Design-Based Finite-Sample Analysis for Regression Adjustment

Dogyoon Song
University of California, Davis

Abstract

In randomized experiments, regression adjustment can improve the precision of average treatment effect (ATE) estimation using covariates without requiring a correctly specified outcome model. Although well studied in low-dimensional settings, its behavior in high-dimensional regimes, where the number of covariates p may exceed the number of observations n, remains underexplored. Moreover, existing analyses are largely asymptotic, providing limited guidance for finite-sample inference. We develop a design-based, non-asymptotic framework for analyzing the regression-adjusted ATE estimator under complete randomization. This yields finite-sample-valid confidence intervals with explicit, instance-adaptive widths, even when $p > n$. While these intervals rely on oracle (population-level) quantities, we also outline data-driven envelopes computable from observed data. Our approach hinges on a refined swap sensitivity analysis of an estimator: stochastic fluctuation is controlled via a variance-adaptive Doob martingale and Freedman's inequality, and design bias is bounded by Stein's method of exchangeable pairs. The analysis elucidates how covariate geometry governs concentration and bias of the adjusted estimator, suggesting when and how regression adjustment can be effective.

1 INTRODUCTION

Randomized experiments are a central tool for causal inference across scientific domains; see, e.g., Fisher [1935], Neyman [1923], Kempthorne [1952] for foundational developments and Imbens and Rubin [2015], Rosenberger and Lachin [2015], Ding [2024] for modern textbook treatments. In a treatment–control experiment with units $i \in [n]$ and potential outcomes $\{(y_i^{(1)}, y_i^{(0)})\}_{i=1}^n$, the average treatment effect (ATE) $\tau = n^{-1}\sum_{i=1}^n (y_i^{(1)} - y_i^{(0)})$ is often the target of estimation. Two inferential viewpoints are common. Following Neyman [1923], the *design-based* (finite-population) perspective treats the potential outcomes as fixed, with randomness solely arising from the experimenter's treatment assignment [Aronow et al., 2014, Dasgupta et al., 2015, Fogarty, 2018, Mukerjee et al., 2018, Li and Ding, 2020]. By contrast, the superpopulation view models units as i.i.d. draws from a hypothetical infinite population [Tsiatis et al., 2008, Cattaneo et al., 2018]. We adopt the design-based, finite-population perspective to foreground and study the role of randomization without positing an outcome-generating model.

A natural baseline for ATE estimation is the difference-in-means (DiM) estimator $\hat{\tau}_{\mathrm{DiM}}$, which is unbiased under complete randomization with a known variance formula [Neyman, 1923]. When pre-treatment covariates $x_i \in \mathbb{R}^p$ are predictive of outcomes, regression adjustment (RA) can improve precision by incorporating them. Classical analysis of covariance (ANCOVA) often assumes a constant treatment effect, under which the ordinary least squares (OLS) of Y_i on (T_i, x_i) with a treatment indicator T_i is a traditional approach [Fisher, 1935, Kempthorne, 1952, Hinkelmann and Kempthorne, 2007]. Freedman [2008] highlighted three concerns: (i) RA can be less efficient than DiM under heterogeneity; (ii) usual OLS standard errors can be inconsistent; and (iii) the OLS estimator has $O_p(n^{-1})$ bias. In fixed dimensions these issues are resolved by including all treatment–covariate interactions (equivalently, two arm-wise OLS fits): Lin [2013] shows the resulting estimator is consistent, asymptotically normal, and at least as efficient as DiM; Chang et al. [2024] provide exact bias corrections with and without interaction.

Modern experiments often entail many covariates, and fixed-p asymptotics may fail to capture the behavior of RA. Recent work extends design-based asymp-

Proceedings of the 29th International Conference on Artificial Intelligence and Statistics (AISTATS) 2026, Tangier, Morocco. PMLR: Volume 300. Copyright 2026 by the author(s).

totics to increasing-p regimes, e.g., asymptotic normality up to $p = o(n^{1/2})$, and then up to $p = o(n^{2/3})$ with a bias correction [Lei and Ding, 2021], or further up to $p = o(n^{3/4})$ with cross-fitting [Chiang et al., 2023]. A complementary line of work considers high-dimensional settings and provides guarantees under sparsity or semiparametric structure [Bloniarz et al., 2016, Wager et al., 2016, Chernozhukov et al., 2018]. Nevertheless, a *design-based*, *finite-sample* theory for RA that remains valid in high-dimensional and interpolating regimes—without additional ad hoc techniques or modeling assumptions—remains scarce. Moreover, practitioners typically face a single finite dataset, so asymptotic assurances alone may not offer actionable uncertainty quantification.

Contributions. In this paper, we develop a design-based, non-asymptotic framework for analyzing ATE estimators under complete randomization. Our contributions in this paper are four-fold.

(i) *Design concentration and bias bounds* (Section 3). Viewing an ATE estimator as a function $f(\mathcal{S})$ of the treatment set $\mathcal{S}$, we derive a *variance-adaptive* concentration inequality using swap-sensitivity analysis coupled with a Doob martingale and Freedman's inequality. We also obtain a *design-bias* upper bound via Stein's method of exchangeable pairs on the Johnson graph. Both results apply to general set functions beyond ATE estimators.

(ii) *Finite-sample-valid inference* (Section 4.1). Combining these bounds yields an instance-adaptive error radius $\varepsilon_{\mathcal{S}}$ such that, for any $\delta \in (0, 1)$,

$$\Pr\big(|\hat{\tau}(\mathcal{S}) - \tau| > \varepsilon_{\mathcal{S}}\big) \leq \delta,$$

producing finite-sample-valid $(1 - \delta)$ confidence intervals under complete randomization, without requiring asymptotic limits (Theorem 5).

(iii) *Oracle swap sensitivity and covariate Q-geometry* (Section 4.2). Specializing to OLS adjustment, we derive a closed-form *paired deletion–insertion* identity (Theorem 7) for the *oracle* one-swap change $\Delta_{ij}\hat{\tau}_{\mathrm{OLS}}(\mathcal{S})$ with respect to exchanging units $i \leftrightarrow j$ by applying rank-one pseudoinverse updates to the ratio-of-quadratics representation of the intercepts (Proposition 1). The resulting formulas suggest that the covariate Q-geometry governs sensitivity via leverages and standardized residual energy.

(iv) *Toward data-driven proxies* (Appendix A). We isolate the observable Q-geometric factors that enter these oracle envelopes and discuss how they may serve as practical proxies; establishing fully data-driven insertion envelopes is left as future work.

Related Work. Regression adjustment, dating back to Fisher's ANCOVA, has long been used to improve precision in randomized experiments by leveraging pre-treatment covariates [Fisher, 1935, Kempthorne, 1952, Imbens and Rubin, 2015]. Freedman [2008] criticized pitfalls of RA (efficiency loss under heterogeneity, inconsistent OLS standard errors, and $O_p(n^{-1})$ bias), while Lin [2013] advocated fitting all treatment–covariate interactions (equivalently, arm-wise OLS with an intercept) to achieve asymptotic normality and non-inferiority to the difference-in-means (DiM) baseline. Recent refinements provide exact bias corrections and sharpen inference in complete randomization [Chang et al., 2024, Fogarty, 2018].

Beyond the classical fixed-p regime, recent work establishes design-based asymptotics for RA as p grows with n, including asymptotic normality up to $p = o(n^{1/2})$, extensions to $p = o(n^{2/3})$ with a degree-0 bias correction, and to $p = o(n^{3/4})$ via cross-fitting, under leverage regularity conditions [Lei and Ding, 2021, Chiang et al., 2023]. This line of work is further extended to establish asymptotic normality up to $p = o(n^{\frac{d+2}{d+3}})$ with degree-d corrections for all $d \geq 0$ [Song, 2025]. High-dimensional methods under sparsity or semiparametric structure (e.g., lasso and double ML) also provide complementary guarantees, albeit typically in superpopulation or structured settings [Bloniarz et al., 2016, Chernozhukov et al., 2018]. However, finite-sample, design-based guarantees that remain valid when $p \geq n$, without additional modeling assumptions, are still scarce; this paper aims to bridge this gap.

For the analysis, we use Freedman's inequality [Freedman, 1975] together with a Doob martingale tailored to sampling without replacement, to obtain a variance-adaptive concentration bound. We bound the design bias via Stein's method of exchangeable pairs on the Johnson graph of treatment assignment sets [Ross, 2011, Chen et al., 2010, Delsarte, 1973, Levin and Peres, 2017]. For OLS-RA, we derive exact swap sensitivity formulas from rank-one pseudoinverse updates applied to a ratio-of-quadratics representation of the OLS intercepts, exposing the associated covariate geometry [Greville, 1960, Ben-Israel and Greville, 2003].

Organization. Section 2 introduces the setup and design-based preliminaries for ATE estimation and RA. Section 3 decomposes the estimation error into stochastic fluctuation and design bias, yielding finite-sample bounds for each (Propositions 3 and 4). Section 4 establishes an oracle finite-sample-valid inferential guarantee (Theorem 5); Section 4.2 then specializes to OLS-RA and examines how covariate geometry affects RA. Section 5 and Appendix D report numerical experiments, and Section 6 concludes.

2 SETUP AND BACKGROUND

Notation. Let $\mathbb{R}$ denote the set of real numbers, and $\mathbb{Z}$ the set of integers. Let $\mathbb{Z}_{\geq 0} := \{z \in \mathbb{Z} : z \geq 0\}$ and $\mathbb{Z}_{>0} := \{z \in \mathbb{Z} : z > 0\}$. For $n \in \mathbb{Z}_{>0}$, write $[n] := \{1, \ldots, n\}$. A vector $v \in \mathbb{R}^n$ is identified with its column vector representation $v \in \mathbb{R}^{n\times 1}$. For a matrix M, $M^\top$ denotes its transpose, and M^{-1} its inverse (when it exists). For readability, we use calligraphic letters to denote sets (e.g., $\mathcal{S}$). We write $\mathbf{1}_n := (1, \cdots, 1) \in \mathbb{R}^n$ and $\mathbb{I}_E$ for the indicator that equals 1 if and only if E is true (and 0 otherwise).

2.1 Average Treatment Effect Estimation

Finite Population and Potential Outcomes. We consider a finite population of n units. Each unit $i \in [n]$ has two potential outcomes $y_i^{(\alpha)}$ for $\alpha \in \{0, 1\}$, where $y_i^{(1)}$ is the treated outcome and $y_i^{(0)}$ is the control outcome. Collect $y^{(\alpha)} = (y_1^{(\alpha)}, \ldots, y_n^{(\alpha)}) \in \mathbb{R}^n$.

Randomized Experiments. A completely randomized design selects a treatment set $\mathcal{S}_1 \subset [n]$ of a fixed size n_1 uniformly at random; the control set is $\mathcal{S}_0 := [n] \setminus \mathcal{S}_1$, with $n_0 = |\mathcal{S}_0| = n - n_1$. Let $T_i := \mathbb{I}\{i \in \mathcal{S}_1\} \in \{0, 1\}$ denote the treatment indicator. Then the observed outcome is

$$Y_i = T_i y_i^{(1)} + (1 - T_i) y_i^{(0)}. \tag{1}$$

Randomness arises *only* from the treatment assignment (that is, $y^{(1)}, y^{(0)}$ are deterministic).

Remark 2.1. Note that in our setup, the so-called stable unit treatment value assumption (SUTVA) is implicitly imposed by (1), which posits (i) "no interference" (each unit's potential outcomes do not depend on other units' assignments) and (ii) "consistency" (there are no other versions of treatment) conditions.

For each $\alpha \in \{0, 1\}$, let $S_\alpha \in \{0, 1\}^{n_\alpha \times n}$ be the selection matrix such that $(S_\alpha)_{ij} = 1$ if and only if j is the i-th smallest element in $\mathcal{S}_\alpha$. Define $y_\alpha = S_\alpha y^{(\alpha)}$, the *observed* subvector of $y^{(\alpha)}$ corresponding to $\mathcal{S}_\alpha$.

Average Treatment Effect (ATE) Estimation. The finite-population average treatment effect is

$$\tau = \frac{1}{n} \sum_{i=1}^{n} (y_i^{(1)} - y_i^{(0)}).$$

The difference-in-means (DiM) estimator of ATE is

$$\hat{\tau}_{\text{DiM}} = \frac{1}{n_1} \sum_{i \in \mathcal{S}_1} y_i^{(1)} - \frac{1}{n_0} \sum_{i \in \mathcal{S}_0} y_i^{(0)}. \tag{2}$$

Together with complete randomization, SUTVA makes the finite-population ATE identifiable and the difference-in-means (DiM) estimator unbiased.

2.2 Regression Adjustment

Covariates, Design and Normalization. Each unit i carries pre-treatment covariates $x_i \in \mathbb{R}^p$. The design matrix refers to $X = \begin{bmatrix} x_1 & \cdots & x_n \end{bmatrix}^\top \in \mathbb{R}^{n\times p}$.

Assumption 0. *Suppose that*

1. $\mathbf{1}_n^\top X = 0$, *i.e., columns of* X *are centered, and*
2. $X^\top X = nI_p$, *i.e., columns of* X *are orthogonal and have norm* $\sqrt{n}$.

Remark 2.2. Assumption 0 is only used for Section 2.

For $\alpha \in \{0, 1\}$, let $(\mu_\alpha, \beta_\alpha)$ denote the *population* ordinary least squares (OLS) coefficients from regressing $y^{(\alpha)}$ on $[\mathbf{1}_n, X]$, and define the population residuals

$$r^{(\alpha)} := y^{(\alpha)} - \mu_\alpha \mathbf{1}_n - X\beta_\alpha, \qquad \alpha \in \{0, 1\}. \tag{3}$$

By OLS orthogonality, $\mathbf{1}_n^\top r^{(\alpha)} = 0$ and $X^\top r^{(\alpha)} = 0$. Computing the population OLS coefficients $(\mu_\alpha, \beta_\alpha)$ is only a notional device to aid in our analysis; it requires access to the potential outcome $y^{(\alpha)}$ for the entire population, but we only observe $y_\alpha = S_\alpha y^{(\alpha)}$.

Arm-wise OLS with Intercept. For each $\alpha \in \{0, 1\}$, let $(\hat{\mu}_\alpha, \hat{\beta}_\alpha)$ be the arm-wise OLS coefficients obtained by fitting observed data:

$$(\hat{\mu}_\alpha, \hat{\beta}_\alpha) \in \arg\min_{(\mu,\beta)} \left\{ \frac{1}{n_\alpha} \sum_{i \in \mathcal{S}_\alpha} \left(y_i^{(\alpha)} - \mu - x_i^\top \beta \right)^2 \right\}.$$

For each $\alpha \in \{0, 1\}$, define $X_\alpha = S_\alpha X$ as the submatrix of X that retains only the rows with indices in $\mathcal{S}_\alpha$, and define arm-wise means

$$\bar{x}_\alpha := \frac{1}{n_\alpha} X_\alpha^\top \mathbf{1}_{n_\alpha} \qquad \text{and} \qquad \bar{y}_\alpha := \frac{1}{n_\alpha} \mathbf{1}_{n_\alpha}^\top y_\alpha.$$

Define the arm-wise (row-) centering projection matrix $C_\alpha := I_{n_\alpha} - \frac{1}{n_\alpha} \mathbf{1}_{n_\alpha} \mathbf{1}_{n_\alpha}^\top$ so $C_\alpha \mathbf{1}_{n_\alpha} = 0$. Then $C_\alpha X_\alpha = X_\alpha - \mathbf{1}_{n_\alpha} \bar{x}_\alpha^\top$, and the arm-wise sample covariance

$$\Sigma_\alpha := \frac{1}{n_\alpha} \sum_{i \in \mathcal{S}_\alpha} (x_i - \bar{x}_\alpha)(x_i - \bar{x}_\alpha)^\top = \frac{1}{n_\alpha} X_\alpha^\top C_\alpha X_\alpha.$$

In the classical regime where Σ_α is invertible, standard OLS-with-intercept algebra gives

$$\hat{\beta}_\alpha = \Sigma_\alpha^{-1} \Big(\frac{1}{n_\alpha} X_\alpha^\top C_\alpha y_\alpha \Big), \qquad \hat{\mu}_\alpha = \bar{y}_\alpha - \bar{x}_\alpha^\top \hat{\beta}_\alpha.$$

Substituting $y_\alpha = \mu_\alpha \mathbf{1}_{n_\alpha} + X_\alpha \beta_\alpha + S_\alpha r^{(\alpha)}$ and using $C_\alpha \mathbf{1}_{n_\alpha} = 0$ and $X_\alpha^\top C_\alpha \mathbf{1}_{n_\alpha} = 0$ yields

$$\hat{\beta}_\alpha = \beta_\alpha + \Sigma_\alpha^{-1} \Big(\frac{1}{n_\alpha} X_\alpha^\top C_\alpha S_\alpha r^{(\alpha)} \Big). \tag{4}$$

Finally,

$$\hat{\mu}_\alpha - \mu_\alpha = \underbrace{\frac{1}{n_\alpha} \mathbf{1}_{n_\alpha}^\top S_\alpha r^{(\alpha)}}_{\text{residual mean in arm } \alpha} - \bar{x}_\alpha^\top \big(\hat{\beta}_\alpha - \beta_\alpha \big). \tag{5}$$

Regression Adjustment. Recall the population residuals in (3). Since $\mathbf{1}_n^\top X = 0$ and $\mathbf{1}_n^\top r^{(\alpha)} = 0$,

$$\tau = \frac{1}{n}\mathbf{1}_n^\top y^{(1)} - \frac{1}{n}\mathbf{1}_n^\top y^{(0)} = \mu_1 - \mu_0.$$

The regression-adjusted ATE estimator via OLS (or, the OLS-adjusted ATE estimator) refers to a natural plug-in estimator

$$\hat{\tau}_{\mathrm{OLS}} \coloneqq \hat{\mu}_1 - \hat{\mu}_0. \tag{6}$$

2.3 Asymptotics of Regression Adjustment

For fixed p, OLS adjustment with treatment-covariate interactions (equivalently, arm-wise OLS with intercept as in (6)) is a classical method that is asymptotically non-inferior to DiM under standard regularity. In particular, Lin [2013, Theorem 1] show

$$\sqrt{n}(\hat{\tau}_{\mathrm{OLS}} - \tau) \overset{d}{\Rightarrow} \mathcal{N}\big(0, \sigma_{\mathrm{res}}^2\big), \quad \sigma_{\mathrm{res}}^2 \coloneqq \frac{v_1}{\rho} + \frac{v_0}{1-\rho} - v_\tau$$

where $v_\alpha \coloneqq \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^n \big(r_i^{(\alpha)}\big)^2$ for each $\alpha \in \{0,1\}$, and $v_\tau \coloneqq \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^n \big(r_i^{(1)} - r_i^{(0)}\big)^2$. For p growing with n, Lei and Ding [2021] establish asymptotic normality (AN) for OLS–RA up to $p = o(n^{1/2})$ (modulo logs), and extend AN to $p = o(n^{2/3})$ with de-biasing.

At the core of these classical fixed-p analyses is the decomposition obtained by combining (4)–(5) and $\tau = \mu_1 - \mu_0$:

$$\hat{\tau}_{\mathrm{OLS}} - \tau = \underbrace{\frac{1}{n_1}\mathbf{1}_{n_1}^\top S_1 r^{(1)} - \frac{1}{n_0}\mathbf{1}_{n_0}^\top S_0 r^{(0)}}_{=:\hat{\tau}_{\mathrm{DiM}}^{\mathrm{res}};\ \text{Residual DiM}} - \underbrace{\big(\mathsf{R}_1 - \mathsf{R}_0\big)}_{\text{Remainder}} \tag{7}$$

where the arm-wise remainder for each $\alpha \in \{0,1\}$ is

$$\mathsf{R}_\alpha \coloneqq \bar{x}_\alpha^\top\big(\hat{\beta}_\alpha - \beta_\alpha\big) = \bar{x}_\alpha^\top \Sigma_\alpha^{-1}\Big(\frac{1}{n_\alpha} X_\alpha^\top C_\alpha S_\alpha r^{(\alpha)}\Big).$$

A sampling-without-replacement central limit theorem (e.g., Hájek [1960]) yields $\sqrt{n}\cdot\hat{\tau}_{\mathrm{DiM}}^{\mathrm{res}} \overset{d}{\Rightarrow} \mathcal{N}\big(0, \sigma_{\mathrm{res}}^2\big)$; if the remainder is $o_P(n^{-1/2})$, the AN of $\hat{\tau}_{\mathrm{OLS}} - \tau$ follows. Thus, ensuring $\mathsf{R}_1 - \mathsf{R}_0 = o_P(n^{-1/2})$ provides a sufficient condition for asymptotic normality.

Gaps and Motivation. Despite strong asymptotic results for regression adjustment, several gaps remain.

Existing guarantees based on the residual–DiM decomposition (7) are fundamentally *asymptotic*: they establish pointwise AN and consistency of variance estimators, but do not by themselves furnish *finite-sample, design-based* confidence intervals with explicit, verifiable constants. Moreover, in high-dimensional or interpolating regimes, OLS can nearly interpolate outcomes, making residuals small and the residual–DiM route numerically unstable (or infeasible); then the remainder $\mathsf{R}_1 - \mathsf{R}_0$ need not be negligible without additional control. Alternatives that rely on sparsity, cross-fitting/double ML, or superpopulation modeling provide asymptotic guarantees but again lack finite-sample, design-based error control.

To our knowledge, there remains a paucity of *non-asymptotic*, finite-population results that (a) operate under complete randomization, (b) remain valid when p is comparable to or exceeds the arm sizes, and (c) yield ready-to-use confidence intervals with explicit constants.

3 FINITE-SAMPLE ANALYSIS OF CONCENTRATION AND BIAS

We depart from the residual-DiM decomposition (7) and work directly with the minimum-norm OLS estimator: least squares (LS) with an unpenalized intercept and the minimum-ℓ_2-norm slope among all LS minimizers. This estimator admits a closed-form intercept as a ratio of quadratic forms, valid in both classical and over-parameterized regimes.

Definition 3.1. *Given $X \in \mathbb{R}^{n\times p}$ and $y \in \mathbb{R}^n$, the minimum-norm OLS estimator (with unpenalized intercept) of (X, y), denoted by $\mathsf{OLS}(X, y)$, is the unique minimizer*

$$\begin{aligned} (\hat{\mu}, \hat{\beta}) &\in \operatorname*{arg\,min}_{(\mu,\beta)\in\mathcal{S}_{X,y}} \|\beta\|_2^2, \qquad \textit{where} \\ \mathcal{S}_{X,y} &\coloneqq \operatorname*{arg\,min}_{(\mu,\beta)\in\mathbb{R}\times\mathbb{R}^p} \|y - X\beta - \mu\mathbf{1}_n\|_2^2. \end{aligned} \tag{8}$$

Remark 3.1. When $\mathcal{S}_{X,y}$ is a singleton (e.g., if $n \geq p$ and $\mathbf{1}_n \notin \mathrm{col}(X)$), the formulation (8) reduces to the standard OLS with intercept.

