ToM regression adjustment for stratified experiments. To compare the approaches, for propensity $p = a/k$ define the within-arm partialling $\check{h}_{i1} = h_i - a^{-1} \sum_{i \in g} D_i h_i$ and $\check{h}_{i0} = h_i - (k-a)^{-1} \sum_{i \in g} (1-D_i) h_i$. Their estimator takes the form $\hat{\theta}_{\text{LL}} = \hat{\theta} - \hat{\gamma}'_{\text{LL}}(\bar{h}_1 - \bar{h}_0)$. In our notation, their adjustment coefficient $\hat{\gamma}_{\text{LL}} = \hat{S}_{hh}^{-1} \hat{S}_{hY}$ has

$$\hat{S}_{hh} = E_n \left[\frac{D_i \check{h}_{i1} \check{h}'_{i1}}{p^2} \frac{a}{a-1} + \frac{(1-D_i) \check{h}_{i0} \check{h}'_{i0}}{(1-p)^2} \frac{k-a}{k-a-1} \right]$$

and similarly for \hat{S}_{hY} . Their approach is infeasible if $a = 1$ or $a = k - 1$. For example, this prohibits its use in matched pairs and matched triples experiments.

3.4.5 *Main result* The main result of this section shows that all three estimators above are asymptotically equivalent and efficient in the sense of Definition 3.5.

THEOREM 3.23. *Suppose Assumptions 3.1 and 3.14 hold. If $D_{1:n} \sim \text{Loc}(\psi, p)$, then $\hat{\theta}_{\text{PL}} - \hat{\theta}_G = o_p(n^{-1/2})$ and $\hat{\theta}_{\text{PL}} - \hat{\theta}_{\text{TM}} = o_p(n^{-1/2})$. We have $\sqrt{n}(\hat{\theta}_{\text{PL}} - \text{ATE}) \Rightarrow \mathcal{N}(0, V^*)$ with the optimal variance*

$$V^* = \text{Var}(c(X)) + \min_{\gamma} E[\text{Var}(b - \gamma h|\psi)] + E \left[\frac{\sigma_1^2(X)}{p} + \frac{\sigma_0^2(X)}{1-p} \right].$$

Methods for asymptotically exact inference on the ATE using these estimators are discussed in Section 4 below. Our simulations and empirical results show that the partialled Lin, Group OLS, and ToM estimators behave very similarly in finite samples.

3.5 Further adjustment for stratification variables

In this section, we provide modified versions of the previous estimators that allow further adjustment for covariates $z(\psi)$ that are functions of the stratification variables. As discussed above, this cannot improve first-order efficiency but may improve finite sample performance by correcting for any remaining imbalances in ψ not controlled by the stratification.

Denote $z_i = z(\psi_i)$. For each estimator $\hat{\theta}_k$ above with $k \in \{\text{FE, PL, G, TM}\}$, we define a modified estimator of the form $\hat{\tau}_k = \hat{\theta}_k - \hat{\alpha}'_k(\bar{z}_1 - \bar{z}_0)c_p$. For the fixed effects estimator, define $\hat{\tau}_{\text{FE}}$ to be the coefficient on D_i in the regression $Y_i \sim (1, D_i, \check{h}_i, z_i)$. For the partialled Lin estimator, define $\hat{\tau}_{\text{PL}}$ to be the coefficient on D_i in the regression $Y_i \sim (1, \check{h}_i, z_i) + D_i(1, \check{h}_i, z_i)$. Define the modified ToM estimator to be as in Equation (3.10), with (\check{h}_i, z_i) in place of \check{h}_i . Finally, define the modified group OLS estimator $\hat{\tau}_G = \hat{\theta}_G - \hat{\alpha}'_G(\bar{z}_1 - \bar{z}_0)c_p$, with $\hat{\alpha}_G = \hat{\alpha}_{\text{PL}}$. Our next theorem shows that these estimators are asymptotically equivalent to the original versions of each estimator that do not adjust for $z(\psi)$. However, the simulations in Sections 5 and 6 show that they may perform better in small experiments.

THEOREM 3.24. *Suppose Assumptions 3.1 and 3.14 hold, as well as $\text{Var}(z) > 0$ and $E[|z|_2^2] < \infty$. Then if $D_{1:n} \sim \text{Loc}(\psi, p)$, we have $\hat{\tau}_k = \hat{\theta}_k + o_p(n^{-1/2})$ for $k \in \{\text{FE, PL, G, TM}\}$. Each estimator has the form $\hat{\tau}_k = \hat{\theta}_k - \hat{\alpha}'_k(\bar{z}_1 - \bar{z}_0)c_p$ with $\hat{\alpha}_{\text{FE}} \xrightarrow{p} \text{argmin}_\alpha \text{Var}(f - \alpha'z)$ for f as in Theorem 3.15 and $\hat{\alpha}_{\text{PL}}, \hat{\alpha}_G, \hat{\alpha}_{\text{TM}} \xrightarrow{p} \text{argmin}_\alpha \text{Var}(b - \alpha'z)$.*

From the second statement of the theorem, we can interpret the modified estimators as taking a conservative approach that ignores stratification on ψ and adjusts for imbalances in $z(\psi)$ as if the experiment were completely randomized.