Proposition 1. *Let $X \in \mathbb{R}^{n\times p}$ and $y \in \mathbb{R}^n$, and let $(\hat{\mu}, \hat{\beta}) = \mathsf{OLS}(X, y)$. Then*

$$\hat{\mu} = \frac{\mathbf{1}_n^\top Q_X\, y}{\mathbf{1}_n^\top Q_X\, \mathbf{1}_n}, \tag{9}$$

with

$$Q_X = \begin{cases} M_X \coloneqq I_n - XX^\dagger, & \textit{if } \mathbf{1}_n \notin \mathrm{col}(X), \\ K_X \coloneqq (XX^\top)^\dagger, & \textit{if } \mathbf{1}_n \in \mathrm{col}(X). \end{cases} \tag{10}$$

Remark 3.2. $XX^\dagger$ is the orthogonal projection matrix onto $\mathrm{col}(X)$, so-called the "hat matrix." $XX^\top$ is the Gram matrix of the row vectors $\{x_1, \ldots, x_n\}$. The denominator in (9) is strictly positive, i.e., $\mathbf{1}_n^\top Q_X \mathbf{1}_n > 0$, because Q_X is positive semidefinite and $\mathbf{1}_n \notin \ker(Q_X)$ in both cases of (10). See Appendix B.1 for a proof of Proposition 1.

Recall that $\hat{\tau}_{\mathrm{OLS}} := \hat{\mu}_1 - \hat{\mu}_0$ with arm-wise intercepts, cf. (6). Thus, $\hat{\tau}_{\mathrm{OLS}}$ is a function of the random arm index sets $\mathcal{S}_\alpha$ via $\hat{\mu}_\alpha$, which depend on $\mathcal{S}_\alpha$ through Q_{X_α} and y_α. Intuitively, if $\hat{\tau}_{\mathrm{OLS}}$ is not very sensitive to a change in $\mathcal{S}_\alpha$, then $\hat{\tau}_{\mathrm{OLS}}$ concentrates around $\mathbb{E}[\hat{\tau}_{\mathrm{OLS}}]$, which may still differ from the true ATE τ though.

Section Overview. In what follows, we develop a design-based, non-asymptotic analysis for ATE estimators under complete randomization. Let $n \in \mathbb{Z}_{\geq 0}$, $m \in [n]$, and $f : \binom{[n]}{m} \to \mathbb{R}$ (e.g., $f = \hat{\tau}_{\mathrm{OLS}} - \tau$). We seek a high-probability upper bound for $|f(\mathcal{S})|$. Note

$$|f(\mathcal{S})| \leq |f(\mathcal{S}) - \mathbb{E}f(\mathcal{S})| + |\mathbb{E}f(\mathcal{S})|.$$

Instantiating $f = \hat{\tau}_{\mathrm{OLS}} - \tau$, the first term $f(\mathcal{S}) - \mathbb{E}f(\mathcal{S}) = \hat{\tau} - \mathbb{E}\hat{\tau}$ represents the stochastic fluctuation of the ATE estimator $\hat{\tau}$, and the second term $\mathbb{E}f(\mathcal{S}) = \mathbb{E}\hat{\tau} - \tau$ is the bias of $\hat{\tau}$. In what follows, we describe our approach to control these two separately:

- Section 3.1: a high-probability bound on $\big|f(\mathcal{S}) - \mathbb{E}f(\mathcal{S})\big|$ via concentration based on swap-sensitivity analysis with a variance-adaptive Doob martingale,
- Section 3.2: a deterministic upper bound on $\big|\mathbb{E}f(\mathcal{S})\big|$ via Stein's method of exchangeable pairs.

Together, these yield a high-probability upper bound on the estimation error $\big|\hat{\tau}_{\mathrm{OLS}} - \tau\big|$, enabling finite-sample-valid confidence intervals.

3.1 Concentration via Sensitivity Analysis

We derive a high-probability bound for $|f(\mathcal{S}) - \mathbb{E}f(\mathcal{S})|$ under $\mathcal{S} \sim \mathrm{Unif}\binom{[n]}{m}$.

Swap Sensitivity. Recall the treatment set $\mathcal{S}_1 \subset [n]$ is drawn uniformly at random from $\binom{[n]}{n_1}$, and the control set $\mathcal{S}_0 = [n] \setminus \mathcal{S}_1$. The ATE estimators, such as $\hat{\tau}_{\mathrm{DiM}}$ in (2) and $\hat{\tau}_{\mathrm{OLS}}$ in (6), are functions of the random subset $\mathcal{S}_1$. We quantify the sensitivity of such a random function to a swap operation $(i \leftrightarrow j)$.

Definition 3.2. *Let $n \in \mathbb{Z}_{\geq 0}$, $m \in [n]$, and $f : \binom{[n]}{m} \to \mathbb{R}$. For any $\mathcal{S} \in \binom{[n]}{m}$ and any $(i,j) \in \mathcal{S} \times ([n] \setminus \mathcal{S})$, the* (i,j)-swap sensitivity of f at $\mathcal{S}$ *is*

$$\Delta_{ij} f(\mathcal{S}) := f\big(\mathcal{S}^{(i\leftrightarrow j)}\big) - f(\mathcal{S}) \tag{11}$$

where

$$\mathcal{S}^{(i\leftrightarrow j)} := \big(\mathcal{S} \setminus \{i\}\big) \cup \{j\}.$$

Example 1. *For the difference-in-means estimator,*

$$\Delta_{ij}\hat{\tau}_{\mathrm{DiM}} = \frac{1}{n_1}\big(y_j^{(1)} - y_i^{(1)}\big) + \frac{1}{n_0}\big(y_j^{(0)} - y_i^{(0)}\big).$$

Computing the swap sensitivity of $\hat{\tau}_{\mathrm{OLS}}$ is more involved; see Proposition 6 and Theorem 7 for this.

Naive Bounded Differences Are Suboptimal. A (McDiarmid-type) permutation bounded differences argument controls $|\Delta_{ij} f(\mathcal{S})|$ by a global Lipschitz envelope c_{BD}, yielding tail bounds of the form $\exp\{-\Theta(t^2/(m\, c_{\mathrm{BD}}^2))\}$. This is typically suboptimal because c_{BD} reflects a worst-case range (e.g., $\max_{i,j} |y_i^{(\alpha)} - y_j^{(\alpha)}|$) rather than the variance scale, and it ignores the correlations/cancellations across the swaps that connect realizable sets $\mathcal{S}_1 \leftrightarrow \mathcal{S}_1'$. In contrast, in the asymptotic analysis of the DiM, the dispersion is governed by the variance proxy underlying $\sigma_{\mathrm{res}}^2 = \frac{v_1}{\rho} + \frac{v_0}{1-\rho} - v_\tau$, often much smaller than worst-case ranges. Therefore, a variance-adaptive construction is needed for a useful concentration bound.

Assignment-exposure Martingale. Encode complete randomization by a uniform permutation Π of $[n]$, with $\mathcal{S}_1 = \{\Pi(1), \ldots, \Pi(n_1)\}$. Define the *assignment-exposure filtration and martingale*

$$\mathcal{F}_t := \sigma\big(\Pi(1), \ldots, \Pi(t)\big), \qquad t = 0, 1, \ldots, n_1, \tag{12}$$

$$M_t := \mathbb{E}[f(\mathcal{S}_1) \mid \mathcal{F}_t]. \tag{13}$$

Then M_t is a martingale with respect to $\mathcal{F}_t$, and $D_t := M_t - M_{t-1}$ is a martingale difference sequence. Furthermore, $\mathbb{E}[D_t \mid \mathcal{F}_{t-1}] = 0$ and

$$f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1) = \sum_{t=1}^{n_1} D_t.$$

Define the set of treated indices revealed up to step $t-1$, and the remaining pool at $t-1$: for $t \in [n_1]$,

$$\mathcal{S}_{t-1}^{\mathrm{past}} := \{\Pi(1), \ldots, \Pi(t-1)\}, \qquad \mathcal{R}_{t-1} := [n] \setminus \mathcal{S}_{t-1}^{\mathrm{past}}. \tag{14}$$

Let

$$m_t := |\mathcal{R}_{t-1}| = n - (t-1), \qquad \ell_t := n_1 - (t-1),$$

denote the number of total remaining units and the number of treated units not revealed yet, respectively.

Remark 3.3. Although $\mathcal{F}_t$ reveals treated indices only, conditioning on $\mathcal{F}_{t-1}$ also updates the law of (projected) terminal control set, which is uniform over the n_0-subsets of $\mathcal{R}_{t-1}$. Thus both $\mathbb{E}[\hat{\mu}_1(\mathcal{S}_1) \mid \mathcal{F}_t]$ and $\mathbb{E}[\hat{\mu}_0(\mathcal{S}_0) \mid \mathcal{F}_t]$ evolve with t. For example, for DiM, $\mathbb{E}[\hat{\mu}_0 \mid \mathcal{F}_{t-1}]$ is the average of $y^{(0)}$ over $\mathcal{R}_{t-1}$ and shifts by a hypergeometric update when $I = \Pi(t)$ is revealed.

Next, we express the increment D_t as the average of a single paired swap around a feasible proxy set. Let $I := \Pi(t)$. Let J be uniform on $\mathcal{R}_{t-1} \setminus \{I\}$ and $\mathcal{T}$ uniform on $\binom{\mathcal{R}_{t-1}\setminus\{I,J\}}{\ell_t - 1}$, conditionally on $(\mathcal{F}_{t-1}, I)$. We define the proxy set of $\mathcal{S}_1$ based on information up to $t-1$

$$\mathcal{S}_{t-1}^{\mathrm{prox}}(I, \mathcal{T}) := \mathcal{S}_{t-1}^{\mathrm{past}} \cup \{I\} \cup \mathcal{T}. \tag{15}$$

Proposition 2. *Fix $t \in \{1, \dots, n_1\}$ and let $I = \Pi(t)$. Conditional on $(\mathcal{F}_{t-1}, I)$, draw*

$$J \sim \mathrm{Unif}\big(\mathcal{R}_{t-1} \setminus \{I\}\big), \quad \mathcal{T} \sim \mathrm{Unif}\binom{\mathcal{R}_{t-1} \setminus \{I, J\}}{\ell_t - 1}.$$

Then the martingale increment

$$\begin{aligned} D_t &= \mathbb{E}[f(\mathcal{S}_1) \mid \mathcal{F}_t] - \mathbb{E}[f(\mathcal{S}_1) \mid \mathcal{F}_{t-1}] \\ &= -\alpha_t \cdot \mathbb{E}\Big[\, \Delta_{IJ}\, f\big(\mathcal{S}_{t-1}^{\mathrm{prox}}(I, \mathcal{T})\big) \;\big|\; \mathcal{F}_{t-1},\, I \,\Big], \end{aligned} \tag{16}$$

where $\alpha_t = \frac{m_t - \ell_t}{m_t} = \frac{n_0}{n-t+1} \in (0, 1)$.

In (16), the expectation is taken with respect to $(J, \mathcal{T})$. See Appendix B.2 for the proof of Proposition 2.

Remark 3.4. At step t, conditioned on the past $(\mathcal{S}_{t-1}^{\mathrm{past}})$ and the current candidate $(I = \Pi(t))$, D_t quantifies the average effect of the swap that "replace a uniformly drawn random competitor J by I" inside a uniformly drawn random feasible completion set $\mathcal{T}$.

Concentration via Freedman's Inequality. For each $t = 1, \dots, n_1$ and each $i \in \mathcal{R}_{t-1}$, define a $\mathcal{F}_{t-1}$-measurable function

$$\zeta_t(i) := \mathbb{E}\left[\Delta_{iJ}\, f\left(\mathcal{S}_{t-1}^{\mathrm{prox}}(i, \mathcal{T})\right) \;\middle|\; \mathcal{F}_{t-1}\right], \tag{17}$$

where the expectation is over J and $\mathcal{T}$ as in (16). By Proposition 2, $D_t = -\alpha_t\, \zeta_t(I)$ with $I \mid \mathcal{F}_{t-1} \sim \mathrm{Unif}(\mathcal{R}_{t-1})$. Define (oracle) parameters using (16):

$$v_t^\star := \mathrm{Var}(D_t \mid \mathcal{F}_{t-1}) = \alpha_t^2\, \mathrm{Var}\big(\zeta_t(I)\big), \tag{18}$$

$$r_t^\star := \sup |D_t| \;=\; \alpha_t \max_{i \in \mathcal{R}_{t-1}} |\zeta_t(i)|. \tag{19}$$

By construction, $\zeta_t(\cdot)$ and α_t are $\mathcal{F}_{t-1}$-measurable, and thus, $v_t^\star$ and $r_t^\star$ are $\mathcal{F}_{t-1}$-measurable (predictable).

Applying Freedman's martingale concentration inequality [Freedman, 1975, Thm. 1.6] to the sum $\sum_{t=1}^{n_1} D_t$ yields the following probabilistic tail bound.

Proposition 3. *Let $(M_t, \mathcal{F}_t)$ be the assignment-exposure martingale in (12), (13). Then for any $\varepsilon \geq 0$,*

$$\Pr\big(f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1) \geq \varepsilon\big) \leq \exp\left(-\frac{\varepsilon^2/2}{V^\star + R^\star \varepsilon/3}\right), \tag{20}$$

where $V^\star := \sum_{t=1}^{n_1} v_t^\star$ and $R^\star := \max_{1 \leq t \leq n_1} r_t^\star$ with $(v_t^\star, r_t^\star)$ as in (18), (19).

The bound (20) is self-normalized and uses predictable $(v_t^\star, r_t^\star)$ along the exposure filtration; see Appendix B.3 for a detailed derivation.

Remark 3.5. Applying (20) to $\pm(f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1))$ and taking a union bound yields that for any $\delta \in (0, 1)$,

$$\Pr\left(|f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1)| \leq \sqrt{2V^\star \mathsf{L}_\delta} + \frac{R^\star}{3}\, \mathsf{L}_\delta\right) \geq 1 - \delta, \tag{21}$$

where $\mathsf{L}_\delta := \log \frac{2}{\delta}$.

Discussion. Relative to permutation bounded differences, the bound in (20) is *variance-adaptive*. $V^\star$ aggregates *local* conditional variances of paired swaps along the reveal process, while $R^\star$ is a *local* range. These parameters are *oracle* because they depend on the true per-swap effects $\Delta_{ij} f(\cdot)$. Given full potential-outcome information, they are computable via (18)–(19); for OLS-RA, Section 4.2 and Appendix A describe the corresponding covariate geometry that governs swap sensitivities. The construction is symmetric in treatment/control exposures (ending at n_1 or n_0); both yield valid bounds for the same f, and one may take the smaller radius in practice.

3.2 Bias Control via Exchangeable Pairs

Continuing with a generic function $f : \binom{[n]}{m} \to \mathbb{R}$, we derive an upper bound for the design bias $\big|\mathbb{E}f(\mathcal{S}_1)\big|$ via Stein's method of exchangeable pairs; see Ross [2011] for an introduction, and Chen et al. [2010], Reinert and Röllin [2009] for comprehensive treatments of Stein's method and the exchangeable-pairs framework.

Throughout this section, expectations $\mathbb{E}[\cdot]$ and inner products $\langle f, g\rangle := \mathbb{E}[f(\mathcal{S})g(\mathcal{S})]$ are taken with respect to the uniform law π on $\binom{[n]}{m}$, unless stated otherwise.

Exchangeable Pair on Johnson Graph. For $n \in \mathbb{Z}_{\geq 0}$ and $m \in [n]$, the Johnson graph $J(n, m) = \big(\mathcal{V}_{n,m}, \mathcal{E}_{n,m}\big)$ is an undirected graph with the vertex set $\mathcal{V}_{n,m} = \binom{[n]}{m} = \big\{\mathcal{S} \subseteq [n] : |\mathcal{S}| = m\big\}$ and the edge set $\mathcal{E}_{n,m} = \big\{\{v_1, v_2\} \in \mathcal{V}_{n,m} \times \mathcal{V}_{n,m} : |v_1 \cap v_2| = m-1\big\}$; edges connect sets that differ by a single swap. For $\mathcal{S} \in \mathcal{V}_{n,m}$, let $\mathcal{N}(\mathcal{S})$ denote the set of $m(n-m)$ neighbors of $\mathcal{S}$. Let P denote the Markov kernel (transition matrix) on $J(n, m)$ that jumps from $\mathcal{S}$ to a neighbor $\mathcal{S}' \in \mathcal{N}(\mathcal{S})$ uniformly: $P(\mathcal{S}, \mathcal{S}') = \frac{1}{m(n-m)}$ if $\mathcal{S}' \in \mathcal{N}(\mathcal{S})$ and 0 otherwise. The uniform distribution π is stationary for P, and P is reversible, i.e., $\pi(\mathcal{S})P(\mathcal{S}, \mathcal{S}') = \pi(\mathcal{S}')P(\mathcal{S}', \mathcal{S})$.

Let $\mathcal{S} \sim \mathrm{Unif}\binom{[n]}{m}$, and let $\mathcal{S}'$ be obtained by $\mathcal{S}' \mid \mathcal{S} \sim P(\mathcal{S}, \cdot)$. Then $(\mathcal{S}, \mathcal{S}')$ is exchangeable. Given a function $f : \binom{[n]}{m} \to \mathbb{R}$, since $\mathcal{S}'|\mathcal{S} \sim \mathrm{Unif}\big(\mathcal{N}(\mathcal{S})\big)$,

$$\begin{aligned} \mathbb{E}[f(\mathcal{S}') - f(\mathcal{S}) \mid \mathcal{S}] &= \mathbb{E}[f(\mathcal{S}') \mid \mathcal{S}] - f(\mathcal{S}) \\ &= (Pf)(\mathcal{S}) - f(\mathcal{S}). \end{aligned} \tag{22}$$

Below, we use (22) to derive a bias upper bound.

Markov Chain and Carré du Champ. Consider a Markov chain on $J(n, m)$ with the (discrete-time) generator $L = P - \mathrm{Id}$, so $(Lf)(\mathcal{S}) = (Pf)(\mathcal{S}) - f(\mathcal{S}) = \mathbb{E}[f(\mathcal{S}') - f(\mathcal{S}) \mid \mathcal{S}]$. The carré du champ (Dirichlet form) is defined

$$\Gamma(f)(\mathcal{S}) := \frac{1}{2}\, \mathbb{E}\big[\, (f(\mathcal{S}') - f(\mathcal{S}))^2 \;\big|\; \mathcal{S}\,\big].$$

Note that $\Gamma(f)(\mathcal{S}) \geq 0$ and $\mathbb{E}\Gamma(f) = 0$ if and only if f is constant on $\binom{[n]}{m}$. By Jensen's inequality (equivalently, by Cauchy–Schwarz for reversible kernels), for all $\mathcal{S}$,

$$\begin{aligned} \big|(Lf)(\mathcal{S})\big| &\leq \sqrt{\mathbb{E}\big[(f(\mathcal{S}') - f(\mathcal{S}))^2 \mid \mathcal{S}\big]} \\ &= \sqrt{2\,\Gamma(f)(\mathcal{S})}. \end{aligned} \tag{23}$$

(See e.g., Bakry et al. [2013, Chapter 1] for more background and details on L and Γ.)

Stein Regression. Recall the left-hand side of (22). The exchangeable-pairs "regression with remainder" posits the conditional linearization

$$\mathbb{E}\big[f(\mathcal{S}') - f(\mathcal{S}) \mid \mathcal{S}\big] = -\lambda\, f(\mathcal{S}) + \psi(\mathcal{S})$$

or equivalently, $Lf = -\lambda f + \psi$, where $\lambda \geq 0$ (the "drift") and the "remainder" ψ are to be determined. Taking expectations and using $\mathbb{E}[f(\mathcal{S}') - f(\mathcal{S})] = 0$ (exchangeability) yields the Stein identity $\lambda\,\mathbb{E}f = \mathbb{E}\psi$.

Next, we choose the value of λ to compute (or bound) $\mathbb{E}f$. If $\mathrm{Var}(f) = 0$, then f is constant and $|\mathbb{E}f| = |f(\mathcal{S})|$ for the realized $\mathcal{S}$; this degenerate case is handled separately below. We henceforth assume $\mathrm{Var}(f) > 0$. Let $g \coloneqq f - \mathbb{E}f$ and define

$$\lambda^\star \coloneqq \arg\min_{\lambda \geq 0}\ \mathbb{E}\big(Lf + \lambda g\big)^2 = -\frac{\langle Lf,\, g\rangle}{\langle g,\, g\rangle}, \tag{24}$$

$$\begin{aligned} \psi_c^\star &\coloneqq Lf + \lambda^\star g, \\ \psi_u^\star &\coloneqq Lf + \lambda^\star f. \end{aligned} \tag{25}$$

We use $\psi_c^\star$ to choose λ, and use $\psi_u^\star$ to control $\mathbb{E}f$.

Observe that $\langle \psi_c^\star, g\rangle = 0$. Moreover, since $\mathbb{E}Lf = 0$,

$$\mathbb{E}\psi_u^\star = \mathbb{E}(Lf + \lambda^\star f) = \lambda^\star \mathbb{E}f.$$

By reversibility and the Dirichlet identity, $\lambda^\star > 0$; see (26) below. Applying Cauchy–Schwarz inequality,

$$|\mathbb{E}f| = \frac{|\mathbb{E}\psi_u^\star|}{\lambda^\star} \leq \frac{\sqrt{\mathbb{E}(\psi_u^\star)^2}}{\lambda^\star}.$$

(Oracle) Identification of $\lambda^\star$ and Bias Bound. Reversibility implies the discrete Dirichlet identity

$$\mathbb{E}\,\Gamma(h) = -\langle Lh, h\rangle \qquad \text{for all } h,$$

see, e.g., Bakry et al. [2013, Ch. 1]. Since L annihilates constants, $Lg = Lf$, and $\mathrm{Var}(f) = \langle g, g\rangle$. Moreover, $\mathbb{E}\Gamma(f) = -\langle Lf, f\rangle = -\langle Lg, f\rangle = -\langle Lg, f - \mathbb{E}f\rangle = -\langle Lg, g\rangle$. Using the Rayleigh-quotient form in (24),

$$\lambda^\star = -\frac{\langle Lf,\, g\rangle}{\langle g,\, g\rangle} = -\frac{\langle Lg,\, g\rangle}{\langle g,\, g\rangle} = \frac{\mathbb{E}\Gamma(f)}{\mathrm{Var}(f)} \overset{(*)}{>} 0. \tag{26}$$

Here, (*) holds because $\mathrm{Var}(f) > 0$ implies $\mathbb{E}\Gamma(f) \neq 0$.

By instantiating $f(\mathcal{S}_1) = \hat{\tau}(\mathcal{S}_1) - \tau$, we obtain the following bias bound.