4. INFERENCE

In this section, we provide asymptotically exact confidence intervals for the ATE in stratified experiments using generic linearly adjusted estimators. Overcoverage is known to be a problem for inference based on the usual Eicker–Huber–White (EHW) variance estimator in stratified experiments. For example, Bai, Romano, and Shaikh (2021) shows that the EHW variance estimators for $Y \sim 1 + D + h$ and the fixed-effects regression $Y \sim D + h + z^n$ are asymptotically conservative for matched pairs designs if $h = 0$. To the best of our knowledge, we give the first asymptotically exact inference methods for covariate-adjusted ($h \neq 0$) ATE estimation under general stratified designs. Our main inference result applies to any estimator of the form $\hat{\theta} - \hat{\gamma}'(\bar{h}_1 - \bar{h}_0)c_p + o_p(n^{-1/2})$. In particular, this enables asymptotically exact inference on the ATE using any of the estimators in this paper. Our confidence intervals are shorter than those produced by EHW in the simulations and empirical application below, taking full advantage of the efficiency gains from both stratification and covariate adjustment.

To define our inference methods, consider such an estimator $\hat{\theta}(\hat{\gamma}) = \hat{\theta} - \hat{\gamma}'(\bar{h}_1 - \bar{h}_0)c_p$ with $\hat{\gamma} \xrightarrow{p} \gamma$. Define the augmented potential outcomes $Y_i^a(d) = Y_i(d) - c_p \hat{\gamma}' h_i$ for $d \in$

$\{0, 1\}$ and the augmented outcome $Y_i^a = Y_i - c_p \widehat{\gamma}' h_i$. Then apparently

$$\widehat{\theta}(\widehat{\gamma}) = \bar{Y}_1 - \bar{Y}_0 - c_p \widehat{\gamma}' (\bar{h}_1 - \bar{h}_0) = \bar{Y}_1^a - \bar{Y}_0^a. \tag{4.1}$$

Our strategy is to apply the inference results of Cytrynbaum (2023) for difference of means estimation $\widehat{\theta} = \bar{Y}_1 - \bar{Y}_0$ to the difference of augmented potential outcomes $\bar{Y}_1^a - \bar{Y}_0^a$. To do so, let \mathcal{G}_n denote the set of groups in Definition 2.1. For each $g \in \mathcal{G}_n$, define the group centroid $\bar{\psi}_g = |g|^{-1} \sum_{i \in g} \psi_i$. Let $\nu : \mathcal{G}_n \rightarrow \mathcal{G}_n$ be a bijective matching between groups satisfying $\nu(g) \neq g, \nu^2 = \text{Id}$, and the homogeneity condition

$$\frac{1}{n} \sum_{g \in \mathcal{G}_n} |\bar{\psi}_g - \bar{\psi}_{\nu(g)}|_2^2 = o_p(1). \tag{4.2}$$

In practice, ν is obtained by simply matching the group centroids $\bar{\psi}_g$ into pairs using the Derigs (1988) matching algorithm. Let $\mathcal{G}_n^\nu = \{g \cup \nu(g) : g \in \mathcal{G}_n\}$ be the unions of paired groups formed by this matching. Define $a(g) = \sum_{i \in g} D_i$ and $k(g) = |g|$. Finally, define the variance estimator components:

$$\begin{aligned} \widehat{v}_1 &= n^{-1} \sum_{g \in \mathcal{G}_n^\nu} \frac{1}{a(g) - 1} \sum_{i \neq j \in g} \frac{Y_i^a Y_j^a D_i D_j (1 - p)}{p^2}, \\ \widehat{v}_0 &= n^{-1} \sum_{g \in \mathcal{G}_n^\nu} \frac{1}{(k - a)(g) - 1} \sum_{i \neq j \in g} \frac{Y_i^a Y_j^a (1 - D_i)(1 - D_j) p}{(1 - p)^2}, \\ \widehat{v}_{10} &= n^{-1} \sum_{g \in \mathcal{G}_n} \frac{k}{a(k - a)}(g) \sum_{i, j \in g} Y_i^a Y_j^a D_i (1 - D_j). \end{aligned}$$

Next, define the variance estimator:

$$\widehat{V} = \text{Var}_n \left(\frac{(D_i - p) Y_i^a}{p - p^2} \right) - \widehat{v}_1 - \widehat{v}_0 - 2\widehat{v}_{10}. \tag{4.3}$$

Our inference strategy begins with the sample variance of the adjusted estimator, which is consistent for the asymptotic variance of $\widehat{\theta}_{\text{adj}}$ under an i.i.d. design, but too large under stratified designs. We correct this sample variance using the estimators above, which measure how well the stratification variables predict augmented outcomes in local regions of the covariate space. This section’s main result shows that \widehat{V} is consistent for the limiting variance of Theorem 3.4, enabling asymptotically exact inference on the ATE using adjusted estimators.

THEOREM 4.1 (Inference). *Under the conditions of Theorem 3.4, if $D_{1:n} \sim \text{Loc}(\psi, p)$, then $\widehat{V} = V + o_p(1)$.*

By Theorem 4.1 and our previous asymptotic results in Theorem 3.4, the confidence interval $\widehat{C} = [\widehat{\theta}(\widehat{\gamma}) \pm \widehat{V}^{1/2} c_{1-\alpha/2} / \sqrt{n}]$ with $c_\alpha = \Phi^{-1}(\alpha)$ is asymptotically exact in the sense that $P(\text{ATE} \in \widehat{C}) = 1 - \alpha + o(1)$.

5. SIMULATIONS

In this section, we use simulations to test the finite-sample performance of the estimators studied above. We consider quadratic outcome models of the form

$$Y_i(d) = \psi'_i Q_d \psi_i + \psi'_i L_d + c_d \cdot u(X_i) + \epsilon_i^d \quad E[\epsilon_i^d | X_i] = 0$$

for $d \in \{0, 1\}$. The component $u_i = u(X_i)$ represents covariate signal that is independent of the stratification variables $\psi(X_i)$. After implementing the design $D_{1:n} \sim \text{Loc}(\psi, p)$, we receive access to scalar covariates h_i that are correlated with both ψ_i and $Y_i(d)$. In particular, suppose that $h_i = \psi'_i Q_h \psi_i + \psi'_i L_h + u_i$ with $E[u_i | \psi_i] = 0$. In the following simulations, we let $\psi_i \sim N(0, I_m)$, $u_i \sim N(0, 1)$, and $\epsilon_i^d \sim N(0, 1/10)$ with $(\psi_i, u_i, \epsilon_i^d)$ jointly independent. We use treatment proportions $p = 2/3$ unless otherwise specified. With $m \equiv \dim(\psi)$, let $A \in \mathbb{R}^{m \times m}$ have $A_{ij} = 1$ for $i \neq j$ and $A_{ii} = 0$. We simulate the following DGPs:

Model 1: Quadratic coefficients $Q_h = (1/m^2)A$ and $Q_0 = Q_1 = (1/m)A$. Linear coefficients $L_0 = \mathbb{1}_m, L_1 = 2\mathbb{1}_m, L_h = \mathbb{1}_m$. Regressor signal $c_1 = c_0 = -3$.

Model 2: As in Model 1 but $c_0 = -4$ and $c_1 = -1$.

Model 3: As in Model 2 but $p = 1/2$.

Model 4: As in Model 1 but $c_0 = 2$ and $c_1 = 4$.

Model 5: As in Model 1 but $c_0 = 2$ and $c_1 = 4$ and $p = 1/2$.

Model 6: As in Model 1 but $Q_h = (1/100)A$.

We begin by comparing the efficiency properties of different linearly adjusted estimators. **Unadj** refers to simple difference of means (unadjusted). The **Lin** estimator is studied in Theorem 3.2. **Naive** refers to the noninteracted regression $Y \sim (1, D, h)$, (Theorem A.4). **FE** refers to the fixed-effects estimator (Theorem 3.15) and **Plin** the partialled Lin estimator (Theorem 3.23). **GO** refers to Group OLS and **ToM** refers to Tyranny-of-the-Minority estimation (Theorem 3.23). **Strata controls** refer to modified versions of each of the previous estimators that further adjust for parametric strata controls $z(\psi)$, as discussed in Section 3.5. In our simulations, we set $z(\psi) = \psi$. **Ad** refers to an adaptive⁵ estimator that sets $\hat{\theta}_{\text{adj}} = \hat{\theta}_L$ if $\widehat{V}(\hat{\gamma}_L) \leq \widehat{V}(\hat{\gamma}_{\text{PL}})$ and $\hat{\theta}_{\text{adj}} = \hat{\theta}_{\text{PL}}$ otherwise,⁶ including parametric controls $z(\psi) = \psi$ in both cases.

Table 1 studies finite sample efficiency. We present the mean squared error (MSE) ratio, relative to unadjusted estimation, for each of the adjusted estimators above. The bottom line of the table reports the *excess risk* R_k of each estimator k relative to the optimal estimator. To define this, let $\text{MSE}_{k,s}$ be the relative MSE of estimator k in simulation

⁵This estimator is pointwise asymptotically equivalent to $\hat{\theta}_{\text{PL}}$. Issues with post model-selection inference (e.g., Leeb and Pötscher (2005)) appear to be less worrying here, since even under the fixed alternative $\gamma^* \neq \gamma_L$, the Lin estimator is still \sqrt{n} -consistent and asymptotically unbiased.

⁶We could also use a cross-fit version of $\widehat{V}(\gamma)$ to reduce bias. However, the in-sample criterion performed quite well in our simulations.