Proposition 4 (Oracle Stein bias bound)**.** *Let $f : \binom{[n]}{m} \to \mathbb{R}$ and let $(\mathcal{S}, \mathcal{S}')$ be the one-swap exchangeable pair on $J(n, m)$. If $\mathrm{Var}(f) > 0$, then, under complete randomization and the kernel P above,*

$$|\mathbb{E}f(\mathcal{S})| \leq \frac{\sqrt{\mathbb{E}\big(Lf + \lambda^\star f\big)^2}}{\lambda^\star}, \qquad \lambda^\star = \frac{\mathbb{E}\Gamma(f)}{\mathrm{Var}(f)}. \tag{27}$$

The right-hand side is an oracle quantity: it requires the full distributional knowledge about f over all treatment assignments. In particular, for $f(\mathcal{S}_1) = \hat{\tau}(\mathcal{S}_1) - \tau$ instantiated as the deviation of an ATE estimator, Proposition 4 yields an oracle bias bound for $\hat{\tau}$. Although Proposition 4 follows from the discussions in the main text so far, we provide a formal proof in Appendix B.4 for completeness.

Remark 3.6. By Poincaré (spectral-gap) inequality, $\mathbb{E}\,\Gamma(h) \geq \mathrm{gap}_{n,m}\,\mathrm{Var}(h)$ for all h, where $\mathrm{gap}_{n,m}$ is the spectral gap of the random walk by P on $J(n, m)$. Consequently, for $\lambda^\star = \mathbb{E}\Gamma(f)/\mathrm{Var}(f)$ chosen above, we have $\lambda^\star \geq \mathrm{gap}_{n,m}$ [Levin and Peres, 2017, Chapter 13]. It is well known that $\mathrm{gap}_{n,m} = 1 - \lambda_2(P) = \frac{1}{m} + \frac{1}{n-m} = \frac{n}{m(n-m)}$; see, e.g., Delsarte [1973].

4 MAIN RESULTS

In this section we apply the machinery from Section 3 to $f(\mathcal{S}) = \hat{\tau}(\mathcal{S}) - \tau$, where τ is the ATE and $\hat{\tau}$ an estimator. Section 4.1 gives an *oracle* finite-sample bound in terms of concentration and Stein parameters. Section 4.2 specializes these quantities to OLS–RA via a closed-form swap-sensitivity identity.

4.1 Oracle Finite-sample Analysis

Let $f(\mathcal{S}) = \hat{\tau}(\mathcal{S}) - \tau$. With $(M_t, \mathcal{F}_t)$ as in (13), recall from (18)–(19) that for $t = 1, \ldots, n_1$,

$$v_t^\star \coloneqq \mathrm{Var}(M_t - M_{t-1} \mid \mathcal{F}_{t-1}), \qquad r_t^\star \coloneqq \sup |M_t - M_{t-1}|. \tag{28}$$

Set the (oracle) concentration parameters

$$V^\star \coloneqq \sum_{t=1}^{n_1} v_t^\star \qquad \text{and} \qquad R^\star \coloneqq \max_{1 \leq t \leq n_1} r_t^\star. \tag{29}$$

With the one-swap kernel P on $J(n, n_1)$,

$$\Gamma(f)(\mathcal{S}) = \frac{1}{2n_1 n_0} \sum_{i \in \mathcal{S}} \sum_{j \notin \mathcal{S}} \big(\Delta_{ij} f(\mathcal{S})\big)^2, \tag{30}$$

where $\mathcal{S}' | \mathcal{S} \sim P(\mathcal{S}, \cdot)$ and $\Delta_{ij} f(\mathcal{S}) \coloneqq f(\mathcal{S}^{(i \leftrightarrow j)}) - f(\mathcal{S})$, cf. (11). Averaging (30) over $\mathcal{S} \sim \mathrm{Unif}\binom{[n]}{n_1}$ gives the population quantity $\mathbb{E}\Gamma(f)$.

Following (27), define the oracle bias parameter

$$B^\star := \frac{\sqrt{\mathbb{E}\big(Lf+\lambda^\star f\big)^2}}{\lambda^\star}, \quad \text{where} \quad \lambda^\star = \frac{\mathbb{E}\Gamma(f)}{\mathrm{Var}(f)}, \tag{31}$$

with the convention that if $\mathrm{Var}(f)=0$ then $B^\star := |\mathbb{E}f|$ (since $\Gamma(f)\equiv 0$ in this case).

Theorem 5 (Oracle finite-sample CI)**.** *Let $\mathcal{S}_1 \sim \mathrm{Unif}\binom{[n]}{n_1}$ and $f(\mathcal{S}) = \hat\tau(\mathcal{S}) - \tau$. For any $\delta\in(0,1)$,*

$$\Pr\left(|f(\mathcal{S}_1)| \le \sqrt{2V^\star \mathsf{L}_\delta} + \frac{R^\star}{3}\mathsf{L}_\delta + B^\star\right) \ge 1-\delta, \tag{32}$$

where $\mathsf{L}_\delta := \log\frac{2}{\delta}$. $V^\star, R^\star, B^\star$ are as in (29) and (31).

Proof sketch of Theorem 5. Observe $|f(\mathcal{S})| \le |f(\mathcal{S}) - \mathbb{E}f(\mathcal{S})| + |\mathbb{E}f(\mathcal{S})|$ and combine Propositions 3 and 4. See Appendix C.1 for a complete proof. □

Remarks on Bias Bound $B^\star$. The quantities in (31) are population-level oracle quantities computable from the full potential-outcome information over the randomization distribution, not from the instantiated data alone. While $\mathbb{E}\Gamma(f)$ determines $\lambda^\star$, the bias bound is governed by the uncentered Stein residual $Lf+\lambda^\star f$. See Example 2 in Appendix A.1 for the DiM illustration. For DiM, $f = \hat\tau_{\mathrm{DiM}} - \tau$ is an exact Johnson-graph eigenfunction, $Lf = -\lambda_{\mathrm{lin}} f$ with $\lambda_{\mathrm{lin}} = \frac{n}{n_1 n_0}$, and $\lambda^\star = \lambda_{\mathrm{lin}}$. Hence $Lf+\lambda^\star f \equiv 0$, so $B^\star = 0$, capturing the design unbiasedness. For estimators beyond DiM, such as OLS–RA, $B^\star$ measures the size of the corresponding uncentered Stein residual, and provides an oracle bound on design bias.

4.2 Swap Sensitivity Analysis for OLS-RA

Since D_t is a conditional average of one-swap sensitivities (Proposition 2), the parameters $V^\star, R^\star$ in (29) reduce to population-level swap-sensitivity analysis. For OLS–RA, the change $\Delta_{ij} f(\mathcal{S}) = f(\mathcal{S}^{(i\leftrightarrow j)}) - f(\mathcal{S})$ admits a convenient algebraic form via paired deletion–insertion and rank-one pseudoinverse updates.

Definition 4.1. *Let $X\in\mathbb{R}^{n\times p}$ and $y^{(1)}, y^{(0)}\in\mathbb{R}^n$. For $\alpha\in\{0,1\}$, $\mathcal{S}\in\binom{[n]}{m}$, $i\in\mathcal{S}$, and $j\notin\mathcal{S}$, define the armwise atomic changes due to deletion/insertion as*

$$\begin{aligned}\Delta_\alpha^{\mathrm{del}}(i;\mathcal{S}) &:= \hat\mu_\alpha\big(\mathcal{S}\setminus\{i\}\big) - \hat\mu_\alpha(\mathcal{S}),\\ \Delta_\alpha^{\mathrm{ins}}(j;\mathcal{S}) &:= \hat\mu_\alpha\big(\mathcal{S}\cup\{j\}\big) - \hat\mu_\alpha(\mathcal{S}),\end{aligned} \tag{33}$$

where $\hat\mu_\alpha(\mathcal{S}) = \frac{\mathbf{1}_m^\top Q_{X_\mathcal{S}}\, y_\mathcal{S}^{(\alpha)}}{\mathbf{1}_m^\top Q_{X_\mathcal{S}}\,\mathbf{1}_m}$, cf. (9).

For OLS–RA, the swap sensitivity of $f(\mathcal{S}_1) = \hat\tau_{\mathrm{OLS}}(\mathcal{S}_1) - \tau$ equals that of $\hat\tau_{\mathrm{OLS}}(\mathcal{S}_1) = \hat\mu_1(\mathcal{S}_1) - \hat\mu_0(\mathcal{S}_0)$ because τ is constant. The one-swap change admits an exact paired deletion-insertion decomposition.

Proposition 6. *For any $\mathcal{S}_1\in\binom{[n]}{n_1}$ and $(i,j)\in\mathcal{S}_1\times\mathcal{S}_0$,*

$$\begin{aligned}\Delta_{ij}\hat\tau_{\mathrm{OLS}}(\mathcal{S}_1) = &\Big[\Delta_1^{\mathrm{del}}(i;\mathcal{S}_1) + \Delta_1^{\mathrm{ins}}(j;\mathcal{S}_1\setminus\{i\})\Big]\\ &- \Big[\Delta_0^{\mathrm{del}}(j;\mathcal{S}_0) + \Delta_0^{\mathrm{ins}}(i;\mathcal{S}_0\setminus\{j\})\Big].\end{aligned} \tag{34}$$

(34) is an algebraic telescoping identity, and no properties of the OLS estimator are needed beyond that $\hat\mu_\alpha(\cdot)$ is well-defined. See Appendix C.2 for a proof.

By Proposition 1, the atomic deletion and insertion changes admit explicit, computable formulas through one-row updates of the corresponding arm-wise Q-matrices. To state the result, we introduce notation. For a positive semidefinite matrix Q, define the Q-inner product and Q-norm as follows: for vectors u, v,

$$\langle u,v\rangle_Q := u^\top Q v, \qquad \|u\|_Q := \sqrt{u^\top Q u}.$$

With $X\in\mathbb{R}^{n\times p}$ and $y^{(1)}, y^{(0)}\in\mathbb{R}^n$ fixed, we let $X_\mathcal{S}\in\mathbb{R}^{|\mathcal{S}|\times p}$ and $y_\mathcal{S}^{(\alpha)}\in\mathbb{R}^{|\mathcal{S}|}$ denote the row-submatrix/subvector of X, $y^{(\alpha)}$ with indices restricted to $\mathcal{S}$. To avoid clutter, we write $\langle\cdot,\cdot\rangle_\mathcal{S} := \langle\cdot,\cdot\rangle_{Q_{X_\mathcal{S}}}$ and $\|\cdot\|_\mathcal{S} := \|\cdot\|_{Q_{X_\mathcal{S}}}$. With $s := |\mathcal{S}|$, denote by $e_i\in\mathbb{R}^s$ the i-th standard basis vector indexed on $\mathcal{S}$. Define

$$\begin{aligned}\Phi_i^{\mathrm{del}}(\mathcal{S}) &:= \|\mathbf{1}\|_\mathcal{S}^2\,\|e_i\|_\mathcal{S}^2 - \langle e_i,\mathbf{1}\rangle_\mathcal{S}^2,\\ \Phi_j^{\mathrm{ins}}(\mathcal{S}) &:= \|\mathbf{1}\|_\mathcal{S}^2\,\sigma_j^2(\mathcal{S}) \;+\; \big(1-\alpha_j(\mathcal{S})\big)^2,\end{aligned}$$

where $\sigma_j^2(\mathcal{S}) := \|x_j\|_2^2 - \|X_\mathcal{S} x_j\|_\mathcal{S}^2$ and $\alpha_j(\mathcal{S}) := \langle\mathbf{1}, X_\mathcal{S} x_j\rangle_\mathcal{S}$.

Theorem 7 (Atomic OLS updates: regular cases)**.** *Let $X\in\mathbb{R}^{n\times p}$ and $y^{(1)}, y^{(0)}\in\mathbb{R}^n$. For any $\alpha\in\{0,1\}$, $\mathcal{S}\in\binom{[n]}{m}$, and $(i,j)\in\mathcal{S}\times\mathcal{S}^c$, the atomic changes in (33) are exactly computable from the ratio representation (9) by one-row Moore–Penrose update formulas.*

1. *In the regular deletion case, if $\Phi_i^{\mathrm{del}}(\mathcal{S})>0$, then*

$$\Delta_\alpha^{\mathrm{del}}(i;\mathcal{S}) = -\frac{\langle e_i,\mathbf{1}\rangle_\mathcal{S}\langle e_i,\tilde y_\mathcal{S}^{(\alpha)}\rangle_\mathcal{S}}{\Phi_i^{\mathrm{del}}(\mathcal{S})},$$

   where $\tilde y_\mathcal{S}^{(\alpha)} := y_\mathcal{S}^{(\alpha)} - \hat\mu_\alpha(\mathcal{S})\mathbf{1}$.

2. *In the regular K-branch insertion case, where both $\mathcal{S}$ and $\mathcal{S}\cup\{j\}$ are in the K-branch and $\sigma_j^2(\mathcal{S})>0$,*

$$\Delta_\alpha^{\mathrm{ins}}(j;\mathcal{S}) = \frac{1-\alpha_j(\mathcal{S})}{\Phi_j^{\mathrm{ins}}(\mathcal{S})}\Big(y_j^{(\alpha)} - \hat\mu_\alpha(\mathcal{S}) - \langle X_\mathcal{S} x_j, \tilde y_\mathcal{S}^{(\alpha)}\rangle_\mathcal{S}\Big).$$

Appendix C.3 gives the full branch-specific identities and proof of Theorem 7. Together with the paired telescope (34), these identities yield oracle one-swap effects $\Delta_{ij}\hat\tau_{\mathrm{OLS}}(\mathcal{S})$. Consequently, $\Gamma(f)$, Lf, and the predictable increments D_t—and hence $(V^\star, R^\star, B^\star)$—are oracle-computable from $(X, y^{(0)}, y^{(1)})$; see Appendix A.2 for more details.

Discussion. Theorem 5 yields a design-based, finite-sample CI for τ centered at $\hat{\tau}$. The pathwise fluctuation terms $(V^\star, R^\star)$ are *instance-adaptive*, through the realized assignment $\mathcal{S}_1$ and reveal order via the predictable increments D_t, while $B^\star$ is a population (oracle) term determined by the Stein residual in (31). Given $(X, y^{(0)}, y^{(1)})$, the exact swap identity (34) and Theorem 7 permit oracle evaluation of $V^\star, R^\star$ and $B^\star$.

For OLS–RA, Theorem 7 shows that swap sensitivity is governed by the geometry of the arm-wise Q-matrices, cf. Appendix A.3. The symmetry between deletion and insertion becomes clear when insertion is viewed as reverse deletion. Deleting i from $\mathcal{S}$ is governed by the Q-leverage of i in the current set $\mathcal{S}$, whereas inserting j into $\mathcal{S}$ is governed by the Q-leverage of j in the augmented set $\mathcal{S} \cup \{j\}$. In regular cases, both atomic changes obey the same leverage–residual template:

$$|\Delta| = \text{standardized residual} \times \frac{\sqrt{\text{leverage}}}{1 - \text{leverage}}.$$

For deletion,

$$\left|\Delta_\alpha^{\mathrm{del}}(i; \mathcal{S})\right| = R_{\alpha,i}^{\mathrm{del}}(\mathcal{S}) \, \frac{\sqrt{\ell_i(\mathcal{S})}}{1 - \ell_i(\mathcal{S})},$$

where

$$\ell_i(\mathcal{S}) := \frac{\langle e_i, \mathbf{1} \rangle_{Q_\mathcal{S}}^2}{\|e_i\|_{Q_\mathcal{S}}^2 \|\mathbf{1}\|_{Q_\mathcal{S}}^2}, \qquad R_{\alpha,i}^{\mathrm{del}}(\mathcal{S}) := \frac{|\langle e_i, \tilde{y}_\mathcal{S}^{(\alpha)} \rangle_\mathcal{S}|}{\|e_i\|_\mathcal{S} \|\mathbf{1}\|_\mathcal{S}}.$$

For insertion, let $\ell_j^+(\mathcal{S}) := \ell_j(\mathcal{S} \cup \{j\})$ denote the leverage of the incoming unit in the augmented set, and let $R_{\alpha,j}^+(\mathcal{S})$ be the residual defined similarly; see (43) in Appendix A.3. Whenever the corresponding regular deletion identity applies in $\mathcal{S} \cup \{j\}$,

$$\left|\Delta_\alpha^{\mathrm{ins}}(j; \mathcal{S})\right| = R_{\alpha,j}^+(\mathcal{S}) \, \frac{\sqrt{\ell_j^+(\mathcal{S})}}{1 - \ell_j^+(\mathcal{S})}.$$

Thus, moderate deletion/insertion leverage and small standardized residuals imply reduced per-swap effects. Compared to DiM, whose swap sensitivity scales with raw potential-outcome contrasts (cf. Example 1), OLS–RA can benefit when covariates reduce residuals without inflating deletion or insertion Q-leverages.

Remark 4.1 (Data-driven proxy). $\Delta_\alpha^{\mathrm{ins}}(j; \mathcal{S})$ depends on the incoming potential outcome $y_j^{(\alpha)}$, which is counterfactual when j lies in the opposite arm. Appendix A.3 isolates the observable geometric factors.

5 EXPERIMENTS

Although this paper is primarily theoretical, Appendix D reports oracle-setting Monte Carlo experiments that quantify the finite-sample behavior of the confidence intervals in Theorem 5. For OLS–RA, the oracle quantities $(V^\star, R^\star, B^\star)$ are approximated by Monte Carlo estimators $(\widehat{V}^\star, \widehat{R}^\star, \widehat{B}^\star)$; for DiM, $V^\star$ and $R^\star$ are available in closed form under the remaining-pool reveal law. Details about the MC procedure, results and discussion are deferred to Appendix D.

Overall, the proposed finite-sample (FS) CIs are valid but conservative. For DiM, the exact remaining-pool oracle quantities yield over-coverage at small n, so Experiment 1 should be read primarily as a validity and conservatism sanity check for the general design-based finite-sample bound. For OLS–RA, the direct MC estimates $(\widehat{V}^\star, \widehat{R}^\star, \widehat{B}^\star)$ produce empirical coverage essentially equal to 1.000 across the reported cells, while the FS widths remain about 2.2–2.6 times the empirical inter-percentile range (IPR), which itself peaks near the arm-wise interpolation threshold and then decreases as p grows, consistent with the Q-geometric sensitivity picture in Section 4.2. Diagnostics suggest nontrivial slack in the current FS bound, motivating future work on tighter assignment-adaptive range and bias proxies that preserve finite-sample validity.

6 CONCLUSION

We developed a design-based, non-asymptotic framework for analyzing regression adjustment under complete randomization, yielding finite-sample-valid confidence intervals for the ATE. Our analysis controls stochastic fluctuation through a variance-adaptive Doob martingale and Freedman's inequality, bounds design bias via Stein's method of exchangeable pairs, and yields an oracle guarantee (Theorem 5) that applies in both classical and high-dimensional/interpolating regimes. For OLS–RA, paired deletion–insertion identities and rank-one pseudoinverse updates reveal how deletion leverage, insertion leverage, and standardized residual energy govern swap sensitivity, clarifying when RA can improve on the DiM baseline and how covariate Q-geometry shapes its precision gains.

Several limitations motivate future work. First, the current guarantees are oracle: computing $(V^\star, R^\star, B^\star)$ requires full potential outcomes. Developing data-driven surrogates based only on (X, T, Y), with finite-sample guarantees, is an important next step. Second, the current Freedman-type concentration and bias bounds can be conservative; sharper geometry-aware proxies, or alternatives to Freedman-type bounds, may tighten the intervals while retaining validity. Third, extending the martingale and exchangeable-pairs constructions beyond complete randomization—for example to stratified, clustered, rerandomized, or covariate-adaptive designs—would broaden applicability.

References

Peter M Aronow, Donald P Green, and Donald KK Lee. Sharp bounds on the variance in randomized experiments. *Annals of Statistics*, 42(3):850–871, 2014.

Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and geometry of Markov diffusion operators*, volume 348. Springer Science & Business Media, 2013.

Adi Ben-Israel and Thomas NE Greville. *Generalized Inverses: Theory and Applications*. Springer, 2003.

Adam Bloniarz, Hanzhong Liu, Cun-Hui Zhang, Jasjeet S Sekhon, and Bin Yu. Lasso adjustments of treatment effect estimates in randomized experiments. *Proceedings of the National Academy of Sciences*, 113(27):7383–7390, 2016.

Matias D Cattaneo, Michael Jansson, and Whitney K Newey. Inference in linear regression models with many covariates and heteroscedasticity. *Journal of the American Statistical Association*, 113(523): 1350–1361, 2018.

Haoge Chang, Joel A Middleton, and PM Aronow. Exact bias correction for linear adjustment of randomized controlled trials. *Econometrica*, 92(5):1503–1519, 2024.

Louis HY Chen, Larry Goldstein, and Qi-Man Shao. *Normal Approximation by Stein's Method*. Springer Science & Business Media, 2010.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018.

Harold D Chiang, Yukitoshi Matsushita, and Taisuke Otsu. Regression adjustment in randomized controlled trials with many covariates. *arXiv preprint arXiv:2302.00469*, 2023.

Tirthankar Dasgupta, Natesh S Pillai, and Donald B Rubin. Causal inference from 2k factorial designs by using potential outcomes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 77(4):727–753, 2015.

Philippe Delsarte. An algebraic approach to the association schemes of coding theory. *Philips Res. Rep. Suppl.*, 10:vi+–97, 1973.

Peng Ding. *A First Course in Causal Inference*. Chapman and Hall/CRC, 2024.

R.A. Fisher. *The Design of Experiments*. The Design of Experiments. Oliver and Boyd, 1935. URL https://books.google.com/books?id=-EsNAQAAIAAJ.

Colin B Fogarty. Regression-assisted inference for the average treatment effect in paired experiments. *Biometrika*, 105(4):994–1000, 2018.

David A Freedman. On tail probabilities for martingales. *the Annals of Probability*, pages 100–118, 1975.

David A Freedman. On regression adjustments to experimental data. *Advances in Applied Mathematics*, 40(2):180–193, 2008.

Thomas NE Greville. Some applications of the pseudoinverse of a matrix. *SIAM Review*, 2(1):15–22, 1960.

Jaroslav Hájek. Limiting distributions in simple random sampling from a finite population. *A Magyar Tudományos Akadémia Matematikai Kutató Intézetének közlemenyei*, 5(3):361–374, 1960.

Klaus Hinkelmann and Oscar Kempthorne. *Design and Analysis of Experiments, Volume 1: Introduction to Experimental Design*, volume 1. John Wiley & Sons, 2007.

Guido W Imbens and Donald B Rubin. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press, 2015.

Oscar Kempthorne. *The Design and Analysis of Experiments*. A Wiley Publication in Mathematical Statistics. John Wiley & Sons, New York, 1952.

Lihua Lei and Peng Ding. Regression adjustment in completely randomized experiments with a diverging number of covariates. *Biometrika*, 108(4):815–828, 2021.

David A Levin and Yuval Peres. *Markov Chains and Mixing Times*, volume 107. American Mathematical Society, 2017.

Xinran Li and Peng Ding. Rerandomization and regression adjustment. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(1): 241–268, 2020.

Winston Lin. Agnostic notes on regression adjustments to experimental data: Reexamining Freedman's critique. *The Annals of Applied Statistics*, 7(1):295 – 318, 2013. doi: 10.1214/12-AOAS583. URL https://doi.org/10.1214/12-AOAS583.

Rahul Mukerjee, Tirthankar Dasgupta, and Donald B Rubin. Using standard tools from finite population sampling to improve causal inference for complex experiments. *Journal of the American Statistical Association*, 113(522):868–881, 2018.

Jerzy Neyman. Sur les applications de la théorie des probabilités aux expériences agricoles: Essai des principes. *Roczniki Nauk Rolniczych*, 10:1–51, 1923.

Excerpts translated and reprinted in English in *Statistical Science*, Vol. 5, No. 4 (1990), pp. 465–472, by D. M. Dabrowska and T. P. Speed.

Gesine Reinert and Adrian Röllin. Multivariate normal approximation with stein's method of exchangeable pairs under a general linearity condition. *Annals of Probability*, 37(6):2150–2173, 2009.

William F Rosenberger and John M Lachin. *Randomization in Clinical Trials: Theory and Practice*. John Wiley & Sons, 2015.

Nathan Ross. Fundamentals of Stein's method. *Probability Surveys*, 88:210–293, 2011.

Dennis Shen, Dogyoon Song, Peng Ding, and Jasjeet S Sekhon. Algebraic and statistical properties of the ordinary least squares interpolator. *arXiv preprint arXiv:2309.15769*, 2023.

Dogyoon Song. Neumann-series corrections for regression adjustment in randomized experiments. *arXiv preprint arXiv:2511.08539*, 2025.

Anastasios A Tsiatis, Marie Davidian, Min Zhang, and Xiaomin Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: a principled yet flexible approach. *Statistics in Medicine*, 27(23):4658–4677, 2008.

Stefan Wager, Wenfei Du, Jonathan Taylor, and Robert J Tibshirani. High-dimensional regression adjustments in randomized experiments. *Proceedings of the National Academy of Sciences*, 113(45): 12673–12678, 2016.

Checklist

1. For all models and algorithms presented, check if you include:
   (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [**Yes**]
   (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [**Yes**]
   (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [**Not Applicable**]

2. For any theoretical claim, check if you include:
   (a) Statements of the full set of assumptions of all theoretical results. [**Yes**]
   (b) Complete proofs of all theoretical results. [**Yes**]
   (c) Clear explanations of any assumptions. [**Yes**]

3. For all figures and tables that present empirical results, check if you include:
   (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [**Yes**]
   (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [**Not Applicable**]
   (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [**Yes**]
   (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [**Yes**]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
   (a) Citations of the creator If your work uses existing assets. [**Not Applicable**]
   (b) The license information of the assets, if applicable. [**Not Applicable**]
   (c) New assets either in the supplemental material or as a URL, if applicable. [**Not Applicable**]
   (d) Information about consent from data providers/curators. [**Not Applicable**]
   (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [**Not Applicable**]

5. If you used crowdsourcing or conducted research with human subjects, check if you include:
   (a) The full text of instructions given to participants and screenshots. [**Not Applicable**]
   (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [**Not Applicable**]
   (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [**Not Applicable**]

Supplementary Material for "Design-Based Finite-Sample Analysis for Regression Adjustment"

The supplementary material is organized as follows. Section A offers additional discussion, including an oracle Stein bias example for DiM (Section A.1), further details of the swap–sensitivity analysis for OLS–RA (Section A.2), and geometry-based envelopes computable without counterfactuals (Sections A.3–A.4). Section B contains deferred proofs for the results in Section 3. Section C presents the proofs of the main results from Section 4. Finally, Section D reports numerical experiments that complement the theory in the main text.

A ADDITIONAL DISCUSSION

A.1 Example: Oracle Bias Bound for DiM

Example 2. *As a sanity check, we compute the oracle bias parameter $B^\star$ in (31) for $\hat{\tau}_{\mathrm{DiM}}$.*

Define

$$a_i := \frac{y_i^{(1)}}{n_1} + \frac{y_i^{(0)}}{n_0}, \qquad \bar{a} := \frac{1}{n}\sum_{i=1}^{n} a_i.$$

Since

$$\hat{\tau}_{\mathrm{DiM}}(\mathcal{S}) - \tau = \sum_{i\in\mathcal{S}} (a_i - \bar{a}),$$

the DiM error $f(\mathcal{S}) = \hat{\tau}_{\mathrm{DiM}}(\mathcal{S}) - \tau$ is a linear function on the Johnson graph. Hence, averaging one-swap changes gives

$$\begin{aligned}(Lf)(\mathcal{S}) &= \frac{1}{n_1 n_0}\sum_{i\in\mathcal{S}}\sum_{j\notin\mathcal{S}}\big\{f(\mathcal{S}^{(i\leftrightarrow j)}) - f(\mathcal{S})\big\} \\ &= \frac{1}{n_1 n_0}\sum_{i\in\mathcal{S}}\sum_{j\notin\mathcal{S}}(a_j - a_i) \\ &= -\frac{n}{n_1 n_0} f(\mathcal{S}).\end{aligned}$$

Thus f is an exact eigenfunction of the Johnson-walk generator with

$$\lambda_{\mathrm{lin}} = \frac{n}{n_1 n_0}.$$

Moreover,

$$\mathrm{Var}(f) = \frac{n_1 n_0}{n-1}\left(\frac{\mathrm{Var}_1}{n_1^2} + \frac{\mathrm{Var}_0}{n_0^2} + \frac{2\,\mathrm{Cov}_{\mathrm{diff}}}{n_1 n_0}\right), \tag{35}$$

where

$$\mathrm{Var}_\alpha := \frac{1}{n}\sum_{i=1}^{n}\big(y_i^{(\alpha)} - \bar{y}^{(\alpha)}\big)^2 \quad (\alpha\in\{0,1\}), \qquad \mathrm{Cov}_{\mathrm{diff}} := \frac{1}{n}\sum_{i=1}^{n}\big(y_i^{(1)} - \bar{y}^{(1)}\big)\big(y_i^{(0)} - \bar{y}^{(0)}\big).$$

Next, recall that the carré du champ for Johnson walk (simple random walk generated by uniform transition kernel on the Johnson graph), cf. (30), is

$$\Gamma(f)(\mathcal{S}) = \frac{1}{2\,n_1 n_0}\sum_{i\in\mathcal{S}}\sum_{j\notin\mathcal{S}}\big(\Delta_{ij} f(\mathcal{S})\big)^2.$$

For $f = \hat{\tau}_{\mathrm{DiM}} - \tau$, observe (cf. Example 1) that

$$\Delta_{ij}\hat{\tau}_{\mathrm{DiM}} = \frac{1}{n_1}\big(y_j^{(1)} - y_i^{(1)}\big) + \frac{1}{n_0}\big(y_j^{(0)} - y_i^{(0)}\big). \tag{36}$$

Averaging $\Gamma(f)(\mathcal{S})$ over $\mathcal{S} \sim \mathrm{Unif}\binom{[n]}{n_1}$ (complete randomization) equals averaging over all ordered pairs (i,j) with $i \neq j$, i.e., for any function $g : [n] \times [n] \to \mathbb{R}$,

$$\mathbb{E}_S\left[\frac{1}{n_1 n_0}\sum_{i\in S}\sum_{j\notin S} g(i,j)\right] = \frac{1}{n(n-1)}\sum_{i\neq j} g(i,j). \tag{37}$$

Lastly, observe the finite-population identities (straightforward expansions):

$$\frac{1}{n(n-1)}\sum_{i\neq j}\big(a_j - a_i\big)^2 = \frac{2n}{n-1}\,\mathrm{Var}_n(a),$$

$$\frac{1}{n(n-1)}\sum_{i\neq j}\big(a_j - a_i\big)\big(b_j - b_i\big) = \frac{2n}{n-1}\,\mathrm{Cov}_n(a,b),$$

where $\mathrm{Var}_n(a) = \frac{1}{n}\sum_{i=1}^n (a_i - \bar{a})^2$ and $\mathrm{Cov}_n(a,b) = \frac{1}{n}\sum_{i=1}^n (a_i - \bar{a})(b_i - \bar{b})$.

Therefore,

$$\begin{aligned}
\mathbb{E}_{\mathcal{S}}[\Gamma(\hat{\tau}_{\mathrm{DiM}} - \tau)(\mathcal{S})] &= \mathbb{E}_{\mathcal{S}}\left[\frac{1}{2\,n_1 n_0}\sum_{i\in\mathcal{S}}\sum_{j\notin\mathcal{S}}\big(\Delta_{ij}\hat{\tau}_{\mathrm{DiM}}(\mathcal{S})\big)^2\right] && \because \Delta_{ij}(\hat{\tau}_{\mathrm{DiM}} - \tau) = \Delta_{ij}\hat{\tau}_{\mathrm{DiM}} \ \&\ (36)\\
&= \frac{1}{2n(n-1)}\sum_{i,j\in[n]:i\neq j}\big(\Delta_{ij}\hat{\tau}_{\mathrm{DiM}}(\mathcal{S})\big)^2 && \because (37)\\
&= \frac{1}{2}\cdot\frac{1}{n(n-1)}\sum_{i\neq j}\left[\frac{\big(y_j^{(1)} - y_i^{(1)}\big)^2}{n_1^2} + \frac{\big(y_j^{(0)} - y_i^{(0)}\big)^2}{n_0^2} + \frac{2\big(y_j^{(1)} - y_i^{(1)}\big)\big(y_j^{(0)} - y_i^{(0)}\big)}{n_1 n_0}\right]\\
&= \frac{n}{n-1}\left(\frac{\mathrm{Var}_1}{n_1^2} + \frac{\mathrm{Var}_0}{n_0^2} + \frac{2\,\mathrm{Cov}_{\mathrm{diff}}}{n_1 n_0}\right).
\end{aligned} \tag{38}$$

Combining (35) with (38) yields

$$\lambda^\star = \frac{\mathbb{E}\Gamma(f)}{\mathrm{Var}(f)} = \frac{n}{n_1 n_0} = \lambda_{\mathrm{lin}}.$$

Therefore, under the corrected definition (31),

$$B^\star = \frac{\sqrt{\mathbb{E}\big(Lf + \lambda^\star f\big)^2}}{\lambda^\star} = 0.$$

Thus, for DiM estimator, the Stein bias parameter $B^\star$ captures its design unbiasedness.

A.2 Swap Sensitivity Analysis of OLS-RA

Here we assemble the exact swap sensitivity for OLS–RA from the atomic deletion/insertion identities (Theorem 7) using the paired deletion–insertion decomposition (Proposition 6). Throughout, $\mathcal{S}_1$ is the realized treatment set, $\mathcal{S}_0 = [n] \setminus \mathcal{S}_1$, and we consider a paired swap $(i \leftrightarrow j)$ with $i \in \mathcal{S}_1$, $j \in \mathcal{S}_0$.

Recollecting Notation. For a set $\mathcal{S}$, write $Q_{\mathcal{S}} \equiv Q_{X_{\mathcal{S}}}$ and define the $Q_{\mathcal{S}}$–inner product and norm by $\langle u, v\rangle_{\mathcal{S}} := u^\top Q_{\mathcal{S}} v$ and $\|u\|_{\mathcal{S}} := \sqrt{u^\top Q_{\mathcal{S}} u}$. Let $\mathbf{1}$ be the all–ones vector of length $|\mathcal{S}|$. For arm $\alpha \in \{0,1\}$,

$$\hat{\mu}_\alpha(\mathcal{S}) = \frac{\langle \mathbf{1}, y_{\mathcal{S}}^{(\alpha)}\rangle_{\mathcal{S}}}{\|\mathbf{1}\|_{\mathcal{S}}^2}, \qquad \text{and} \qquad \tilde{y}_{\mathcal{S}}^{(\alpha)} := y_{\mathcal{S}}^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S})\mathbf{1}.$$

For $j \notin \mathcal{S}$, define $\sigma_j^2(\mathcal{S}) := \|x_j\|_2^2 - \|X_{\mathcal{S}} x_j\|_{\mathcal{S}}^2$ and $\alpha_j(\mathcal{S}) := \langle \mathbf{1}, X_{\mathcal{S}} x_j\rangle_{\mathcal{S}}$. We define

$$\Phi_i^{\mathrm{del}}(\mathcal{S}) := \|\mathbf{1}\|_{\mathcal{S}}^2\,\|e_i\|_{\mathcal{S}}^2 - \langle e_i, \mathbf{1}\rangle_{\mathcal{S}}^2, \qquad \text{and} \qquad \Phi_j^{\mathrm{ins}}(\mathcal{S}) := \|\mathbf{1}\|_{\mathcal{S}}^2\,\sigma_j^2(\mathcal{S}) \ + \ \big(1 - \alpha_j(\mathcal{S})\big)^2,$$

Paired Deletion–Insertion Decomposition. For $i \in \mathcal{S}_1$, $j \in \mathcal{S}_0$, let $\mathcal{S}_1^- = \mathcal{S}_1 \setminus \{i\}$ and $\mathcal{S}_0^- = \mathcal{S}_0 \setminus \{j\}$. Then the paired swap $(i \leftrightarrow j)$ can be decomposed into *two* row operations in each arm:

$$\mathcal{S}_1 \xrightarrow{\text{delete } i} \mathcal{S}_1^- := \mathcal{S}_1 \setminus \{i\} \xrightarrow{\text{insert } j} \mathcal{S}_1' := \mathcal{S}_1^- \cup \{j\},$$
$$\mathcal{S}_0 \xrightarrow{\text{delete } j} \mathcal{S}_0^- := \mathcal{S}_0 \setminus \{j\} \xrightarrow{\text{insert } i} \mathcal{S}_0' := \mathcal{S}_0^- \cup \{i\},$$

and

$$\Delta_{ij}\hat{\tau}_{\mathrm{OLS}}(\mathcal{S}_1) = \underbrace{\Delta_1^{\mathrm{del}}(i;\mathcal{S}_1) + \Delta_1^{\mathrm{ins}}(j;\mathcal{S}_1^-)}_{\text{treatment arm}} - \underbrace{\Big(\Delta_0^{\mathrm{del}}(j;\mathcal{S}_0) + \Delta_0^{\mathrm{ins}}(i;\mathcal{S}_0^-)\Big)}_{\text{control arm}}. \tag{39}$$

This is a purely algebraic telescope (see Proposition 6 and its proof in Appendix C.2).

Atomic Identities and Unit-wise Pairing. By Theorem 7, for any $\alpha \in \{0,1\}$, $\mathcal{S} \in \binom{[n]}{m}$, $i \in \mathcal{S}$, and $j \notin \mathcal{S}$, the regular deletion formula is

$$\Delta_\alpha^{\mathrm{del}}(i;\mathcal{S}) = -\frac{\langle e_i, \mathbf{1}\rangle_{\mathcal{S}} \langle e_i, \tilde{y}_{\mathcal{S}}^{(\alpha)}\rangle_{\mathcal{S}}}{\Phi_i^{\mathrm{del}}(\mathcal{S})}, \tag{40}$$

whenever the regular deletion update applies and $\Phi_i^{\mathrm{del}}(\mathcal{S}) > 0$. For insertion, Theorem 7 gives the exact block-ratio formula in general, with the simplified K- and M-branch regular cases stated in (62) and (63).

Define the unit-wise contributions

$$c_i := \Delta_1^{\mathrm{del}}(i;\mathcal{S}_1) \;-\; \Delta_0^{\mathrm{ins}}(i;\mathcal{S}_0^-), \qquad d_j := \Delta_1^{\mathrm{ins}}(j;\mathcal{S}_1^-) \;-\; \Delta_0^{\mathrm{del}}(j;\mathcal{S}_0).$$

Substituting the deletion identity (40) together with the appropriate branch-specific insertion identity from Lemma C.5 into (39) yields the symmetric exact form

$$\Delta_{ij}\hat{\tau}_{\mathrm{OLS}}(\mathcal{S}_1) = c_i + d_j. \tag{41}$$

A Remark on Oracle vs. Data-driven Analyses. Because insertion terms involve incoming counterfactual outcomes $y_j^{(1)}$ and $y_i^{(0)}$, the exact c_i, d_j (and hence $\Delta_{ij}\hat{\tau}_{\mathrm{OLS}}$) are oracle quantities. Appendix A.3 identifies the Q-geometric factors—deletion leverage, insertion leverage, and standardized residual scales—that govern these terms. Fully data-driven insertion envelopes additionally require observable bounds for the incoming residuals.

A.3 Geometric Envelopes for Deletion and Insertion

The exact swap sensitivity in (41) combines atomic deletion/insertion changes in closed form, cf. (40), (61), (62), and (63). This subsection records geometry-driven envelopes for these terms and clarifies which parts are oracle and which would require additional data-driven proxies.

Deletion and Insertion Leverages. For $\mathcal{A} \subset [n]$ and $u \in \mathcal{A}$, define the Q-leverage

$$\ell_u(\mathcal{A}) := \frac{\langle e_u, \mathbf{1}\rangle_{Q_{X_{\mathcal{A}}}}^2}{\|e_u\|_{Q_{X_{\mathcal{A}}}}^2 \|\mathbf{1}\|_{Q_{X_{\mathcal{A}}}}^2} \in [0,1].$$

Deletion uses the current leverage $\ell_i(\mathcal{S})$. Insertion uses the augmented leverage

$$\ell_j^+(\mathcal{S}) := \ell_j(\mathcal{S} \cup \{j\}),$$

i.e., the leverage of the incoming unit after it has been inserted.

In the regular K-branch, this augmented leverage has the pre-insertion expression

$$\ell_j^{K,+}(\mathcal{S}) = \frac{(1-\alpha_j(\mathcal{S}))^2}{\|\mathbf{1}\|_{Q_{\mathcal{S}}}^2 \sigma_j^2(\mathcal{S}) + (1-\alpha_j(\mathcal{S}))^2},$$

where

$$\alpha_j(\mathcal{S}) = \langle \mathbf{1}, X_{\mathcal{S}} x_j\rangle_{Q_{\mathcal{S}}}, \qquad \sigma_j^2(\mathcal{S}) = \|x_j\|_2^2 - \|X_{\mathcal{S}} x_j\|_{Q_{\mathcal{S}}}^2.$$

In the M-branch row-span case $c_j = 0$, the augmented insertion leverage is

$$\ell_j^{M,+}(\mathcal{S}) = \frac{(1-\eta_j)^2}{\|\mathbf{1}\|_{Q_\mathcal{S}}^2(1+h_j) + (1-\eta_j)^2},$$

with h_j, η_j as in (63). If $c_j \neq 0$, then the M-branch insertion effect is zero.

Remark A.1 (Relation to classical hat leverages). The quantity $\ell_u(\mathcal{A})$, which we refer to as the "Q-leverage" (or simply "leverage"), is a directional cosine in Q-geometry and is not equal to the diagonal of the OLS hat matrix with intercept, which is commonly referred to as the "leverage score" in the regression analysis. In the special case $Q_\mathcal{A} = M_{X_\mathcal{A}}$ and $1^\top X_\mathcal{A} = 0$,

$$\ell_u(\mathcal{A}) \;=\; \frac{1}{|\mathcal{A}|\,(1-h_{uu}(X_\mathcal{A}))}, \qquad \text{where} \qquad h_{uu}(X_\mathcal{A}) = \big[X_\mathcal{A}(X_\mathcal{A}^\top X_\mathcal{A})^{-1}X_\mathcal{A}^\top\big]_{uu},$$

linking our Q-leverage to the familiar hat diagonals h_{ii}. We use the term "leverage" to denote this Q-cosine and retain the classical term "hat leverage" when referring to $h_{uu}(X_\mathcal{A})$.

Geometric Envelopes. From the regular deletion identity (40),

$$\Delta_\alpha^{\mathrm{del}}(i;\mathcal{S}) = -\frac{\langle e_i, \mathbf{1}\rangle_\mathcal{S}\langle e_i, \tilde{y}_\mathcal{S}^{(\alpha)}\rangle_\mathcal{S}}{\|\mathbf{1}\|_\mathcal{S}^2\,\|e_i\|_\mathcal{S}^2 - \langle e_i, \mathbf{1}\rangle_\mathcal{S}^2}.$$

Taking absolute values gives the envelope:

$$\big|\Delta_\alpha^{\mathrm{del}}(i;\mathcal{S})\big| = R_{\alpha,i}^{\mathrm{del}}(\mathcal{S})\frac{\sqrt{\ell_i(\mathcal{S})}}{1-\ell_i(\mathcal{S})}, \qquad R_{\alpha,i}^{\mathrm{del}}(\mathcal{S}) := \frac{|\langle e_i, \tilde{y}_\mathcal{S}^{(\alpha)}\rangle_\mathcal{S}|}{\|e_i\|_\mathcal{S}\|\mathbf{1}\|_\mathcal{S}}.$$

Insertion has the same structure when viewed as reverse deletion. Let $\mathcal{S}^+ = \mathcal{S}\cup\{j\}$. Since

$$\Delta_\alpha^{\mathrm{ins}}(j;\mathcal{S}) = -\Delta_\alpha^{\mathrm{del}}(j;\mathcal{S}^+),$$

the augmented-set leverage gives, in regular cases,

$$\big|\Delta_\alpha^{\mathrm{ins}}(j;\mathcal{S})\big| = R_{\alpha,j}^{+}(\mathcal{S})\frac{\sqrt{\ell_j^+(\mathcal{S})}}{1-\ell_j^+(\mathcal{S})}, \tag{42}$$

where

$$R_{\alpha,j}^{+}(\mathcal{S}) := \frac{|\langle e_j, \widetilde{y}_{\mathcal{S}^+}^{(\alpha)}\rangle_{Q_{X_{\mathcal{S}^+}}}|}{\|e_j\|_{Q_{X_{\mathcal{S}^+}}}\|\mathbf{1}\|_{Q_{X_{\mathcal{S}^+}}}}. \tag{43}$$

Thus deletion and insertion obey the same leverage–residual template; only the set defining the geometry changes from $\mathcal{S}$ to $\mathcal{S}^+$.