TABLE 1. Ratio of MSEs (%), adjusted vs. unadjusted estimation.

$(n, \dim(\psi))$	No strata controls								Strata controls $z(\psi)$						
	Model	Unadj	Naive	Lin	FE	Plin	GO	ToM	Naive	Lin	FE	Plin	GO	ToM	Ad
(600, 2)	1	100	113	102	49	48	49	48	36	35	35	37	37	36	34
	2	100	126	102	64	57	58	57	60	46	52	47	47	47	45
	3	100	116	116	38	38	38	38	48	48	36	36	37	36	37
	4	100	27	31	31	27	27	27	26	26	38	32	33	32	26
	5	100	28	28	18	18	18	18	21	21	19	19	19	19	19
	6	100	100	100	100	11	11	11	11	7	7	9	9	9	9
(1200, 2)	1	100	114	103	44	44	44	44	35	34	31	33	33	33	32
	2	100	126	102	60	56	56	56	61	47	50	47	46	47	45
	3	100	116	116	38	38	38	38	48	48	37	37	37	37	37
	4	100	26	30	29	25	25	25	23	24	36	30	30	30	24
	5	100	28	28	17	17	17	17	20	20	17	18	17	18	18
	6	100	101	101	9	9	9	9	7	7	8	8	8	8	7
(1200, 5)	1	100	142	127	85	84	84	84	25	24	41	46	55	46	24
	2	100	145	123	94	86	87	86	45	34	57	54	62	54	34
	3	100	137	137	81	81	81	81	40	40	54	54	57	54	40
	4	100	27	31	31	27	27	27	25	20	54	45	49	45	20
	5	100	32	32	24	24	24	24	18	18	38	38	39	38	18
	6	100	138	138	67	67	67	67	15	15	36	36	39	37	15
R_k	73	65	60	17	15	15	15	5	2	10	8	10	8	0.2	

s . Then we set $R_k = (1/S) \sum_s (\text{MSE}_{k,s} - \min_j \text{MSE}_{j,s})$, averaging over all simulations in the table. All results are calculated using 2000 Monte Carlo repetitions.

In models 1, 2, and 3, both Naive and Lin style linear adjustment are strictly inefficient relative to unadjusted estimation. These models have marginal covariance $\text{Cov}(Y(d), h) > 0$ but conditional covariance $E[\text{Cov}(Y(d), h|\psi)] < 0$, conditional on the stratification variables. Because of this, the optimal adjustment coefficient $\gamma^* < 0$, while the Naive and Lin regressions estimate positive adjustment coefficients $\gamma_N, \gamma_L > 0$, leading to even worse performance than unadjusted estimation in some cases. For Models 4 and 5, the **Naive** and **Lin** methods are competitive with the generic efficient methods from Section 3.4. This is because in these cases we made it so that $\text{Cov}(Y(d), h) \approx E[\text{Cov}(Y(d), h|\psi)]$, so that “by chance” γ^* is close to γ_N and γ_L . However, the parametric coefficients γ_N and γ_L are estimated more precisely than the semiparametric object $\gamma^* = E[\text{Var}(h|\psi)]^{-1} E[\text{Cov}(h, b|\psi)]$. For Model 6, **Lin** with $z(\psi) = \psi$ controls is (approximately) optimal by Theorem 3.9, since $E[w|\psi]$ is (approximately) linear in ψ .

Summarizing our findings, the **Lin**, **Plin**, and **Naive** estimators with parametric **Strata controls** $z(\psi) = \psi$ had low excess risk across specifications, while the **Ad** estimator was the most efficient overall. The **Naive** and **Lin** estimators without strata controls or within-stratum partialling had large MSE. The **Plin**, **GO**, and **ToM** estimators had similar MSE across model specifications. These generic methods performed the best in regimes with large n , small $\dim(\psi)$, and nonlinear $E[h|\psi]$. In these cases, the gap $\gamma_L - \gamma^*$ between the suboptimal Lin coefficient and optimal coefficient γ^* domi-

TABLE 2. Properties of inference.