The augmented-set representation (42) provides a symmetrical viewpoint that treats insertion as deletion in $\mathcal{S}^+$. For computation and interpretation, it is also useful to rewrite this same quantity in pre-insertion coordinates. In the regular K-branch insertion case, (62) to be presented in Lemma C.5 can be equivalently written as

$$\big|\Delta_\alpha^{\mathrm{ins}}(j;\mathcal{S})\big| = R_{\alpha,j}^{K,+}(\mathcal{S})\sqrt{\ell_j^{K,+}(\mathcal{S})\big(1-\ell_j^{K,+}(\mathcal{S})\big)}, \tag{44}$$

where

$$R_{\alpha,j}^{K,+}(\mathcal{S}) := \frac{\Big|y_j^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S}) - \langle X_\mathcal{S}x_j, \tilde{y}_\mathcal{S}^{(\alpha)}\rangle_{Q_\mathcal{S}}\Big|}{\sigma_j(\mathcal{S})\|\mathbf{1}\|_{Q_\mathcal{S}}}.$$

Indeed, this follows from

$$\sqrt{\ell_j^{K,+}(\mathcal{S})\big(1-\ell_j^{K,+}(\mathcal{S})\big)} = \frac{|1-\alpha_j(\mathcal{S})|\,\sigma_j(\mathcal{S})\|\mathbf{1}\|_{Q_\mathcal{S}}}{\|\mathbf{1}\|_{Q_\mathcal{S}}^2\sigma_j^2(\mathcal{S}) + (1-\alpha_j(\mathcal{S}))^2}.$$

Similarly, in the M-branch row-span case $c_j = 0$, (63) to be presented in Lemma C.5 gives

$$\left|\Delta_\alpha^{\text{ins}}(j;\mathcal{S})\right| = R_{\alpha,j}^{M,+}(\mathcal{S})\sqrt{\ell_j^{M,+}(\mathcal{S})\big(1-\ell_j^{M,+}(\mathcal{S})\big)}, \tag{45}$$

where

$$R_{\alpha,j}^{M,+}(\mathcal{S}) := \frac{|y_j^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S}) - \beta_j|}{\sqrt{1+h_j}\,\|\mathbf{1}\|_{Q_\mathcal{S}}}.$$

If $c_j \neq 0$, then $\Delta_\alpha^{\text{ins}}(j;\mathcal{S}) = 0$. Degenerate or branch-switching cases are handled by the exact ratio and block identities in Appendix C.3.

Remarks on the Envelopes. The deletion envelope is amplified when i is highly aligned with the intercept direction in the current $Q_\mathcal{S}$-geometry, i.e., when $\ell_i(\mathcal{S})$ is close to one. Insertion is amplified when the incoming unit has high leverage in the augmented geometry $Q_{X_{\mathcal{S}\cup\{j\}}}$, i.e., when $\ell_j^+(\mathcal{S})$ is close to one. Thus, OLS–RA swap effects are small when deletion leverage, insertion leverage, and the corresponding standardized residual scales are controlled.

The insertion residuals in $R_{\alpha,j}^+$, $R_{\alpha,j}^{K,+}$, and $R_{\alpha,j}^{M,+}$ involve the incoming potential outcome $y_j^{(\alpha)}$, which is counterfactual when j lies in the opposite arm. A fully data-driven insertion proxy therefore requires an additional observable envelope for these incoming residuals. Developing such calibrated envelopes is left to future work.

A.4 Envelope-controlled Concentration and Bias

The exact formulas above give a geometry-driven swap envelope for OLS–RA. For $\alpha \in \{0,1\}$, define the local deletion envelope

$$\mathcal{D}_{\alpha,i}(\mathcal{A}) := \begin{cases} R_{\alpha,i}^{\text{del}}(\mathcal{A})\dfrac{\sqrt{\ell_i(\mathcal{A})}}{1-\ell_i(\mathcal{A})}, & \text{if the regular deletion identity applies,} \\ \left|\Delta_\alpha^{\text{del}}(i;\mathcal{A})\right|, & \text{otherwise, using the exact ratio identity.} \end{cases}$$

Similarly, define the local insertion envelope

$$\mathcal{I}_{\alpha,j}(\mathcal{A}) := \begin{cases} R_{\alpha,j}^{+}(\mathcal{A})\dfrac{\sqrt{\ell_j^+(\mathcal{A})}}{1-\ell_j^+(\mathcal{A})}, & \text{if the regular augmented deletion identity applies,} \\ \left|\Delta_\alpha^{\text{ins}}(j;\mathcal{A})\right|, & \text{otherwise, using the exact block identity.} \end{cases}$$

Then we set

$$\mathfrak{D}_\alpha := \sup_{\substack{\mathcal{A}\subset[n],\,|\mathcal{A}|=n_\alpha \\ i\in\mathcal{A}}} \mathcal{D}_{\alpha,i}(\mathcal{A}), \qquad \text{and} \qquad \mathfrak{I}_\alpha := \sup_{\substack{\mathcal{A}\subset[n],\,|\mathcal{A}|=n_\alpha-1 \\ j\notin\mathcal{A}}} \mathcal{I}_{\alpha,j}(\mathcal{A}).$$

In regular K- and M-branch cases, $\mathcal{I}_{\alpha,j}$ can equivalently be computed from the pre-insertion forms (44) and (45). Finally, set

$$\Delta_{\text{geo}} := \mathfrak{D}_1 + \mathfrak{I}_1 + \mathfrak{D}_0 + \mathfrak{I}_0.$$

Lemma A.1 (Geometry-controlled oracle parameters)**.** *For OLS–RA with $f(\mathcal{S}) = \hat{\tau}_{\mathsf{OLS}}(\mathcal{S}) - \tau$, the oracle parameters satisfy*

$$R^\star \le \Delta_{\text{geo}}, \qquad V^\star \le A_n\,\Delta_{\text{geo}}^2, \qquad \mathbb{E}\Gamma(f) \le \tfrac{1}{2}\Delta_{\text{geo}}^2, \tag{46}$$

where $A_n := \sum_{t=1}^{n_1}\alpha_t^2$. Moreover,

$$(B^\star)^2 \le |\mathbb{E}f|^2 + \frac{\Delta_{\text{geo}}^2}{(\lambda^\star)^2} \le |\mathbb{E}f|^2 + \Big(\frac{n_1 n_0}{n}\Delta_{\text{geo}}\Big)^2. \tag{47}$$

Proof of Lemma A.1. By construction of $\mathfrak{D}_\alpha$ and $\mathfrak{I}_\alpha$, together with the paired telescope (39),

$$|\Delta_{ij}\hat{\tau}_{\mathsf{OLS}}(\mathcal{S}_1)| \le \mathfrak{D}_1 + \mathfrak{I}_1 + \mathfrak{D}_0 + \mathfrak{I}_0 = \Delta_{\text{geo}}$$

for every admissible swap $(i,j) \in \mathcal{S}_1 \times \mathcal{S}_0$. The rest follows by propagating this geometric swap envelope through the definitions of $R^\star, V^\star, \Gamma(f)$, and $B^\star$.

By Proposition 2, $D_t = -\alpha_t g_t(I)$, where

$$g_t(i) = \mathbb{E}\left[\Delta_{iJ} f(\mathcal{S}_{t-1}^{\text{prox}}(i,\mathcal{T})) \mid \mathcal{F}_{t-1}\right], \qquad I \mid \mathcal{F}_{t-1} \sim \text{Unif}(\mathcal{R}_{t-1}).$$

The envelope gives $|g_t(i)| \le \Delta_{\text{geo}}$. Hence

$$r_t^\star \le \alpha_t \Delta_{\text{geo}} \le \Delta_{\text{geo}}, \qquad v_t^\star = \alpha_t^2 \text{Var}(g_t(I) \mid \mathcal{F}_{t-1}) \le \alpha_t^2 \Delta_{\text{geo}}^2.$$

Summing over t gives

$$R^\star \le \Delta_{\text{geo}}, \qquad V^\star \le A_n \Delta_{\text{geo}}^2.$$

Moreover, by (30),

$$\Gamma(f)(\mathcal{S}) = \frac{1}{2n_1 n_0} \sum_{i \in \mathcal{S}} \sum_{j \notin \mathcal{S}} \left(\Delta_{ij} f(\mathcal{S})\right)^2 \le \frac{1}{2} \Delta_{\text{geo}}^2,$$

and hence $\mathbb{E}\Gamma(f) \le \frac{1}{2}\Delta_{\text{geo}}^2$.

Finally, since $Lf(\mathcal{S})$ is the average one-swap change of f, $|Lf(\mathcal{S})| \le \Delta_{\text{geo}}$. Let $\psi_c^\star := Lf + \lambda^\star(f - \mathbb{E}f)$ as in (25). By the definition of $\lambda^\star$ as the minimizer of $\mathbb{E}\{Lf + \lambda(f - \mathbb{E}f)\}^2$,

$$\mathbb{E}(\psi_c^\star)^2 \le \mathbb{E}(Lf)^2 \le \Delta_{\text{geo}}^2.$$

Since $Lf + \lambda^\star f = \psi_c^\star + \lambda^\star \mathbb{E}f$ and $\mathbb{E}\psi_c^\star = 0$, we obtain

$$(B^\star)^2 = \frac{\mathbb{E}(Lf + \lambda^\star f)^2}{(\lambda^\star)^2} = |\mathbb{E}f|^2 + \frac{\mathbb{E}(\psi_c^\star)^2}{(\lambda^\star)^2} \le |\mathbb{E}f|^2 + \frac{\Delta_{\text{geo}}^2}{(\lambda^\star)^2}.$$

Using $\lambda^\star \ge \text{gap}_{n,n_1} = n/(n_1 n_0)$ gives (47). □

Corollary A.1 (Typical scaling)**.** *If the geometry-driven envelope above satisfies $\Delta_{\text{geo}} = o(1)$, then*

$$R^\star = O(\Delta_{\text{geo}}), \qquad V^\star = O(A_n \Delta_{\text{geo}}^2), \qquad \mathbb{E}\Gamma(f) = O(\Delta_{\text{geo}}^2).$$

In particular, if $A_n = O(n)$ and $\Delta_{\text{geo}} = O(n^{-1/2})$, then $V^\star = O(1)$. For the corrected Stein bias parameter,

$$B^\star \le |\mathbb{E}f| + \frac{\Delta_{\text{geo}}}{\lambda^\star} \le |\mathbb{E}f| + \frac{n_1 n_0}{n} \Delta_{\text{geo}}.$$

The final inequality is a worst-case Poincaré fallback and can be loose in balanced designs; sharper control depends on the actual Rayleigh quotient $\lambda^\star$ and the centered Stein residual.

Tightness of the Bounds in Lemma A.1. The constants in (46) are sharp up to the Cauchy–Schwarz slack in the standardized residual scales. The bounds are informative when deletion leverage and insertion leverage are moderate and the standardized deletion/insertion residuals are small. Large leverage, large incoming residuals, or nearly singular innovation scales can inflate the local range and variance, possibly reflecting genuine sensitivity rather than merely looseness of the current analysis.

B DEFERRED PROOFS FROM SECTION 3

B.1 Proof of Proposition 1

Proof of Proposition 1. Let $L(\mu,\beta) = \|y - \mu\mathbf{1}_n - X\beta\|_2^2$. The normal equations for least-squares are

$$\nabla_\beta L(\mu,\beta) = -2 \cdot X^\top\big(y - \mu\mathbf{1}_n - X\beta\big) = 0, \tag{48}$$

$$\nabla_\mu L(\mu,\beta) = -2 \cdot \mathbf{1}_n^\top\big(y - \mu\mathbf{1}_n - X\beta\big) = 0. \tag{49}$$

We prove this proposition in three steps.

- **Step 1: Characterization of $\mathcal{S}_{X,y}$ (the set of LS minimizers).** Observe that the normal equations (48)–(49) form first-order (necessary) conditions for optimality for the least squares; that is, if $(\mu, \beta) \in \mathcal{S}_{X,y}$, then these equations must be satisfied.

  For each $\mu \in \mathbb{R}$, the set of β satisfying (48) is the affine set

$$\mathcal{S}(\mu) \coloneqq \left\{ \beta(\mu) + (I_p - X^\dagger X)w \ : \ w \in \mathbb{R}^p \right\}, \qquad \text{where} \qquad \beta(\mu) \coloneqq X^\dagger(y - \mu \mathbf{1}_n). \tag{50}$$

  For any $\beta \in \mathcal{S}(\mu)$, the regression residual is

$$\begin{aligned} r(\mu) &= y - \mu \mathbf{1}_n - X\beta \\ &= y - \mu \mathbf{1}_n - XX^\dagger(y - \mu \mathbf{1}_n) && \because X(I_p - X^\dagger X) = 0 \\ &= M_X\,(y - \mu \mathbf{1}_n). \end{aligned}$$

  Therefore, (49) is equivalent to the *scalar* constraint

$$\mathbf{1}_n^\top M_X\,(y - \mu \mathbf{1}_n) = 0. \tag{51}$$

  - If $\mathbf{1}_n \notin \mathrm{col}(X)$, then $M_X \mathbf{1}_n \neq 0$ and (51) determines a *unique* intercept

$$\mu^\star \ = \ \frac{\mathbf{1}_n^\top M_X y}{\mathbf{1}_n^\top M_X \mathbf{1}_n}. \tag{52}$$

  - If $\mathbf{1}_n \in \mathrm{col}(X)$, then $M_X \mathbf{1}_n = 0$, and (51) holds automatically; *every* $\mu \in \mathbb{R}$ is feasible in this case.

  In summary,

$$\mathcal{S}_{X,y} = \Big\{ (\mu,\, \beta(\mu) + (I_p - X^\dagger X)w) \ : \ \mu \in \mathcal{M},\ w \in \mathbb{R}^p \Big\}, \qquad \text{where} \qquad \mathcal{M} = \begin{cases} \{\mu^\star\}, & \mathbf{1}_n \notin \mathrm{col}(X), \\ \mathbb{R}, & \mathbf{1}_n \in \mathrm{col}(X). \end{cases}$$

- **Step 2: Uniqueness of the minimum-norm OLS solution.** We minimize $\|\beta\|_2^2$ over $(\mu, \beta) \in \mathcal{S}_{X,y}$. Let $\mathsf{N}(X) = \{v \in \mathbb{R}^p : Xv = 0\}$ denote the null space of X. Observe that

$$\beta(\mu) \in \mathrm{range}(X^\top), \qquad (I_p - X^\dagger X)w \in \mathsf{N}(X), \qquad \mathrm{range}(X^\top) \perp \mathsf{N}(X).$$

  For any feasible μ and w,

$$\begin{aligned} \|\beta(\mu) + (I_p - X^\dagger X)w\|_2^2 &= \|\beta(\mu)\|_2^2 + \|(I_p - X^\dagger X)w\|_2^2 \\ &\geq \|\beta(\mu)\|_2^2, \end{aligned}$$

  with equality iff $w = 0$, cf. (50). Thus *for each feasible* μ the minimum-norm slope is unique and equals $\beta(\mu)$. It remains to minimize $\|\beta(\mu)\|_2^2$ over $\mu \in \mathcal{M}$.

  - If $\mathbf{1}_n \notin \mathrm{col}(X)$, then $\mathcal{M} = \{\mu^\star\}$ and the unique minimizer is $(\hat{\mu}, \hat{\beta}) = (\mu^\star, \beta(\mu^\star))$.
  - If $\mathbf{1}_n \in \mathrm{col}(X)$, then for all $\mu \in \mathbb{R}$,

$$\begin{aligned} \|\beta(\mu)\|_2^2 &= (y - \mu \mathbf{1}_n)^\top (X^\dagger)^\top X^\dagger (y - \mu \mathbf{1}_n) \\ &= (y - \mu \mathbf{1}_n)^\top (XX^\top)^\dagger (y - \mu \mathbf{1}_n), \end{aligned} \tag{53}$$

    a strictly convex quadratic in μ because $\mathbf{1}_n \in \mathrm{range}(XX^\top)$ and $(XX^\top)^\dagger$ is positive definite on that range. Hence there is a *unique* minimizer $\hat{\mu}$ and thus a *unique* pair $(\hat{\mu}, \hat{\beta})$ with $\hat{\beta} = \beta(\hat{\mu})$.

  Therefore, the minimum-norm OLS pair $(\hat{\mu}, \hat{\beta})$ is *unique* among all LS minimizers.

- **Step 3: Identification and ratio-of-quadratics formulas.** Lastly, we identify the unique $\hat{\mu}$.

  - In the first case ($\mathbf{1}_n \notin \mathrm{col}(X)$), the unique $\hat{\mu}$ is exactly (52), i.e.

$$\hat{\mu} = \frac{\mathbf{1}_n^\top M_X y}{\mathbf{1}_n^\top M_X \mathbf{1}_n},$$

    with $\hat{\beta} = X^\dagger(y - \hat{\mu}\,\mathbf{1}_n)$. Note $\mathbf{1}_n^\top M_X \mathbf{1}_n = \|P_{\mathsf{N}(X^\top)} \mathbf{1}_n\|_2^2 > 0$.

– In the second case ($\mathbf{1}_n \in \mathrm{col}(X)$), minimizing the strictly convex quadratic objective (53) in μ yields

$$\hat{\mu} = \frac{\mathbf{1}_n^\top (XX^\top)^\dagger y}{\mathbf{1}_n^\top (XX^\top)^\dagger \mathbf{1}_n}, \qquad \hat{\beta} = X^\dagger (y - \hat{\mu}\,\mathbf{1}_n),$$

with $\mathbf{1}_n^\top (XX^\top)^\dagger \mathbf{1}_n = \|(XX^\top)^{\dagger 1/2}\mathbf{1}_n\|_2^2 > 0$.

Combining the two cases gives (9)–(10).

□

B.2 Proof of Proposition 2

Proof of Proposition 2. Given $\mathcal{F}_{t-1}$, let $\mathcal{K}$ — the feasible "completion set" that complements $\mathcal{S}_{t-1}^{\text{past}}$ to form a feasible treatment set $\mathcal{S}_1$ — be distributed as $\mathcal{K} \sim \mathrm{Unif}\binom{\mathcal{R}_{t-1}}{\ell_t}$. Then $\mathbb{E}[f \mid \mathcal{F}_{t-1}] = \mathbb{E}\big[f(\mathcal{S}_{t-1}^{\text{past}} \cup \mathcal{K})\big]$.

Conditioning on $\mathcal{F}_t = \sigma(\mathcal{F}_{t-1}, I)$ forces $I \in \mathcal{K}$, i.e., $\mathcal{K} = \{I\} \cup \mathcal{T}$ with $\mathcal{T} \sim \mathrm{Unif}\binom{\mathcal{R}_{t-1}\setminus\{I\}}{\ell_t - 1}$, so $\mathbb{E}[f \mid \mathcal{F}_t] = \mathbb{E}_{\mathcal{T}}\big[f(\mathcal{S}_{t-1}^{\text{prox}}(I, \mathcal{T}))\big]$.

Write $g(\mathcal{K}) \coloneqq f\big(\mathcal{S}_{t-1}^{\text{past}} \cup \mathcal{K}\big)$. Then

$$\mathbb{E}[f \mid \mathcal{F}_{t-1}] = \mathbb{E}\big[g(\mathcal{K}) \mid \mathcal{F}_{t-1}\big],$$
$$\mathbb{E}[f \mid \mathcal{F}_t] = \mathbb{E}\big[g(\mathcal{K}) \mid \mathcal{F}_{t-1},\, I \in \mathcal{K}\big].$$

Let $p_t \coloneqq \Pr(I \in \mathcal{K} \mid \mathcal{F}_{t-1}) = \ell_t / m_t$. By the law of total expectation,

$$\mathbb{E}[g(\mathcal{K}) \mid \mathcal{F}_{t-1}] = p_t\, \mathbb{E}[g(\mathcal{K}) \mid \mathcal{F}_{t-1},\, I \in \mathcal{K}] + (1 - p_t)\, \mathbb{E}[g(\mathcal{K}) \mid \mathcal{F}_{t-1},\, I \notin \mathcal{K}].$$

Hence,

$$\begin{aligned} D_t &= \mathbb{E}[g(\mathcal{K}) \mid \mathcal{F}_{t-1},\, I \in \mathcal{K}] - \mathbb{E}[g(\mathcal{K}) \mid \mathcal{F}_{t-1}] \\ &= (1 - p_t)\Big(\mathbb{E}[g(\mathcal{K}) \mid \mathcal{F}_{t-1},\, I \in \mathcal{K}] - \mathbb{E}[g(\mathcal{K}) \mid \mathcal{F}_{t-1},\, I \notin \mathcal{K}]\Big) \\ &= \frac{m_t - \ell_t}{m_t}\Big(\mathbb{E}[f \mid \mathcal{F}_{t-1},\, I \in \mathcal{K}] - \mathbb{E}[f \mid \mathcal{F}_{t-1},\, I \notin \mathcal{K}]\Big). \end{aligned}$$

To evaluate the difference $\mathbb{E}[f \mid \mathcal{F}_{t-1},\, I \in \mathcal{K}] - \mathbb{E}[f \mid \mathcal{F}_{t-1},\, I \notin \mathcal{K}]$ appearing above, we now construct a coupling under the $I \notin \mathcal{K}$ branch. Under the condition $I \notin \mathcal{K}$, write $\mathcal{K} = \{J\} \cup \mathcal{T}$ with $J \sim \mathrm{Unif}\big(\mathcal{R}_{t-1} \setminus \{I\}\big)$ and $\mathcal{T} \sim \mathrm{Unif}\big(\binom{\mathcal{R}_{t-1}\setminus\{I,J\}}{\ell_t - 1}\big)$. Pair the configurations $\mathcal{K}_I \coloneqq \{I\} \cup \mathcal{T}$ and $\mathcal{K}_J \coloneqq \{J\} \cup \mathcal{T}$ and note

$$f\big(\mathcal{S}_{t-1}^{\text{past}} \cup \mathcal{K}_I\big) - f\big(\mathcal{S}_{t-1}^{\text{past}} \cup \mathcal{K}_J\big) = -\,\Delta_{I\,J}\, f\big(\mathcal{S}_{t-1}^{\text{prox}}(I, \mathcal{T})\big),$$

Averaging completes the proof and yields (16). □

B.3 Proof of Proposition 3

Let $(\mathcal{F}_t)_{t=0}^{n_1}$ be the assignment–exposure filtration from (12), and for a measurable function f, define for $t = 1, \ldots, n_1$,

$$M_t \coloneqq \mathbb{E}\big[f(\mathcal{S}_1) \mid \mathcal{F}_t\big], \qquad \text{and} \qquad D_t \coloneqq M_t - M_{t-1}.$$

Then $M_0 = \mathbb{E}f(\mathcal{S}_1)$ and $M_{n_1} = f(\mathcal{S}_1)$, so $f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1) = \sum_{t=1}^{n_1} D_t$, with $\mathbb{E}[D_t \mid \mathcal{F}_{t-1}] = 0$. As in (18)–(19), define the predictable one–step variance and range

$$v_t^\star \coloneqq \mathrm{Var}(D_t \mid \mathcal{F}_{t-1}), \qquad \text{and} \qquad r_t^\star \coloneqq \sup |D_t|,$$

and their aggregates $V^\star \coloneqq \sum_{t=1}^{n_1} v_t^\star$ and $R^\star \coloneqq \max_{1 \le t \le n_1} r_t^\star$ (sup in the definition of r_t is taken over the law of $D_t \mid \mathcal{F}_{t-1}$). By construction, $v_t^\star, r_t^\star$ are $\mathcal{F}_{t-1}$–measurable and $|D_t| \le r_t^\star$ almost surely.

B.3.1 Background: Freedman's Inequality

We use the following standard form of Freedman's inequality [Freedman, 1975, Theorem 1.6].