	No strata controls								Strata controls $z(\psi)$						
	Model	Unadj	Naive	Lin	FE	Plin	GO	ToM	Naive	Lin	FE	Plin	GO	ToM	Ad
% Δ CI Length vs. Unadj	1	0	17	11	-5	-5	-5	-5	-49	-50	-34	-29	-26	-29	-50
	2	0	18	10	-3	-4	-4	-4	-33	-41	-25	-25	-22	-25	-41
	3	0	16	16	-6	-6	-6	-6	-36	-36	-24	-24	-24	-24	-36
	4	0	-46	-43	-42	-46	-46	-46	-50	-55	-22	-31	-26	-30	-55
	5	0	-44	-44	-49	-49	-49	-49	-56	-56	-34	-34	-31	-34	-56
	6	0	16	16	-12	-12	-12	-12	-59	-59	-35	-35	-35	-35	-59
Coverage (Exact)	1	0.95	0.95	0.95	0.96	0.96	0.96	0.96	0.95	0.95	0.96	0.96	0.95	0.96	0.95
	2	0.95	0.95	0.95	0.95	0.96	0.96	0.96	0.95	0.96	0.95	0.96	0.95	0.96	0.96
	3	0.95	0.95	0.95	0.96	0.96	0.96	0.96	0.96	0.96	0.96	0.96	0.96	0.96	0.96
	4	0.95	0.95	0.94	0.96	0.95	0.95	0.95	0.95	0.95	0.96	0.96	0.96	0.96	0.95
	5	0.94	0.95	0.95	0.96	0.96	0.96	0.96	0.95	0.95	0.96	0.96	0.97	0.96	0.95
	6	0.95	0.95	0.95	0.97	0.97	0.97	0.97	0.96	0.96	0.97	0.97	0.96	0.97	0.96
Coverage (EHW)	1	0.99	0.95	0.96	0.97	0.99			0.99	0.98	0.99	0.97			
	2	0.99	0.95	0.95	0.94	0.99			0.98	0.93	0.98	0.96			
	3	1.00	0.96	0.95	0.96	1.00			0.99	0.93	0.98	0.97			
	4	0.99	0.99	0.90	0.98	1.00			0.97	0.68	0.98	0.97			
	5	0.99	0.97	0.90	0.98	1.00			0.96	0.65	0.99	0.99			
	6	1.00	0.97	0.96	0.96	1.00			0.99	0.97	1.00	0.99			

nates the additional variability $\text{Var}(\hat{\gamma}^*) > \text{Var}(\hat{\gamma}_L)$ required to estimate γ^* (this variability increases with $\text{dim}(\psi)$). For example, **Plin** with $z(\psi)$ controls performs the best when $(n, \text{dim}(\psi)) = (1200, 2)$, but **Lin** with $z(\psi)$ controls is much better when $\text{dim}(\psi) = 5$. The **Ad** estimator used a variance pretest to choose between **Plin** and **Lin** (including $z(\psi)$ controls), allowing it to perform well in both regimes.

Table 2 reports finite-sample efficiency and coverage properties of the asymptotically exact inference methods developed in Section 4. We let $n = 1200$ and $\text{dim}(\psi) = 5$. The first panel shows % change in confidence interval length relative to unadjusted estimation. All confidence intervals are computed using the method in Theorem 4.1. We see that the relative efficiency of different estimators are reflected by our inference methods. In particular, asymptotically exact inference allows researchers to report shorter confidence intervals when a more efficient adjustment method is used. In the second panel, we show coverage probabilities for our asymptotically exact confidence interval across a range of linearly adjusted estimators. The final panel shows coverage probabilities for confidence intervals based on the usual HC2 variance estimator, where applicable. The HC2-based confidence intervals significantly overcover.

6. EMPIRICAL APPLICATION

In this section, we apply our methods to the experiment in Baysan (2022a),⁷ who estimates the effect of a political information campaign on support for a 2017 Turkish refer-

⁷The data is available from Baysan (2022b).

TABLE 3. Empirical results.

Model	No strata controls							Strata controls $z(\psi)$					
	Unadj	Naive	Lin	FE	Plin	GO	ToM	Naive	Lin	FE	Plin	GO	ToM
$\hat{\theta}_{adj}$	-0.0054	0.0040	0.0047	0.0041	0.0021	0.0034	0.0021	0.0041	0.0040	0.0038	0.0037	0.0031	0.0019
SE	0.0088	0.0074	0.0074	0.0074	0.0078	0.0077	0.0081	0.0074	0.0073	0.0077	0.0076	0.0078	0.0083
HC2	0.0155	0.0075	0.0073	0.0075	0.0149			0.0075	0.0070	0.0736	0.0071		

endum removing checks and balances on executive power. The campaign was administered by the opposition Republican People’s Party (CHP), who opposed the referendum. Randomization was performed at the neighborhood level, stratified on quartiles of CHP vote share in the previous 2015 elections. The main outcome is the “No” vote share in the 2017 referendum.⁸ Due to the cost of administering the campaign, $p = 2/11$ out of $n = 550$ total neighborhoods were treated. In the original analysis, Baysan (2022a) performed noninteracted covariate adjustment (Theorem A.4) for $h(X) =$ number of registered voters, number of valid votes, number of votes for the CHP in 2015, CHP vote share in 2015, voter turnout, and CHP vote share quartile fixed effects.