Lemma B.1 (Freedman). *Let $(M_t, \mathcal{F}_t)$ be a martingale with differences $D_t = M_t - M_{t-1}$ satisfying $\mathbb{E}[D_t \mid \mathcal{F}_{t-1}] = 0$ and $|D_t| \le b$ almost surely. Set $V_t := \sum_{s=1}^{t} \mathbb{E}[D_s^2 \mid \mathcal{F}_{s-1}]$. Then for any $\varepsilon \ge 0$ and $\sigma^2 > 0$,*

$$\Pr\left(\exists t \ge 0 : \; M_t \ge \varepsilon \;\; \text{and} \;\; V_t \le \sigma^2\right) \le \exp\left(-\frac{\varepsilon^2/2}{\sigma^2 + b\,\varepsilon/3}\right).$$

A convenient self-normalized (empirical-Bernstein) corollary follows by plugging $\sigma^2 = V_t$ inside the event and optimizing ε.

Corollary B.1 (Self-normalized one-sided form). *Under Lemma B.1, for any $L > 0$,*

$$\Pr\left(\exists t \ge 0 : \; M_t \;\ge\; \sqrt{2V_t\,L} \;+\; \frac{b}{3}\,L\right) \;\le\; e^{-L}.$$

In particular, for any fixed horizon T,

$$\Pr\left(\; M_T \;\ge\; \sqrt{2V_T\,L} \;+\; \frac{b}{3}\,L\;\right) \;\le\; e^{-L}.$$

Proof (sketch). Fix $L > 0$ and consider $\mathcal{E} = \{\exists t : \; M_t \ge \sqrt{2V_t L} + (b/3)L\}$. For each $t \in \mathcal{E}$, set $\varepsilon = \sqrt{2V_t L} + (b/3)L$ and $\sigma^2 = V_t$ in Lemma B.1 to get $\Pr(\mathcal{E}) \le e^{-L}$. The fixed-time form is the corresponding sub-event. □

B.3.2 Completing Proof of Proposition 3

Proof of Proposition 3. Apply Corollary B.1 to the martingale (M_t) above. Since $|D_t| \le r_t^\star$ a.s. and $r_t^\star$ is $\mathcal{F}_{t-1}$–measurable, we may take $b = R^\star = \max_t r_t^\star$ (conditioning on the realized predictable sequence). Likewise, $V_{n_1} = V^\star$ because $\mathbb{E}[D_t^2 \mid \mathcal{F}_{t-1}] = \mathrm{Var}(D_t \mid \mathcal{F}_{t-1}) = v_t^\star$. Therefore, for any $L > 0$,

$$\Pr\left(\; M_{n_1} - M_0 \;\ge\; \sqrt{2V^\star\,L} \;+\; \frac{R^\star}{3}\,L\;\right) \;\le\; e^{-L}. \tag{54}$$

Recalling $M_{n_1} - M_0 = f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1)$, (54) yields the one-sided tail form

$$\Pr\left(\; f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1) \;\ge\; \sqrt{2V^\star\,L} \;+\; \frac{R^\star}{3}\,L\;\right) \;\le\; e^{-L}.$$

Applying the same bound to $-f$ and taking a union bound gives, for any $\delta \in (0,1)$ with $L_\delta := \log \frac{2}{\delta}$,

$$\Pr\left(\; |f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1)| \;\le\; \sqrt{2V^\star\,L_\delta} \;+\; \frac{R^\star}{3}\,L_\delta\;\right) \;\ge\; 1 - \delta,$$

which is the stated concentration inequality in Proposition 3. □

Remark B.1 (Predictability and sharpness). The parameters $(V^\star, R^\star)$ are *predictable* (i.e., $\mathcal{F}_{t-1}$–measurable stepwise) and enter the bound *inside the event*, which is legitimate because the proof conditions on the realized predictable sequence and then applies Freedman with those bounds (a standard empirical-Bernstein device). This retains pathwise sharpness relative to range-based, non-adaptive envelopes.

B.4 Proof of Proposition 4

Although Proposition 4 follows from the discussions in the main text (Section 3.2), we provide a formal proof here for completeness.

Proof of Proposition 4. Throughout, expectations $\mathbb{E}[\cdot]$ and L^2 inner products $\langle f, g\rangle := \mathbb{E}[f(\mathcal{S})g(\mathcal{S})]$ are with respect to the uniform law π on $\binom{[n]}{m}$, for which the one-swap kernel P on $J(n,m)$ is reversible. Recall the generator $L = P - \mathrm{Id}$ and the carré du champ

$$\Gamma(f)(\mathcal{S}) \;=\; \frac{1}{2}\,\mathbb{E}\big[(f(\mathcal{S}') - f(\mathcal{S}))^2 \mid \mathcal{S}\big],$$

as defined in (22)–(23).

Assume $\mathrm{Var}(f) > 0$. Let $g := f - \mathbb{E}f$, and let $\lambda^\star$ be defined according to (24). By the Dirichlet identity and the computation in (26),

$$\lambda^\star = \frac{\mathbb{E}\Gamma(f)}{\mathrm{Var}(f)} > 0.$$

Define the uncentered Stein remainder

$$\psi_u^\star := Lf + \lambda^\star f.$$

Since $\mathbb{E}Lf = 0$, we have

$$\mathbb{E}\psi_u^\star = \lambda^\star \mathbb{E}f.$$

Therefore, by Cauchy–Schwarz inequality, we obtain (27) as follows, completing the proof:

$$|\mathbb{E}f| = \frac{|\mathbb{E}\psi_u^\star|}{\lambda^\star} \le \frac{\sqrt{\mathbb{E}(\psi_u^\star)^2}}{\lambda^\star} = \frac{\sqrt{\mathbb{E}\big(Lf + \lambda^\star f\big)^2}}{\lambda^\star}.$$

□

C DEFERRED PROOFS FROM SECTION 4

C.1 Proof of Theorem 5

Proof of Theorem 5. Let $f(\mathcal{S}) = \hat{\tau}(\mathcal{S}) - \tau$ and recall the assignment-exposure martingale $(M_t, \mathcal{F}_t)$ from (12)–(13) with increments $D_t = M_t - M_{t-1}$. By the triangle inequality,

$$|f(\mathcal{S}_1)| \le |f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1)| + |\mathbb{E}f(\mathcal{S}_1)|.$$

For the fluctuation term $|f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1)|$, apply Proposition 3 in its two-sided form (21), cf. Remark 3.5. With $V^\star, R^\star$ defined in (29) as the aggregates of the predictable parameters $v_t^\star = \mathrm{Var}(D_t \mid \mathcal{F}_{t-1})$ and $r_t^\star = \sup|D_t|$ from (18)–(19), we obtain:

$$\Pr\left(|f(\mathcal{S}_1) - \mathbb{E}f(\mathcal{S}_1)| \le \sqrt{2\,V^\star\,\mathsf{L}_\delta} + \frac{R^\star}{3}\,\mathsf{L}_\delta\right) \ge 1 - \delta, \qquad \text{where} \qquad \mathsf{L}_\delta = \log\frac{2}{\delta}. \tag{55}$$

For the bias term $|\mathbb{E}f(\mathcal{S}_1)|$, recall the Stein bias bound from Proposition 4.

- If $\mathrm{Var}(f) = 0$, then f is constant under the design, so $|f(\mathcal{S}_1)| = |\mathbb{E}f|$. In this degenerate case we interpret $B^\star$ as $|\mathbb{E}f|$.
- Otherwise, if $\mathrm{Var}(f) > 0$, then Proposition 4 yields

$$|\mathbb{E}f(\mathcal{S}_1)| \le B^\star = \frac{\sqrt{\mathbb{E}\big(Lf + \lambda^\star f\big)^2}}{\lambda^\star}, \qquad \lambda^\star = \frac{\mathbb{E}\Gamma(f)}{\mathrm{Var}(f)}. \tag{56}$$

Combining the two displays (55) and (56) yields (32). Note the bias bound is deterministic (no probability), so the overall success probability remains $1 - \delta$ from (55). □

C.2 Proof of Proposition 6

Proof of Proposition 6. Fix $\mathcal{S}_1 \in \binom{[n]}{n_1}$ and $(i,j) \in \mathcal{S}_1 \times \mathcal{S}_0$. Let $\mathcal{S}_1' = \mathcal{S}_1^{(i\leftrightarrow j)} = (\mathcal{S}_1 \setminus \{i\}) \cup \{j\}$ and $\mathcal{S}_0' = \mathcal{S}_0^{(j\leftrightarrow i)} = (\mathcal{S}_0 \setminus \{j\}) \cup \{i\}$ denote the swapped treatment and control sets, respectively. By definition of the regression-adjusted estimator,

$$\begin{aligned}\Delta_{ij}\,\hat{\tau}_{\mathrm{OLS}}(\mathcal{S}_1) &= \hat{\tau}_{\mathrm{OLS}}(\mathcal{S}_1') - \hat{\tau}_{\mathrm{OLS}}(\mathcal{S}_1) \\ &= \Big(\hat{\mu}_1(\mathcal{S}_1') - \hat{\mu}_1(\mathcal{S}_1)\Big) - \Big(\hat{\mu}_0(\mathcal{S}_0') - \hat{\mu}_0(\mathcal{S}_0)\Big).\end{aligned}$$

Insert the intermediate sets $\mathcal{S}_1 \setminus \{i\}$ and $\mathcal{S}_0 \setminus \{j\}$ and telescope:

$$\hat{\mu}_1(\mathcal{S}_1') - \hat{\mu}_1(\mathcal{S}_1) = \underbrace{\big(\hat{\mu}_1(\mathcal{S}_1 \setminus \{i\}) - \hat{\mu}_1(\mathcal{S}_1)\big)}_{\Delta_1^{\mathrm{del}}(i;\,\mathcal{S}_1)} + \underbrace{\big(\hat{\mu}_1\big((\mathcal{S}_1 \setminus \{i\}) \cup \{j\}\big) - \hat{\mu}_1(\mathcal{S}_1 \setminus \{i\})\big)}_{\Delta_1^{\mathrm{ins}}(j;\,\mathcal{S}_1\setminus\{i\})},$$

and likewise,

$$\hat{\mu}_0(\mathcal{S}_0') - \hat{\mu}_0(\mathcal{S}_0) = \underbrace{\big(\hat{\mu}_0(\mathcal{S}_0 \setminus \{j\}) - \hat{\mu}_0(\mathcal{S}_0)\big)}_{\Delta_0^{\mathrm{del}}(j;\,\mathcal{S}_0)} + \underbrace{\big(\hat{\mu}_0\big((\mathcal{S}_0 \setminus \{j\}) \cup \{i\}\big) - \hat{\mu}_0(\mathcal{S}_0 \setminus \{j\})\big)}_{\Delta_0^{\mathrm{ins}}(i;\,\mathcal{S}_0\setminus\{j\})}.$$

Substituting these identities into the expression for $\Delta_{ij}\,\hat{\tau}_{\mathrm{OLS}}(\mathcal{S}_1)$ yields

$$\Delta_{ij}\,\hat{\tau}_{\mathrm{OLS}}(\mathcal{S}_1) = \Big[\Delta_1^{\mathrm{del}}(i;\,\mathcal{S}_1) + \Delta_1^{\mathrm{ins}}(j;\,\mathcal{S}_1 \setminus \{i\})\Big] - \Big[\Delta_0^{\mathrm{del}}(j;\,\mathcal{S}_0) + \Delta_0^{\mathrm{ins}}(i;\,\mathcal{S}_0 \setminus \{j\})\Big],$$

which is the equation display (34). □

C.3 Proof of Theorem 7

C.3.1 Background: Column Append and Deletion Identities for Pseudoinverse

Here we record the standard rank-one column-append/deletion formulas of Greville [1960]. We refer interested readers to a textbook by Ben-Israel and Greville [2003, Chapters 5 and 7] for more background and detailed expositions. The statements below are quoted with minor rephrasing for clarity.

Lemma C.1 (Column append identity; Greville [1960, Section 4], summarized and rephrased)**.** *Let $A \in \mathbb{R}^{n\times m}$ and $a \in \mathbb{R}^n$. Set $d \coloneqq A^\dagger a$ and $c = a - Ad = (I_n - AA^\dagger)a$. Then*

$$\begin{bmatrix} A & a \end{bmatrix}^\dagger = \begin{bmatrix} A^\dagger - db \\ b \end{bmatrix} \qquad \textit{where} \qquad b = \begin{cases} c^\dagger, & \textit{if } c \neq 0, \\ (1 + d^\top d)^{-1} d^\top A^\dagger, & \textit{if } c = 0, \end{cases} \tag{57}$$

where for a nonzero column vector v, $v^\dagger \equiv v^\top/(v^\top v)$ denotes its Moore–Penrose pseudoinverse.

Proof of Lemma C.1. This statement is a summary of Greville [1960, Section 4] as a lemma, with minor rephrasing. The recursive expression in (57) is obtained by combining Eqs. (8), (9), (11), and (16) or (22), depending on whether $c = 0$ or not. □

The next column deletion identity, while not explicitly stated in Greville [1960], is an immediate algebraic reverse of Lemma C.1.

Lemma C.2 (Column deletion identity)**.** *Let $\widetilde{A} \coloneqq \begin{bmatrix} A & a \end{bmatrix} \in \mathbb{R}^{n\times(m+1)}$ and partition $\widetilde{A}^\dagger = \begin{bmatrix} B \\ b \end{bmatrix}$ with $B \in \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^{1\times n}$. Then*

$$A^\dagger = B + db \qquad \textit{where} \qquad d = \begin{cases} B\left(a - b^\dagger\right), & \textit{if } bA = 0, \\ (1 - ba)^{-1} Ba, & \textit{if } bA \neq 0. \end{cases} \tag{58}$$

Proof of Lemma C.2. By Lemma C.1, writing $d \coloneqq A^\dagger a$ and $c \coloneqq (I_n - AA^\dagger)a$, we have

$$\widetilde{A}^\dagger = \begin{bmatrix} A^\dagger - d\,b \\ b \end{bmatrix}, \qquad \text{with} \quad b = \begin{cases} c^\dagger, & c \neq 0, \\ (1 + d^\top d)^{-1} d^\top A^\dagger, & c = 0. \end{cases}$$

Hence $B = A^\dagger - d\,b$, i.e., $A^\dagger = B + d\,b$. It remains to express $d = A^\dagger a$ in terms of B, a, and b.

- *Case 1:* $bA = 0$. This corresponds to the $c \neq 0$ branch in Lemma C.1, where $b = c^\dagger$ and thus $c = b^\dagger$. Therefore, $bA = c^\dagger A = 0$. Using $a = Ad + c$,

$$\begin{aligned} d &= A^\dagger a = A^\dagger A d = A^\dagger (a - c) \\ &\overset{(*)}{=} (A^\dagger - db)\,(a - c) = B\,(a - c) \\ &= B\,(a - b^\dagger). \end{aligned}$$

  Here, (*) follows from the observation that since $a = Ad + c$ and $b = c^\dagger$, $ba = b(Ad + c) = 0 + bc = 1$, and therefore, $b(a - c) = ba - bc = 1 - 1 = 0$.

- *Case 2:* $bA \neq 0$. Multiplying $B = A^\dagger - db$ by a from the right:

$$Ba = A^\dagger a - d\,ba = d\,(1 - ba),$$

  hence $d = (1 - ba)^{-1} Ba$. Here, $c = 0$, and Lemma C.1 yields $b = (1 + d^\top d)^{-1} d^\top A^\dagger$; so $ba = (1 + d^\top d)^{-1} d^\top d \in [0, 1)$ and thus, $1 - ba = (1 + d^\top d)^{-1} > 0$ and the inverse $(1 - ba)^{-1}$ is well defined.

Substituting the corresponding expression for d into $A^\dagger = B + db$ yields (58). We note that the identities invoked are exactly those consolidated in Lemma C.1 from Greville [1960, Section 4, Eqs. (8), (9), (11), (16), (22)]. □

Row insertion/deletion counterparts follow by transposition. In Section C.3.2, we apply these row-wise versions with $A = X_\alpha^\top$ and $a = x_i^\top$ (or $a = x_j^\top$) when deleting (or appending) a row (=unit) in X_α.

C.3.2 Row Deletion and Row Insertion Lemmas

Throughout this section, we use the Q-formulation of the OLS intercept $\hat{\mu}$ from Proposition 1. For $X \in \mathbb{R}^{n \times p}$ and $y \in \mathbb{R}^n$, the minimum-norm OLS intercept is

$$\hat{\mu} = \frac{\mathbf{1}_n^\top Q_X\, y}{\mathbf{1}_n^\top Q_X\, \mathbf{1}_n}, \qquad \text{where} \qquad Q_X = \begin{cases} M_X := I_n - XX^\dagger, & \text{if } \mathbf{1}_n \notin \mathrm{col}(X), \\ K_X := (XX^\top)^\dagger, & \text{if } \mathbf{1}_n \in \mathrm{col}(X). \end{cases}$$

Here, the denominator $\mathbf{1}^\top Q_X \mathbf{1} > 0$ in both cases.

Notation. For a symmetric positive semidefinite matrix Q, we write $\langle u, v \rangle_Q := u^\top Q v$ and $\|u\|_Q := \sqrt{u^\top Q u}$. Given $\mathcal{S} \subset [n]$ and arm $\alpha \in \{0, 1\}$, we abbreviate

$$Q_{\mathcal{S}} := Q_{X_{\mathcal{S}}}, \qquad \hat{\mu}_\alpha(\mathcal{S}) := \frac{\langle \mathbf{1},\, y_{\mathcal{S}}^{(\alpha)} \rangle_{Q_{\mathcal{S}}}}{\|\mathbf{1}\|_{Q_{\mathcal{S}}}^2}, \qquad \tilde{y}_{\mathcal{S}}^{(\alpha)} := y_{\mathcal{S}}^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S}) \cdot \mathbf{1},$$

so that $\langle \mathbf{1},\, \tilde{y}_{\mathcal{S}}^{(\alpha)} \rangle_{Q_{\mathcal{S}}} = 0$. When the dimension is clear from context, we write $\mathbf{1} = \mathbf{1}_n$ (or $\mathbf{1}_{|\mathcal{S}|}$) and drop the subscript.

Rank-1 Ratio Update for Bookkeeping. We isolate a purely algebraic identity for how the intercept ratio changes under a rank–one perturbation of Q. This will be applied with the row-deletion/row-insertion updates.

Lemma C.3. *For any symmetric matrix Q and vector y, define*

$$\mu_y(Q) := \begin{cases} \frac{\mathbf{1}^\top Q y}{\mathbf{1}^\top Q \mathbf{1}}, & \text{if } \mathbf{1}^\top Q \mathbf{1} > 0, \\ 0, & \text{if } \mathbf{1}^\top Q \mathbf{1} = 0. \end{cases}, \qquad \text{and} \qquad y_Q^{\mathrm{c}} := y - \mu_y(Q) \cdot \mathbf{1}.$$

For any scalar κ and vector v, let $Q' = Q + \kappa (Qv)(Qv)^\top$. Then

$$\mu_y(Q') - \mu_y(Q) = \frac{\kappa \langle \mathbf{1}, v \rangle_Q\ \langle v,\, y_Q^{\mathrm{c}} \rangle_Q}{\mathbf{1}^\top Q' \mathbf{1}}.$$

Proof of Lemma C.3. Observe that

$$\mathbf{1}^\top Q' y = \mathbf{1}^\top Q y + \kappa \langle \mathbf{1}, v\rangle_Q \, \langle v, y\rangle_Q, \qquad \text{and} \qquad \mathbf{1}^\top Q' \mathbf{1} = \mathbf{1}^\top Q \mathbf{1} + \kappa \langle \mathbf{1}, v\rangle_Q^2.$$

Therefore,

$$\begin{aligned}
\mu_y(Q') - \mu_y(Q) &= \frac{\mathbf{1}^\top Q y + \kappa \langle \mathbf{1}, v\rangle_Q \langle v, y\rangle_Q}{\mathbf{1}^\top Q \mathbf{1} + \kappa \langle \mathbf{1}, v\rangle_Q^2} - \frac{\mathbf{1}^\top Q y}{\mathbf{1}^\top Q \mathbf{1}} \\
&= \frac{\kappa \langle \mathbf{1}, v\rangle_Q \, \left(\langle v, y\rangle_Q - \mu_y(Q) \langle v, \mathbf{1}\rangle_Q \right)}{\mathbf{1}^\top Q' \mathbf{1}} \\
&= \frac{\kappa \langle \mathbf{1}, v\rangle_Q \, \left\langle v, y - \mu_y(Q) \cdot \mathbf{1} \right\rangle_Q}{\mathbf{1}^\top Q' \mathbf{1}} \\
&= \frac{\kappa \langle \mathbf{1}, v\rangle_Q \, \langle v, \, y_Q^{\mathrm{c}} \rangle_Q}{\mathbf{1}^\top Q' \mathbf{1}}.
\end{aligned}$$

□

Remark C.1 (Useful variant without Q-geometry). The proof of Lemma C.3 goes through verbatim if Q' is updated by a general rank-one form $Q' = Q + \kappa\, r r^\top$ (not necessarily $r = Qv$), but with $\langle \mathbf{1}, v\rangle_Q$ and $\langle v, \cdot\rangle_Q$ replaced by $\mathbf{1}^\top r$ and $r^\top(\cdot)$.

Row Deletion: Exact Identity. We apply Lemma C.3 to the Greville deletion (Lemma C.2 after transpose) to obtain a closed form for the atomic deletion change of the intercept.

Given $\mathcal{S}$ with $s = |\mathcal{S}|$, let $e_i \in \mathbb{R}^s$ denote the coordinate vector of unit i within $\mathcal{S}$ (i.e., the indicator of i's position in $\mathcal{S}$). For $i \in \mathcal{S}$, define

$$\Phi_i^{\mathrm{del}}(\mathcal{S}) := \|\mathbf{1}\|_{Q_\mathcal{S}}^2 \, \|e_i\|_{Q_\mathcal{S}}^2 - \langle e_i, \mathbf{1}\rangle_{Q_\mathcal{S}}^2 \geq 0.$$

Lemma C.4 (Row deletion identity)**.** *Let $X \in \mathbb{R}^{n\times p}$, $y^{(1)}, y^{(0)} \in \mathbb{R}^n$, $\alpha \in \{0,1\}$, $\mathcal{S} \in \binom{[n]}{m}$, and $i \in \mathcal{S}$. Let $R_i \in \{0,1\}^{(m-1)\times m}$ be the row-deletion matrix that removes the coordinate of i from vectors indexed by $\mathcal{S}$. Let $Q_\mathcal{S} = Q_{X_\mathcal{S}}$, $\hat{\mu}_\alpha = \hat{\mu}_\alpha(\mathcal{S})$, and*

$$\tilde{y}_\mathcal{S}^{(\alpha)} = y_\mathcal{S}^{(\alpha)} - \hat{\mu}_\alpha \mathbf{1}.$$

Then the deletion change satisfies the exact ratio identity

$$\Delta_\alpha^{\mathrm{del}}(i; \mathcal{S}) = \frac{\mathbf{1}^\top Q_{X_{\mathcal{S}\setminus\{i\}}} R_i \tilde{y}_\mathcal{S}^{(\alpha)}}{\mathbf{1}^\top Q_{X_{\mathcal{S}\setminus\{i\}}} \mathbf{1}}. \tag{59}$$

Moreover, in the regular deletion case where $q_i := e_i^\top Q_\mathcal{S} e_i > 0$ and the zero-padded update satisfies

$$R_i^\top Q_{X_{\mathcal{S}\setminus\{i\}}} R_i = Q_\mathcal{S} - \frac{Q_\mathcal{S} e_i e_i^\top Q_\mathcal{S}}{q_i},$$

if $\Phi_i^{\mathrm{del}}(\mathcal{S}) > 0$, then

$$\Delta_\alpha^{\mathrm{del}}(i; \mathcal{S}) = -\frac{\langle e_i, \mathbf{1}\rangle_{Q_\mathcal{S}} \langle e_i, \tilde{y}_\mathcal{S}^{(\alpha)}\rangle_{Q_\mathcal{S}}}{\Phi_i^{\mathrm{del}}(\mathcal{S})}. \tag{60}$$

Proof of Lemma C.4. The identity (59) follows directly from the definition:

$$\hat{\mu}_\alpha(\mathcal{S} \setminus \{i\}) = \frac{\mathbf{1}^\top Q_{X_{\mathcal{S}\setminus\{i\}}} y_{\mathcal{S}\setminus\{i\}}^{(\alpha)}}{\mathbf{1}^\top Q_{X_{\mathcal{S}\setminus\{i\}}} \mathbf{1}}.$$

Subtracting $\hat{\mu}_\alpha(\mathcal{S})$ from both sides gives (59).

Next, we prove the regular closed form (60). Let $Q = Q_\mathcal{S}$, $q_i = e_i^\top Q e_i$, $a_i = e_i^\top Q \mathbf{1}$, and $b_i = e_i^\top Q \tilde{y}_\mathcal{S}^{(\alpha)}$. By the assumed regular deletion update,

$$Q_-^\uparrow := R_i^\top Q_{X_{\mathcal{S}\setminus\{i\}}} R_i = Q - \frac{Q e_i e_i^\top Q}{q_i}.$$

Since $\mathbf{1}^\top Q\tilde{y}_{\mathcal{S}}^{(\alpha)} = 0$, the numerator in (59), written in the original s-dimensional coordinates, equals

$$\mathbf{1}^\top Q_-^\uparrow \tilde{y}_{\mathcal{S}}^{(\alpha)} = -\frac{a_i b_i}{q_i}.$$

The denominator equals

$$\mathbf{1}^\top Q_-^\uparrow \mathbf{1} = \mathbf{1}^\top Q\mathbf{1} - \frac{a_i^2}{q_i} = \frac{\|\mathbf{1}\|_Q^2 \|e_i\|_Q^2 - \langle e_i, \mathbf{1}\rangle_Q^2}{q_i} = \frac{\Phi_i^{\mathrm{del}}(\mathcal{S})}{q_i}.$$

Dividing the two displays yields (60), which completes the proof. □

Remark C.2 (Explicit appearance of $y_i^{(\alpha)}$). Intuitively, we expect $\Delta_\alpha^{\mathrm{del}}(i;\mathcal{S})$ to depend on the covariate x_i and the outcome $y_i^{(\alpha)}$ at the deleted unit $i \in \mathcal{S}$. The factor

$$\langle e_i,\, y_{\mathcal{S}}^{(\alpha),c}\rangle_{Q_{\mathcal{S}}} = e_i^\top Q_{\mathcal{S}}\big(y_{\mathcal{S}}^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S})\mathbf{1}\big) = (Q_{\mathcal{S}} y_{\mathcal{S}}^{(\alpha)})_i - \hat{\mu}_\alpha(\mathcal{S})\,(Q_{\mathcal{S}}\mathbf{1})_i$$

enters (60) and depends *explicitly* on $y_i^{(\alpha)}$ (together with the neighboring entries weighted by $Q_{\mathcal{S}}$). Thus the deletion sensitivity reflects both the covariate geometry (via $Q_{\mathcal{S}}$ and $\langle e_i, \mathbf{1}\rangle_{Q_{\mathcal{S}}}$) and the outcome at unit i.

Row Insertion: Exact Identity. Next, we combine the Greville column-append identity (Lemma C.1, transposed to rows) with the bookkeeping lemma (Lemma C.3) to derive the insertion change. This makes the dependence on $(x_j, y_j^{(\alpha)})$ explicit and holds uniformly across the two Q-branches in Proposition 1.

Lemma C.5 (Row insertion identity)**.** *Let $X \in \mathbb{R}^{n\times p}$, $y^{(1)}, y^{(0)} \in \mathbb{R}^n$, $\mathcal{S} \in \binom{[n]}{m}$, $\alpha \in \{0,1\}$, and $j \notin \mathcal{S}$. Let $\mathcal{S}^+ = \mathcal{S} \cup \{j\}$, and order $\mathcal{S}^+$ so that j is the last row. Partition*

$$Q_{X_{\mathcal{S}^+}} = \begin{bmatrix} Q_{00}^+ & q_{0j}^+ \\ (q_{0j}^+)^\top & q_{jj}^+ \end{bmatrix}.$$

With

$$r_j^{(\alpha)} := y_j^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S}), \qquad \tilde{y}_{\mathcal{S}}^{(\alpha)} = y_{\mathcal{S}}^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S})\mathbf{1},$$

the atomic insertion change satisfies the exact block-ratio formula

$$\Delta_\alpha^{\mathrm{ins}}(j;\mathcal{S}) = \frac{\mathbf{1}^\top Q_{00}^+ \tilde{y}_{\mathcal{S}}^{(\alpha)} + (q_{0j}^+)^\top \tilde{y}_{\mathcal{S}}^{(\alpha)} + \big(\mathbf{1}^\top q_{0j}^+ + q_{jj}^+\big) r_j^{(\alpha)}}{\mathbf{1}^\top Q_{00}^+ \mathbf{1} + 2\,\mathbf{1}^\top q_{0j}^+ + q_{jj}^+}. \tag{61}$$

In addition, the following simplifications hold.

1. *(Regular K-branch case) Suppose $Q_{X_{\mathcal{S}}} = K_{X_{\mathcal{S}}}$, and define*

$$\sigma_j^2(\mathcal{S}) = \|x_j\|_2^2 - \|X_{\mathcal{S}} x_j\|_{Q_{\mathcal{S}}}^2 > 0.$$

   If $\sigma_j^2(\mathcal{S}) > 0$, then

$$\Delta_\alpha^{\mathrm{ins}}(j;\mathcal{S}) = \frac{\big(1-\alpha_j(\mathcal{S})\big)\Big(y_j^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S}) - \langle X_{\mathcal{S}} x_j, \tilde{y}_{\mathcal{S}}^{(\alpha)}\rangle_{Q_{\mathcal{S}}}\Big)}{\|\mathbf{1}\|_{Q_{\mathcal{S}}}^2 \sigma_j^2(\mathcal{S}) + \big(1-\alpha_j(\mathcal{S})\big)^2}, \tag{62}$$

   where $\alpha_j(\mathcal{S}) = \langle \mathbf{1}, X_{\mathcal{S}} x_j\rangle_{Q_{\mathcal{S}}}$.

2. *(Regular M-branch case) Suppose $Q_{X_{\mathcal{S}}} = M_{X_{\mathcal{S}}}$, and let*

$$A_{\mathcal{S}} := X_{\mathcal{S}}^\top, \qquad d_j := A_{\mathcal{S}}^\dagger x_j, \qquad c_j := x_j - A_{\mathcal{S}} d_j.$$

   (a) If $c_j \neq 0$, then $\Delta_\alpha^{\mathrm{ins}}(j;\mathcal{S}) = 0$.

(b) If $c_j = 0$, then

$$\Delta_\alpha^{\text{ins}}(j;\mathcal{S}) = \frac{(1-\eta_j)\big(y_j^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S}) - \beta_j\big)}{\|\mathbf{1}\|_{Q_\mathcal{S}}^2(1+h_j) + (1-\eta_j)^2}, \tag{63}$$

where

$$h_j = \|d_j\|_2^2, \qquad \eta_j = \mathbf{1}^\top d_j, \qquad \beta_j = d_j^\top \tilde{y}_\mathcal{S}^{(\alpha)}.$$

Proof of Lemma C.5. The block formula (61) follows directly from the ratio representation: if we write $Q^+ = Q_{X_{\mathcal{S}^+}}$ and $y^+ = (y_\mathcal{S}^{(\alpha)}, y_j^{(\alpha)})$, then

$$\hat{\mu}_\alpha(\mathcal{S}^+) - \hat{\mu}_\alpha(\mathcal{S}) = \frac{[\mathbf{1};1]^\top Q^+\big[y_\mathcal{S}^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S})\mathbf{1};\, y_j^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S})\big]}{[\mathbf{1};1]^\top Q^+[\mathbf{1};1]}.$$

Expanding this expression under the displayed block partition of Q^+ gives (61).

Next, we derive the regular K-branch formula. Let $u_j = X_\mathcal{S} x_j$, $Q = K_{X_\mathcal{S}}$, and $\sigma^2 = \|x_j\|_2^2 - u_j^\top Q u_j > 0$. By the row-append Moore–Penrose update, equivalently the transposed $c \neq 0$ branch of Greville's identity in Lemma C.1, the appended Gram pseudoinverse satisfies

$$Q_{X_{\mathcal{S}^+}} = \begin{bmatrix} Q + \dfrac{Q u_j u_j^\top Q}{\sigma^2} & -\dfrac{Q u_j}{\sigma^2} \\ -\dfrac{u_j^\top Q}{\sigma^2} & \dfrac{1}{\sigma^2} \end{bmatrix}.$$

When $X_\mathcal{S} X_\mathcal{S}^\top$ is invertible this reduces to the usual Schur-complement inverse update. Substituting this block matrix into (61), and using $\mathbf{1}^\top Q \tilde{y}_\mathcal{S}^{(\alpha)} = 0$, yields (62).

Lastly, consider the regular M-branch. Let $Q = M_{X_\mathcal{S}}$, $A_\mathcal{S} = X_\mathcal{S}^\top$, $d_j = A_\mathcal{S}^\dagger x_j$, and $c_j = x_j - A_\mathcal{S} d_j$. Greville's row-append formula for the residual-maker matrix gives two cases. If $c_j \neq 0$, then

$$Q_{X_{\mathcal{S}^+}} = \begin{bmatrix} Q & 0 \\ 0 & 0 \end{bmatrix},$$

so (61) gives $\Delta_\alpha^{\text{ins}}(j;\mathcal{S}) = 0$. If $c_j = 0$, then

$$Q_{X_{\mathcal{S}^+}} = \begin{bmatrix} Q + \dfrac{d_j d_j^\top}{1+\|d_j\|_2^2} & -\dfrac{d_j}{1+\|d_j\|_2^2} \\ -\dfrac{d_j^\top}{1+\|d_j\|_2^2} & \dfrac{1}{1+\|d_j\|_2^2} \end{bmatrix}.$$

Substituting this expression into (61) yields (63). In particular, when $X_\mathcal{S}$ has no covariates, $d_j = 0$, and (63) reduces to

$$\Delta_\alpha^{\text{ins}}(j;\mathcal{S}) = \frac{y_j^{(\alpha)} - \hat{\mu}_\alpha(\mathcal{S})}{|\mathcal{S}|+1},$$

as required for insertion into a sample mean. □

C.3.3 Completing Proof of Theorem 7

Proof of Theorem 7. The regular deletion identity in Theorem 7 is exactly (60) from Lemma C.4. The regular K-branch insertion identity is exactly (62) from Lemma C.5. The full exact branch-specific identities are given by (59), (61), (62), and (63). Combining these atomic identities with the paired deletion–insertion telescope in Proposition 6 gives the oracle one-swap sensitivity of OLS–RA. □

D NUMERICAL EXPERIMENTS

This appendix complements the theory by describing *(i)* Monte Carlo (MC) approximation of the oracle quantities $(V^\star, R^\star, B^\star)$ (Section D.1) and *(ii)* small-scale simulations that probe the sharpness of the finite-sample CIs in Theorem 5 across classical and over-parameterized regimes (Section D.2). The Python code used to generate the Monte Carlo results reported here is available at https://github.com/dogyoons/regadj_finitesample.

D.1 Monte Carlo Approximation of Oracle Parameters

This section explains how the oracle quantities $(V^\star, R^\star, B^\star)$ entering Theorem 5 are computed in the experiments. Throughout, we work in the *oracle* setting, meaning that the full finite-population objects $(X, y^{(0)}, y^{(1)})$ are treated as known. The data-driven proxies discussed in Appendices A.3–A.4 are not used here.

Given an ATE estimator $\hat{\tau}$ and a treatment set $\mathcal{S}$, write $f(\mathcal{S}) = \hat{\tau}(\mathcal{S}) - \tau$. The finite-sample radius in Theorem 5 depends on three population quantities:

$$V^\star = \sum_{t=1}^{n_1} v_t^\star, \qquad R^\star = \max_t r_t^\star, \qquad B^\star = \frac{\sqrt{\mathbb{E}(Lf + \lambda^\star f)^2}}{\lambda^\star},$$

where $v_t^\star = \mathrm{Var}(D_t \mid \mathcal{F}_{t-1})$, $r_t^\star = \sup |D_t|$, and $D_t = M_t - M_{t-1}$ is the Doob-martingale increment along the assignment-exposure filtration. By Proposition 2, these quantities are determined by one-swap sensitivities $\Delta_{ij} f$ and their conditional averages.

For OLS–RA, exact evaluation of these expectations is computationally infeasible in general, because it would require summing over many treatment assignments, reveal states, and ordered swap pairs. Accordingly, in Experiment 2 we approximate $(V^\star, R^\star, B^\star)$ by Monte Carlo estimators $(\widehat{V}^\star, \widehat{R}^\star, \widehat{B}^\star)$. For DiM, by contrast, $V^\star$ and $R^\star$ admit closed forms under the remaining-pool reveal law, so we use these exact oracle quantities in Experiment 1. Throughout, OLS–RA swap effects are evaluated using the paired deletion–insertion identities from Proposition 6 and Theorem 7, so no OLS re-fitting is required.

D.1.1 Monte Carlo Estimation of Concentration Parameters

Recollecting Setup and Notation. Fix a realized assignment $(\mathcal{S}_1, \mathcal{S}_0)$ and an independent reveal order Π of the treated units. Recall from (14) that at step t, we let

$$\mathcal{S}_{t-1}^{\mathrm{past}} = \{\Pi(1), \ldots, \Pi(t-1)\}, \qquad \mathcal{R}_{t-1} = [n] \setminus \mathcal{S}_{t-1}^{\mathrm{past}},$$

and let $\ell_t = n_1 - (t-1)$. For a candidate unit $i \in \mathcal{R}_{t-1}$, the proxy treatment set is defined in (15) as

$$\mathcal{S}_{t-1}^{\mathrm{prox}}(i, \mathcal{T}) = \mathcal{S}_{t-1}^{\mathrm{past}} \cup \{i\} \cup \mathcal{T}, \qquad \mathcal{T} \sim \mathrm{Unif}\binom{\mathcal{R}_{t-1} \setminus \{i\}}{\ell_t - 1}.$$

In the implementation, $\mathcal{I}_t \subseteq \mathcal{R}_{t-1}$ is either the full remaining pool or a uniform subsample of size B_i. For each $i \in \mathcal{I}_t$, we draw B_{cond} completion sets $\mathcal{T}_{b,i}$, and for each completion either enumerate all admissible controls or sample B_J admissible controls uniformly without replacement. Here B_i, B_{cond}, and B_J are the MC budgets controlling the candidate set, completion draws, and control draws, respectively.

Stepwise Monte Carlo Estimate. For each $i \in \mathcal{I}_t$, define

$$\widehat{\zeta}_t(i) = \frac{1}{B_{\mathrm{cond}}} \sum_{b=1}^{B_{\mathrm{cond}}} \left[\frac{1}{|\mathcal{J}_{b,i}|} \sum_{J \in \mathcal{J}_{b,i}} \Delta_{iJ} f\big(\mathcal{S}_{t-1}^{\mathrm{prox}}(i, \mathcal{T}_{b,i})\big) \right],$$

where $\mathcal{J}_{b,i} \subseteq \mathcal{R}_{t-1} \setminus (\{i\} \cup \mathcal{T}_{b,i})$ is either the full admissible control set or a uniform subsample of size B_J. We also define the within-i completion variance

$$\widehat{s}_t^2(i) := \widehat{\mathrm{Var}}_{b=1,\ldots,B_{\mathrm{cond}}} \left(\frac{1}{|\mathcal{J}_{b,i}|} \sum_{J \in \mathcal{J}_{b,i}} \Delta_{iJ} f\big(\mathcal{S}_{t-1}^{\mathrm{prox}}(i, \mathcal{T}_{b,i})\big) \right).$$

The reported experiments use the direct plug-in estimators

$$\widehat{V^\star} = \sum_{t=1}^{n_1} \alpha_t^2 \left[\widehat{\mathrm{Var}}_{i \in \mathcal{I}_t}\big(\widehat{\zeta}_t(i)\big) - \frac{1}{|\mathcal{I}_t|} \sum_{i \in \mathcal{I}_t} \frac{\widehat{s}_t^2(i)}{B_{\mathrm{cond}}} \right]_+, \qquad \widehat{R^\star} = \max_{1 \le t \le n_1} \alpha_t \max_{i \in \mathcal{I}_t} |\widehat{\zeta}_t(i)|.$$

Here, $[z]_+ = \max\{z, 0\}$. The optional one-sided upper-confidence bound correction for $R^\star$ is implemented in the code but is not used in the reported tables.

Algorithm 1 MCVarRange$(X, y^{(0)}, y^{(1)}, \mathcal{S}_1; \Pi, B_i, B_{\text{cond}}, B_J)$ — compute $(\widehat{V}^\star, \widehat{R}^\star)$

Require: X, $y^{(0)}$, $y^{(1)}$, realized $\mathcal{S}_1$; reveal permutation Π; budgets $B_i, B_{\text{cond}}, B_J$
Ensure: $\widehat{V^\star}$, $\widehat{R^\star}$

```
V̂* ← 0; R̂* ← 0
for t = 1 to n_1 do
    S_{t-1}^past ← {Π(1), ..., Π(t-1)};   R_{t-1} ← [n] \ S_{t-1}^past
    if |R_{t-1}| ≤ B_i then
        I_t ← R_{t-1}
    else
        I_t ← a uniform subsample of size B_i from R_{t-1}
    end if
    Initialize ζ̂_t(i) ← 0 for all i ∈ I_t
    for i ∈ I_t do
        for b = 1 to B_cond do
            Draw T_{b,i} ~ Unif(R_{t-1}\{i} choose ℓ_t - 1)
            if B_J = 0 or |R_{t-1} \ ({i} ∪ T_{b,i})| ≤ B_J then
                J_{b,i} ← R_{t-1} \ ({i} ∪ T_{b,i})
            else
                J_{b,i} ← a uniform subsample of size B_J from R_{t-1} \ ({i} ∪ T_{b,i})
            end if
            Compute Δ_{iJ} f(S_{t-1}^prox(i, T_{b,i})) for all J ∈ J_{b,i}
            ζ̂_t(i) ← ζ̂_t(i) + (1/B_cond) · (1/|J_{b,i}|) Σ_{J∈J_{b,i}} Δ_{iJ} f(S_{t-1}^prox(i, T_{b,i}))
        end for
    end for
    Compute ŝ_t^2(i) from the B_cond completion draws
    V̂* ← V̂* + α_t^2 [Var_emp{ ζ̂_t(i) : i ∈ I_t } − (1/|I_t|) Σ_{i∈I_t} ŝ_t^2(i)/B_cond]_+
    R̂* ← max(R̂*, α_t max_{i∈I_t} |ζ̂_t(i)|)
end for
return (V̂*, R̂*)
```

Algorithmic Pseudocode. The entire procedure is summarized in Algorithm 1 in a pseudocode format.

D.1.2 Monte Carlo Estimation of Bias Parameter

Choose budget parameters $B_S \in \mathbb{Z}_{>0}$.

1. Draw i.i.d. treatment sets $\mathcal{S}_1^{(b)} \sim \text{Unif}\binom{[n]}{n_1}$ for $b \in [B_S]$, and set $\mathcal{S}_0^{(b)} = [n] \setminus \mathcal{S}_1^{(b)}$.
2. For each $b \in [B_S]$, compute

$$\widehat{\Gamma}_b = \frac{1}{2n_1 n_0} \sum_{i \in \mathcal{S}_1^{(b)}} \sum_{j \in \mathcal{S}_0^{(b)}} \big(\Delta_{ij} f(\mathcal{S}_1^{(b)})\big)^2,$$

$$\widehat{Lf}_b = \frac{1}{n_1 n_0} \sum_{i \in \mathcal{S}_1^{(b)}} \sum_{j \in \mathcal{S}_0^{(b)}} \Delta_{ij} f(\mathcal{S}_1^{(b)}),$$

$$f_b = f(\mathcal{S}_1^{(b)}) = \hat{\tau}(\mathcal{S}_1^{(b)}) - \tau.$$

3. Then set

$$\widehat{\mathbb{E}\Gamma(f)} = \frac{1}{B_S} \sum_{b=1}^{B_S} \widehat{\Gamma}_b, \qquad \widehat{\text{Var}}(f) = \text{Var}_{\text{emp}}\big(\{f_b\}_{b \le B_S}\big),$$

and

$$\widehat{\lambda^\star} = \max\left\{ \text{gap}_{n,n_1}, \frac{\widehat{\mathbb{E}\Gamma(f)}}{\widehat{\text{Var}}(f)} \right\}.$$

Algorithm 2 MCBias$(X, y^{(0)}, y^{(1)}, n_1; B_S, B_{\text{pair}})$ — estimate $B^\star$ for $f(\mathcal{S}) = \hat{\tau}_{\mathsf{OLS}}(\mathcal{S}) - \tau$

Require: X, $y^{(0)}$, $y^{(1)}$, treated size n_1; budgets B_S, B_{pair}
Ensure: $\widehat{\mathbb{E}\Gamma(f)}$, $\widehat{\lambda^\star}$, $\widehat{B^\star}$

```
for b = 1 to B_S do
    Draw S_1^(b) ~ Unif(binom(n, n_1)); set S_0^(b) = [n] \ S_1^(b)
    f_b ← τ̂(S_1^(b)) − τ
    if B_pair ≥ n_1 n_0 then
        Γ̂_b ← (1/(2 n_1 n_0)) Σ_{i∈S_1^(b)} Σ_{j∈S_0^(b)} (Δ_ij f(S_1^(b)))^2
        L̂f_b ← (1/(n_1 n_0)) Σ_{i∈S_1^(b)} Σ_{j∈S_0^(b)} Δ_ij f(S_1^(b))
    else
        Draw B_pair ordered pairs (i_k, j_k) uniformly at random from S_1^(b) × S_0^(b) without replacement
        Γ̂_b ← (1/(2 B_pair)) Σ_{k=1}^{B_pair} (Δ_{i_k j_k} f(S_1^(b)))^2
        L̂f_b ← (1/B_pair) Σ_{k=1}^{B_pair} Δ_{i_k j_k} f(S_1^(b))
    end if
end for
EΓ̂(f) ← (1/B_S) Σ_{b=1}^{B_S} Γ̂_b,   V̂ar(f) ← Var_emp({f_b}_{b≤B_S})
λ̂* ← max{ gap_{n,n_1}, EΓ̂(f) / V̂ar(f) }
B̂* ← [ B_S^{-1} Σ_{b=1}^{B_S} (L̂f_b + λ̂* f_b)^2 ]^{1/2} / λ̂*
return (EΓ̂(f), λ̂*, B̂*)
```

Here, $\widehat{\mathrm{Var}}(f)$ is the sample variance of $f(\mathcal{S}_1^{(b)}) = \hat{\tau}(\mathcal{S}_1^{(b)}) - \tau$ over $b \le B_S$, and $\mathrm{gap}_{n,n_1} = \frac{n}{n_1 n_0}$.