In the first block of Table 3, we replicate the neighborhood-level analysis of Baysan (2022a). $\hat{\theta}_{adj}$ is the point estimate from each adjustment strategy, SE is the asymptotically exact standard error from Section 4, and EHW is the usual robust standard error (HC2). Estimates in the “strata controls $z(\psi)$ ” section include quartile fixed effects, while the leftmost section does not. The results in Section 3.3 show that Lin adjustment with quartile fixed effects is efficient in this case, and indeed this has the smallest estimated standard error. The generic efficient estimators have slightly larger SE. The asymptotically exact standard errors from Section 4 are generally similar to or smaller than EHW, except for the Lin, FE, and Plin estimators with $z(\psi)$ controls. However, our simulation also showed that EHW standard errors may severely undercover in these cases.⁹

Overall, changing the adjustment method did not have an economically meaningful effect on the conclusions of the study, and we recover the null effect of Baysan (2022a) in all cases. The covariate $h_k =$ “CHP vote share in 2015” is highly predictive of $Y =$ “CHP vote share in 2017,” so adjusting for this variable ex post provides a modest variance reduction even after stratifying on 2015 vote share quartiles. However, the estimated optimal coefficient $\gamma_k^* \approx 0.27$ and Lin coefficient $\gamma_{L,k} \approx 0.31$ are quite similar, so (inefficient) Lin adjustment still performs quite well. The other covariates such as $h_j =$ “voter turnout” are very weak predictors of outcomes, so changing the adjustment coefficient on these variables does not matter much.

Next, we ask how each estimator would have performed in the experiment in Baysan (2022a) under counterfactual randomization procedures, such as fine stratification.¹⁰ To

⁸Baysan (2022a) estimates effects of the campaign on vote share at both the ballot box and neighborhood level. We focus on the neighborhood level effects.

⁹We also note that Bai, Tabord-Meehan, and Liu (2024) have found the EHW standard error from a linear regression with block fixed effects to be potentially invalid in a related problem.

¹⁰Algorithms and inference methods for fine stratification with $p \neq 1/2$ have only been developed recently, for example, Bai (2022) and Cytrynbaum (2023).

TABLE 4. Simulated designs.

Model	No strata controls						Strata controls $z(\psi)$						
	Unadj	Naive	Lin	FE	Plin	GO	ToM	Naive	Lin	FE	Plin	GO	ToM
Coarse	Est	0.0000	0.0001	0.0003	0.0001	0.0002	0.0003	0.0001	0.0004	0.0001	0.0003	0.0002	0.0000
	SE	0.0085	0.0076	0.0075	0.0075	0.0077	0.0078	0.0076	0.0074	0.0078	0.0076	0.0079	0.0079
	HC2	0.0144	0.0078	0.0078	0.0077	0.0141		0.0080	0.0077	0.0736	0.0080		
Fine	Est	-0.0001	0.0000	0.0000	0.0000	0.0004	-0.0001	0.0000	0.0002	0.0000	0.0002	0.0004	-0.0001
	SE	0.0077	0.0081	0.0080	0.0076	0.0077	0.0077	0.0075	0.0075	0.0075	0.0075	0.0077	0.0077
	HC2	0.0144	0.0141	0.0142	0.0078	0.0145		0.0078	0.0078	0.0733	0.0078		
Fine $p = 1/2$	Est	-0.0001	-0.0001	-0.0001	0.0001	0.0001	0.0001	0.0000	-0.0001	0.0001	0.0001	0.0001	0.0001
	SE	0.0072	0.0073	0.0073	0.0070	0.0070	0.0070	0.0066	0.0066	0.0066	0.0066	0.0066	0.0066
	HC2	0.0113	0.0111	0.0111	0.0059	0.0114		0.0060	0.0060	0.0565	0.0061		
Fine $\dim(\psi) = 3$	Est	0.0000	0.0001	0.0002	0.0001	0.0002	0.0003	0.0001	0.0002	0.0000	0.0002	0.0001	0.0001
	SE	0.0155	0.0155	0.0155	0.0155	0.0155	0.0155	0.0146	0.0146	0.0146	0.0146	0.0147	0.0147
	HC2	0.0145	0.0145	0.0145	0.0089	0.0145		0.0079	0.0078	0.0740	0.0078		

do so, we follow the nonparametric imputation strategy in Bai (2022), defining potential outcomes $\widehat{Y}_i(d) = Y_i$ if $D_i = d$ and matching imputation $\widehat{Y}_i(d) = Y_{j(i)}(d)$ with $j(i) = \operatorname{argmin}_{j:D_j=d} |X_i - X_j|_2$ if $D_i \neq d$. We let the matching variables X_i include all controls used in the analysis of Baysan (2022a). Given the imputed data $(X_i, \widehat{Y}_i(0), \widehat{Y}_i(1))_{i=1}^n$, we do the following simulation exercise: (1) draw treatment assignments $D_{1:n} \sim \operatorname{Loc}(\psi, p)$, (2) reveal outcomes $\widehat{Y}_i = \widehat{Y}_i(D_i)$, and (3) form each estimator $\widehat{\theta}_{\text{adj}}$. We report average point estimates and standard errors over $N = 2000$ Monte Carlo repetitions of this procedure.

The first block of Table 4 uses this imputation procedure to reproduce the empirical results in Table 3, stratifying by quartiles of CHP vote share and adjusting for exactly the same covariates. The standard errors are very similar to those in the empirical analysis, which provides some validation for this imputation exercise. In the second block of Table 4, we simulate a design with fine stratification on 2015 CHP vote share, rather than just stratifying by quartiles of the vote share as in Baysan (2022a). We used a matched 11-tuples design, letting $D_{1:n} \sim \operatorname{Loc}(\psi, p)$ for $p = 2/11$ and $\psi = (\text{2015 CHP vote share})$. Covariates $h(X)$ are as above, with $z(\psi) = \psi$. In the third block, we simulate a matched pairs design $D_{1:n} \sim \operatorname{Loc}(\psi, 1/2)$. Note that $p = 1/2$ was infeasible in the original experiment due to the high cost of treatment. The last block uses the design $D_{1:n} \sim \operatorname{Loc}(\psi_{\text{alt}}, p)$ for $\psi_{\text{alt}} = (\text{CHP vote share, Num. of registered voters, Num. of valid votes})$, $p = 2/11$, and covariates $h = \text{Turnout}$.

We make some brief observations about this simulation exercise. First, note that the Naive and Lin adjustment are strictly less efficient than unadjusted estimation under simulated fine stratification, consistent with Theorems 3.2 and Section A.4. Lin and partialled Lin with $z(\psi)$ controls are the most efficient. Adjustment for extra covariates h does not significantly improve efficiency relative to the baseline efficiency gain from finely stratifying on ψ and adjusting for $z(\psi)$ ex post. Using a matched pairs design $p = 1/2$ improves efficiency, though the improvement is small considering that this design would require providing the information campaign to 175 extra neighborhoods. Finally, fine stratification on ψ_{alt} significantly reduces efficiency. This is because the extra covariates are not very predictive of outcomes, but stratifying on these covariates force us to use worse matches on the important covariate $\psi = (\text{2015 CHP vote share})$.

7. DISCUSSION AND RECOMMENDATIONS FOR PRACTICE

Stratified randomization and covariate adjustment are both commonly used in the design and analysis of experiments. In general, experimenters should stratify on a few variables $\psi(X)$ expected to be most predictive of outcomes at design time, and plan to adjust for imbalances in the remaining covariates $h(X)$ ex post, as discussed in Section 3.2. Our analysis showed that under stratified randomization, the usual regression adjusted estimators can be inefficient. Motivated by this, we provide feasible alternatives that are asymptotically optimal in the class of linearly adjusted estimators. We conclude by giving some recommendations for empirical practice based on the theory, simulations, and empirical results above.

We recommend that applied researchers use either (1) the Lin estimator with parametric strata controls $z(\psi)$ (e.g., $z(\psi) = \psi$) or (2) the partialled Lin estimator with parametric controls $z(\psi)$, since these estimators performed the best across our simulations and empirical application. Lin with parametric controls $z(\psi)$ is efficient under a rich covariates condition (Section 3.3), while partialled Lin is generically efficient (Section 3.5). Both estimators are robust to treatment effect heterogeneity, while the strata fixed-effects estimator (Theorem 3.15) is not unless $p = 1/2$.

In our simulations, partialled Lin had good finite-sample performance in regimes where n was large relative to $\dim(\psi)$, especially when $E[h|\psi]$ was very nonlinear. Lin with $z(\psi) = \psi$ controls performed better when $\dim(\psi)$ was large relative to n , or if $E[h|\psi]$ was approximately linear. To decide which regime we are in, we suggest model selection using a variance pretest, choosing Lin if $\widehat{V}(\widehat{\gamma}_L) \leq \widehat{V}(\widehat{\gamma}_{PL})$ and partialled Lin otherwise. This adaptive estimator (**Ad** in Section 6) was efficient in both regimes and had good coverage properties. We leave a more general study of such post model-selection estimators in this context to future work.

Regardless of the adjustment strategy, we recommend using the asymptotically exact confidence intervals provided in Section 4. Our simulations showed close to nominal coverage for these confidence intervals across all considered estimators. By contrast, confidence intervals based on the HC2 robust variance estimator often had significant overcoverage.

REFERENCES

- Abadie, Alberto and Guido W. Imbens (2008), “Estimation of the conditional variance in paired experiments.” *Annales d’Economie et de Statistique*, 175–187. [0973]
- Armstrong, Tim (2022), “Asymptotic efficiency bounds for a class of experimental designs.” arXiv preprint. [arXiv:2205.02726](https://arxiv.org/abs/2205.02726). [0976, 0979]
- Bai, Yuehao (2022), “Optimality of matched-pair designs in randomized controlled trials.” *American Economic Review*. [0993, 0995]
- Bai, Yuehao, Liang Jiang, Joseph P. Romano, Azeem M. Shaikh, and Yichong Zhang (2024b), “Covariate adjustment in experiments with matched pairs.” *Journal of Econometrics*. [0973]
- Bai, Yuehao, Joseph P. Romano, and Azeem M. Shaikh (2021), “Inference in experiments with matched pairs.” *Journal of the American Statistical Association*. [0972, 0973, 0988]
- Bai, Yuehao, Max Tabord-Meehan, and Jizhou Liu (2024), “Inference for matched tuples and fully blocked factorial designs.” *Quantitative Economics*. [0993]
- Baysan, Ceren (2022a), “Persistent polarizing effects of persuasion: Experimental evidence from Turkey.” *American Economic Review*. [0972, 0974, 0992, 0993, 0995]
- Baysan, Ceren (2022b), *Data and Code for: Persistent Polarizing Effects of Persuasion: Experimental Evidence From Turkey*. American Economic Association, Nashville, TN. Ann Arbor, MI: Inter-university Consortium for Political and Social Research, 2022-10-19. [0992]

Bugni, Federico A., Ivan A. Canay, and Azeem M. Shaikh (2018), “Inference under covariate-adaptive randomization.” *Journal of the American Statistical Association*. [0973]

Chang, Haoge (2023), “Design-based estimation theory for complex experiments.” arXiv preprint. arXiv:2311.06891. [0973]

Cytrynbaum, Max (2022), “Designing representative and balanced experiments by local randomization.” arXiv preprint. arXiv:2111.08157v1. [0973, 0974, 0975]

Cytrynbaum, Max (2023), “Optimal stratification of survey experiments.” arXiv preprint. arXiv:2111.08157. [0972, 0973, 0974, 0980, 0981, 0989, 0993]

Cytrynbaum, Max (2024), “Supplement to ‘Covariate adjustment in stratified experiments.’” *Quantitative Economics Supplemental Material*, 15, <https://doi.org/10.3982/QE2475>. [0973]

Derigs, Ulrich (1988), “Solving non-bipartite matching problems via shortest path techniques.” *Annals of Operations Research*, 13, 225–261. [0989]

Fogarty, Colin B. (2018), “Regression-assisted inference for the average treatment effect in paired experiments.” *Biometrika*, 105 (4). [0973, 0987]

Hahn, Jinyong (1998), “On the role of the propensity score in efficient semiparametric estimation of average treatment effects.” *Econometrica*. [0976]

Imbens, Guido and Joshua D. Angrist (1994), “Identification and estimation of local average treatment effects.” *Econometrica*, 62. [0973, 0980]

Imbens, Guido W. and Donald B. Rubin (2015), *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press. [0972, 0983, 0986, 0987]

Ke, Zhu, Hanzhong Liu, and Yuehan Yang (2024), “Design-based theory for lasso adjustment in randomized block experiments with a general blocking scheme.” arXiv preprint. arXiv:2109.11271. [0973, 0986]

Leeb, Hannes and Benedikt M. Pötscher (2005), “Model selection and inference: Facts and fiction.” *Econometric Theory*. [0990]

Lin, Winston (2013), “Agnostic notes on regression adjustments to experimental data: Reexamining Freedman’s critique.” *The Annals of Applied Statistics*, 7 (1), 295–318. [0971, 0972, 0976, 0987]

Liu, Hanzhong and Yuehan Yang (2020), “Regression-adjusted average treatment effect estimates in stratified randomized experiments.” *Biometrika*. [0973]

Lu, Xin and Hanzhong Liu (2024), “Tyranny-of-the-minority regression adjustment in randomized experiments.” *Journal of the American Statistical Association*. [0973, 0987]

Negi, Akanksha and Jeffrey M. Wooldridge (2021), “Revisiting regression adjustment in experiments with heterogeneous treatment effects.” *Econometric Reviews*. [0971, 0976]

Reluga, Katarzyna, Ting Ye, and Qingyuan Zhao (2024), “A unified analysis of regression adjustment in randomized experiments.” *Electronic Journal of Statistics*. [0973]

Robins, James M. and Andrea Rotnitzky (1995), “Semiparametric efficiency in multivariate regression models with missing data.” *Journal of the American Statistical Association*, 90, 122–129. [0980]

Robinson, Peter M. (1988), “Root-n-consistent semiparametric regression.” *Econometrica*, 56 (4). [0972]

Ting, Ye, Jun Shao, Yanyao Yi, and Qingyuan Zhao (2022), “Toward better practice of covariate adjustment in analyzing randomized clinical trials.” *Journal of the American Statistical Association*, 00 (0). [0973]

Wang, Xinhe, Tingyu Wang, and Hanzhong Liu (2021), “Rerandomization in stratified randomized experiments.” *Journal of the American Statistical Association*. Working Paper. [0973]

Wei, Ma, Fuyi Tu, and Hanzhong Liu (2022), “Regression analysis for covariate-adaptive randomization: A robust and efficient inference perspective.” *Statistics in Medicine*. [0973]

Co-editor Stéphane Bonhomme handled this manuscript.

Manuscript received 19 September, 2023; final version accepted 19 July, 2024; available online 19 July, 2024.

The replication package for this paper is available at <https://doi.org/10.5281/zenodo.12674972>. The Journal checked the data and codes included in the package for their ability to reproduce the results in the paper and approved online appendices.