4. Finally, estimate the corrected oracle bias parameter by

$$\widehat{B^\star} = \frac{\left[B_S^{-1} \sum_{b=1}^{B_S} \left(\widehat{Lf}_b + \widehat{\lambda^\star} f_b\right)^2\right]^{1/2}}{\widehat{\lambda^\star}}.$$

If $|\mathcal{S}_1^{(b)} \times \mathcal{S}_0^{(b)}|$ is too large, we subsample B_{pair} ordered pairs (i, j) uniformly without replacement to approximate $\widehat{\Gamma}_b$ and $\widehat{Lf}_b$. See Algorithm 2 for a summary of the procedure in pseudocode format.

D.1.3 Computational Complexity and Monte Carlo Accuracy

For OLS–RA, all swap effects are evaluated via the paired deletion–insertion telescope and the branch-aware row-update identities, so no OLS model re-fitting is required. If $r_{S_\alpha} = \mathrm{rank}(X_{S_\alpha})$, the one-time arm-wise preprocessing cost is $O\left(\sum_{\alpha\in\{0,1\}} \left(|S_\alpha|\, p\, r_{S_\alpha} + r_{S_\alpha}^3\right)\right)$.

Thereafter, the dominant cost is the number of sampled swap evaluations. In the sampled-J implementation, Algorithm 1 and Algorithm 2 respectively perform

$$N_\Delta^{\text{var}} = \sum_{t=1}^{n_1} |\mathcal{I}_t|\, B_{\text{cond}}\, B_J^{\text{eff}}, \qquad N_\Delta^{\text{bias}} = B_S \min(B_{\text{pair}}, n_1 n_0)$$

swap evaluations, where $B_J^{\text{eff}} = n_0$ if $B_J = 0$ and $B_J^{\text{eff}} = B_J$ otherwise. Thus B_i, B_{cond}, and B_J control the resolution and MC noise of $(\widehat{V}^\star, \widehat{R}^\star)$, while B_S and B_{pair} control the stability of $\widehat{B}^\star$. The reported runs use moderate budgets to balance computational cost and Monte Carlo variability.

All in all, the total arithmetic count is

$$O\left(\sum_{\alpha\in\{0,1\}} \left(|S_\alpha|\, p\, r_{S_\alpha} + r_{S_\alpha}^3\right)\right) + O\Big((N_\Delta^{\text{var}} + N_\Delta^{\text{bias}}) \cdot c_{\text{swap}}\Big),$$

where $c_{\text{swap}} = O(r_{S_\alpha})$ for deletions and $c_{\text{swap}} \in \{O(r_{S_\alpha} p), O(r_{S_\alpha}|S_\alpha|)\}$ for insertions, depending on whether the cross-terms are precomputed or not. The loops over t, $i \in \mathcal{I}_t$, and (i,j) pairs are parallelizable. With pre-computation, the memory requirement for the cached factors is $O(|S_\alpha| r_{S_\alpha})$.

D.2 Simulation Studies

D.2.1 Setup

Data Generation Process. We choose $n \in \mathbb{Z}_{>0}$ and let $\rho = 0.3$; $n_1 = \text{round}(\rho n)$; $n_0 = n - n_1$. For $\gamma \in \{0, 0.25, 0.5, 0.75, 1.0, 1.25, 1.5\}$, we set $p = \lceil n^\gamma \rceil$ and form $X \in \mathbb{R}^{n \times p}$ in a two-stage process as follows. First of all, we generate a single "master" matrix $\widetilde{X} \in \mathbb{R}^{n \times \lceil n^{1.5} \rceil}$ with i.i.d. entries from the standard Gaussian distribution $\mathcal{N}(0,1)$, and hold it fixed for controlled comparison of covariate dimensionality p across all conditions. Thereafter, for each p, we create X by subsampling the first p columns of $\widetilde{X}$; then column-center and rescale so that $\mathbf{1}^\top X = 0$ and each column of X has norm $\sqrt{n}$.

With X instantiated, we generate potential outcomes $y^{(1)} = X\beta_1^* + \varepsilon^{(1)}$ and $y^{(0)} = X\beta_0^* + \varepsilon^{(0)}$. In the classical regime where Σ_α is invertible, arm-wise OLS with intercept is invariant to adding a common linear signal $X\beta$ to both arms: the fitted slope in each arm shifts by β, while the fitted intercept, and hence $\hat{\tau} - \tau$, remains unchanged (cf. (4) in Section 2.2). In overparameterized regimes, one may analogously use the minimum-ℓ_2-norm OLS fit, replacing Σ_α^{-1} by $\Sigma_\alpha^\dagger$ [Shen et al., 2023]; we do not rely on an exact invariance statement there though. Thus, in our experiments, we set $\beta_1^* = \beta_0^* = 0$ for simplicity and take $\varepsilon^{(1)}$ and $\varepsilon^{(0)}$ to have i.i.d. $\mathcal{N}(0,1)$ entries. Preliminary runs with a common signal gave similar qualitative conclusions; see Remark D.1 below.

We keep $(X, y^{(1)}, y^{(0)})$ fixed while sampling treatment assignments (complete randomization) and compare the OLS–RA estimator to the DiM baseline and empirical measures for diagnostics. For each realized instance $(X, y^{(0)}, y^{(1)})$, we draw N treatment assignments under complete randomization and compute all statistics assignment-wise; we repeat this R times with independent seeds. Unless otherwise noted in each experiment's caption, we use $R = 20$ and $N = 500$.

Remark D.1. We also tested models with a common linear signal by taking a random unit vector β^* and setting $\beta_1^* = \beta_0^* = \theta\,\beta^*$ with $\theta \geq 0$. Preliminary runs with stronger common linear signal (i.e., $y^{(\alpha)} = \theta\, X\beta^\star + \varepsilon^{(\alpha)}$ with $\theta > 0$) gave qualitatively similar results and did not materially change our conclusions to follow, hence they are omitted for brevity. A more extensive investigation is deferred to future work.

Oracle Quantities $V^\star, R^\star, B^\star$ and Monte Carlo Estimation. We consider two ATE estimators:

(i) the difference-in-means estimator $\hat{\tau}_{\text{DiM}}$, cf. (2), and

(ii) the OLS-adjusted estimator $\hat{\tau}_{\text{OLS}}$, cf. (6).

For each estimator, we estimate the concentration parameters $V^\star, R^\star$, cf. (29), and the bias parameter $B^\star$, cf. (31). For DiM, the remaining-pool oracle quantities $V^\star$ and $R^\star$ admit closed forms, and $B^\star = 0$. For OLS–RA, we estimate $(V^\star, R^\star, B^\star)$ by Monte Carlo, yielding $(\widehat{V^\star}, \widehat{R^\star}, \widehat{B^\star})$; see Appendix D.1 for details.

We recall that with the assignment-exposure filtration $(\mathcal{F}_t)$ and martingale $M_t = \mathbb{E}[f(\mathcal{S}_1) \mid \mathcal{F}_t]$, the increment sequence is $D_t \coloneqq M_t - M_{t-1}$. Thus the realized predictable quadratic variation (PQV) and the local range along a reveal order Π are

$$V(\mathcal{S}_1, \Pi) = \sum_{t=1}^{n_1} \mathrm{Var}\big(D_t \mid \mathcal{F}_{t-1}\big), \qquad R(\mathcal{S}_1, \Pi) = \max_{1 \leq t \leq n_1} \operatorname{ess\,sup}\big(\, |D_t| \;\big|\; \mathcal{F}_{t-1}\big),$$

cf. (28). By the reveal–swap identity (16), $D_t = -\,\alpha_t\, \zeta_t(I)$ with $\alpha_t = \frac{n_0}{n-t+1}$ and $\zeta_t(i) = \mathbb{E}_{J,\mathcal{T}}[\Delta_{iJ} f(\mathcal{S}_{t-1}^{\text{prox}}(i, \mathcal{T})) \mid \mathcal{F}_{t-1}]$ (cf. (17)). Consequently,

$$V^\star = \sum_{t=1}^{n_1} \alpha_t^2\, \mathrm{Var}\big(\zeta_t(I)\big), \qquad R^\star = \max_{1 \leq t \leq n_1} \alpha_t \max_{i \in \mathcal{R}_{t-1}} |\zeta_t(i)|,$$

which Algorithm 1 (presented in Appendix D.1) approximates via Rao–Blackwellization over controls J.

Also, recall that the bias parameter is

$$B^\star = \frac{\sqrt{\mathbb{E}\big(Lf + \lambda^\star f\big)^2}}{\lambda^\star}, \qquad \lambda^\star = \frac{\mathbb{E}\Gamma(f)}{\mathrm{Var}(f)},$$

where

$$\Gamma(f)(\mathcal{S}_1) = \frac{1}{2n_1 n_0} \sum_{i\in\mathcal{S}_1} \sum_{j\in\mathcal{S}_0} \big(\Delta_{ij} f(\mathcal{S}_1)\big)^2, \qquad (Lf)(\mathcal{S}_1) = \frac{1}{n_1 n_0} \sum_{i\in\mathcal{S}_1} \sum_{j\in\mathcal{S}_0} \Delta_{ij} f(\mathcal{S}_1).$$

Algorithm 2 (see Appendix D.1) estimates $\mathbb{E}\Gamma(f)$, $\lambda^\star$, and the residual norm $\mathbb{E}(Lf + \lambda^\star f)^2$, yielding the Monte Carlo estimate $\widehat{B^\star}$.

Once the required oracle quantities are available—exactly in the DiM case, and via $(\widehat{V^\star}, \widehat{R^\star}, \widehat{B^\star})$ in the OLS–RA case—for any given $\delta \in (0,1)$, for each instantiated treatment set $\mathcal{S}_1^{(k)}$, we compute the oracle finite-sample CI for τ of desired coverage $1-\delta$ using (32) as

$$\mathrm{CI}_\delta(\mathcal{S}_1^{(k)}) := \Big[\hat\tau - \widehat{r}_\delta(\mathcal{S}_1^{(k)}),\ \hat\tau + \widehat{r}_\delta(\mathcal{S}_1^{(k)})\Big] \qquad \text{where} \qquad \widehat{r}_\delta(\mathcal{S}_1^{(k)}) := \sqrt{2\widehat{V^\star}\,\log\frac{2}{\delta}} + \frac{\widehat{R^\star}}{3}\log\frac{2}{\delta} + \widehat{B^\star}.$$

Remark D.2. Here, the concentration parameter estimates $\widehat{V^\star} = \widehat{V^\star}(\mathcal{S}_1^{(k)})$ and $\widehat{R^\star} = \widehat{R^\star}(\mathcal{S}_1^{(k)})$ depend on the particular treatment assignment $\mathcal{S}_1^{(k)}$, whereas the bias parameter estimate $\widehat{B^\star}$ does not.

Evaluation and Diagnostics. For each instance $k \in [N]$, we compute the ATE estimator $\hat\tau \in \{\hat\tau_{\mathrm{DiM}}, \hat\tau_{\mathrm{OLS}}\}$. Let $\{\hat\tau_k\}_{k=1}^N$ be the N estimates of a given method. At $\delta = 0.05$, we report the following performance metrics:

- *Coverage rate:*

$$\widetilde{\mathrm{CR}}_\delta = \frac{1}{N}\sum_{k=1}^{N} \mathbb{I}\big\{\,|\hat\tau_k - \tau| \le \widehat{r}_\delta(\mathcal{S}_1^{(k)})\,\big\}.$$

- *Aggregate CI radii:* the mean, median, and variance of the CI radius $\widehat{r}_\delta$ across N assignments.
- *Inter-quantile benchmark:* $q_{0.975} - q_{0.025}$, the empirical 97.5-th percentile of $\hat\tau - \tau$ minus 2.5-th percentile of $\hat\tau - \tau$ across the N assignments.

These metrics are used to assess how tight or loose the constructed confidence intervals are.

Aggregation of results for reporting is nested and experiment-specific as follows. For Experiment 1, coverage is averaged across outer draws, while CI widths and inter-percentile widths are summarized by the median of the within-draw means; the corresponding parenthetical entries report the matching within-draw or across-draw variance summaries. For Experiment 2, each outer draw is first summarized over its N assignments, and the reported entries are medians across the R outer draws, with variances across outer draws in parentheses. Throughout, we report *CI widths* (twice the radius), not the radius.

In addition to coverage and width summaries, we report diagnostic quantities that help interpret the Monte Carlo estimates. For DiM, we compare the exact remaining-pool oracle quantities with terminal-assignment diagnostics computed after conditioning on the realized final treatment/control split. For OLS–RA, we report three algorithm-tied diagnostics—V_{PQV}, R_{swap}, and B_{emp}—all computed from the same sampled candidate units, completion sets, and admissible controls that produce $(\widehat{V^\star}, \widehat{R^\star}, \widehat{B^\star})$. Here PQV stands for *predictable quadratic variation*. These diagnostics are not alternative oracle parameters and should not be interpreted as upper or lower bounds for $(V^\star, R^\star, B^\star)$; rather, they are internal reference scales that help identify where slack enters in the Monte Carlo procedure. In particular, the current implementation forms the direct MC estimate $\widehat{R^\star}$ from the averaged stepwise quantities and records R_{swap} from the corresponding raw sampled swap effects.

Computing Environment. All experiments were run on a single Apple Mac mini (M4 Pro, Apple silicon) under macOS using Python (NumPy/SciPy; CPU only).

D.2.2 Experiment 1: Finite-sample versus Asymptotic CIs for Difference-in-Means

The difference-in-means (DiM) estimator is unbiased under complete randomization and admits asymptotic normality with the classical Neyman variance formula. It is also useful here because, for a realized reveal order Π, the remaining-pool oracle quantities $V^\star$ and $R^\star$ admit explicit closed forms.

To state this precisely, we define

$$a_i \coloneqq \frac{y_i^{(1)}}{n_1} + \frac{y_i^{(0)}}{n_0}, \qquad m_t \coloneqq |\mathcal{R}_{t-1}| = n - t + 1, \qquad \bar{a}_{t-1} \coloneqq \frac{1}{m_t} \sum_{u \in \mathcal{R}_{t-1}} a_u.$$

Then, for each $i \in \mathcal{R}_{t-1}$,

$$\zeta_t(i) = \frac{\sum_{u \in \mathcal{R}_{t-1}} a_u - m_t a_i}{m_t - 1} = -\frac{m_t}{m_t - 1}\big(a_i - \bar{a}_{t-1}\big).$$

Since $\alpha_t = n_0/(n-t+1) = n_0/m_t$, it follows that

$$V^\star = \sum_{t=1}^{n_1} \alpha_t^2 \mathrm{Var}_{i \sim \mathrm{Unif}(\mathcal{R}_{t-1})}\big(\zeta_t(i)\big) = \sum_{t=1}^{n_1} \frac{n_0^2}{(m_t-1)^2} \cdot \frac{1}{m_t} \sum_{i \in \mathcal{R}_{t-1}} \big(a_i - \bar{a}_{t-1}\big)^2, \tag{64}$$

$$R^\star = \max_{1 \le t \le n_1} \alpha_t \max_{i \in \mathcal{R}_{t-1}} |\zeta_t(i)| = \max_{1 \le t \le n_1} \frac{n_0}{m_t - 1} \max_{i \in \mathcal{R}_{t-1}} \big|a_i - \bar{a}_{t-1}\big|. \tag{65}$$

Thus, for the DiM estimator, $V^\star$ and $R^\star$ are computable exactly from the remaining-pool means, variances, and maximal deviations of the score sequence $\{a_i\}_{i=1}^n$ along the reveal path.

In Experiment 1, we use the DiM estimator setting as a sanity check for the finite-sample (FS) construction, asking: (i) whether the proposed FS CIs achieve at least nominal coverage; (ii) how conservative their widths are relative to Wald intervals and empirical randomization dispersion; and (iii) how the remaining-pool oracle quantities scale with n.

Setup and Evaluation. We follow the general data-generation protocol in Section D.2.1. DiM does not depend on covariates, so we vary only $n \in \{10, 20, 40, 80, 160, 320, 640\}$ and set $\gamma = 0$ (hence $p = 1$). Here we set $\theta = 0$, so $y^{(\alpha)} = \varepsilon^{(\alpha)}$ and the finite-population estimand is $\tau = n^{-1} \sum_{i=1}^n (y_i^{(1)} - y_i^{(0)})$, which is centered at 0 and fluctuates at scale $O_P(n^{-1/2})$.

For each instantiated finite population of potential outcomes and each complete-randomization assignment, we compute:

- *Finite-sample (FS) CI for DiM* with radius

$$r_{\mathrm{DiM}}(\delta) = \sqrt{2\, V_{\mathrm{rp}}^\star \log \frac{2}{\delta}} + \frac{R_{\mathrm{rp}}^\star}{3} \log \frac{2}{\delta},$$

  where $(V_{\mathrm{rp}}^\star, R_{\mathrm{rp}}^\star)$ are the exact remaining-pool oracle quantities for DiM under the reveal law of Proposition 2; we set $B^\star = 0$ because DiM is design-unbiased.

- *Asymptotic Wald (Neyman) CI* with radius

$$r_{\mathrm{Wald}}(\delta) = z_{1-\delta/2} \sqrt{\widehat{\mathrm{Var}}_{\mathrm{Ney}}}, \qquad \widehat{\mathrm{Var}}_{\mathrm{Ney}} = \frac{s_1^2}{n_1} + \frac{s_0^2}{n_0},$$

  where s_α^2 is the sample variance of observed outcomes in arm α under the current assignment; $z_{1-\delta/2}$ is the standard normal quantile.

- *Empirical inter-percentile width* of $\hat{\tau}_{\mathrm{DiM}} - \tau$, $q_{1-\delta/2} - q_{\delta/2}$. For instance, when $\delta = 0.05$, we compute the empirical 2.5–97.5% widths of $\hat{\tau}_{\mathrm{DiM}} - \tau$.

Table 1: Coverage and width summaries for DiM at $\delta = 0.05$ (Experiment 1). For each n, we report FS and Wald coverage, CI width (i.e., 2×radius), and the empirical 2.5%–97.5% width of $\hat{\tau}_{\mathrm{DiM}} - \tau$, computed from $R = 20$ outer replicates and $N = 500$ assignments per replicate. For coverage and percentile widths, the parenthetical entries are variances across the R outer replicates of the replicate-level summaries. For CI widths, the parenthetical entries are the medians across the R outer replicates of the within-replicate width variances.

n	Finite-sample CI		Asymptotic CI		Inter-percentile range
	Coverage	CI width	Coverage	CI width	$q_{0.975} - q_{0.025}$
10	1.000 (0.000)	4.413 (0.108)	0.945 (0.002)	2.504 (0.601)	1.903 (0.236)
20	1.000 (0.000)	2.933 (0.023)	0.976 (0.000)	1.933 (0.093)	1.426 (0.075)
40	1.000 (0.000)	2.068 (0.006)	0.983 (0.000)	1.423 (0.027)	1.087 (0.023)
80	1.000 (0.000)	1.256 (0.001)	0.987 (0.000)	0.948 (0.006)	0.697 (0.004)
160	0.999 (0.000)	0.885 (0.000)	0.988 (0.000)	0.691 (0.002)	0.532 (0.001)
320	0.998 (0.000)	0.593 (0.000)	0.987 (0.000)	0.481 (0.000)	0.364 (0.000)
640	0.999 (0.000)	0.404 (0.000)	0.990 (0.000)	0.338 (0.000)	0.258 (0.000)

Table 2: Summary of remaining-pool oracle quantities and terminal-assignment diagnostics for DiM (Experiment 1). For each n, we report medians across outer draws, with variances across outer draws in parentheses. Here, the subscript rp in $V^\star_{\mathrm{rp}}$ and $R^\star_{\mathrm{rp}}$ indicates they are computed using the remaining-pool identities (64), (65).

n	Variance		Range	
	$V^\star_{\mathrm{rp}}$	V_{term}	$R^\star_{\mathrm{rp}}$	R_{term}
10	0.257 (0.022)	0.055 (0.001)	0.657 (0.043)	0.386 (0.009)
20	0.143 (0.003)	0.063 (0.001)	0.385 (0.005)	0.267 (0.002)
40	0.082 (0.000)	0.049 (0.000)	0.211 (0.002)	0.153 (0.000)
80	0.033 (0.000)	0.025 (0.000)	0.111 (0.000)	0.082 (0.000)
160	0.018 (0.000)	0.016 (0.000)	0.065 (0.000)	0.048 (0.000)
320	0.009 (0.000)	0.008 (0.000)	0.036 (0.000)	0.027 (0.000)
640	0.004 (0.000)	0.004 (0.000)	0.019 (0.000)	0.014 (0.000)

Empirical Diagnostics for DiM. For DiM, the reveal law of Proposition 2 yields exact remaining-pool oracle quantities $(V^\star, R^\star)$. Separately, we report terminal-assignment diagnostics $(V_{\mathrm{term}}, R_{\mathrm{term}})$, obtained by conditioning on the realized final treatment/control split. Concretely, we define

$$V_{\mathrm{term}} = \sum_{t=1}^{n_1} \alpha_t^2 \mathrm{Var}\big(\mu_i : i \in S_1 \setminus \{\Pi(1), \ldots, \Pi(t-1)\}\big), \qquad \text{and} \qquad R_{\mathrm{term}} = \max_{i \in S_1} \alpha_{\mathrm{pos}(i)} |\mu_i|,$$

where $\mu_i = \frac{\overline{y^{(1)}_{S_0}} - y^{(1)}_i}{n_1} - \frac{y^{(0)}_i - \overline{y^{(0)}_{S_0}}}{n_0}$, and $\mathrm{pos}(i)$ is the reveal position of i in Π. These terminal diagnostics are inexpensive empirical surrogates for assessing conservatism, but they are not the oracle quantities used in the FS CI. They condition on the realized final treatment/control split, whereas the oracle quantities follow the remaining-pool reveal law and therefore account for the randomness of the eventual control set.

Results. Table 1 summarizes coverage and CI widths, and Table 2 compares the exact remaining-pool oracle quantities with terminal-assignment diagnostics.

The FS CIs achieve at least nominal coverage (which is 0.95 here) across all reported n, but they are conservative. The FS widths exceed the Wald widths throughout; for example, at $n = 10$, the FS width is 4.413 versus Wald width 2.504, while the empirical 2.5%–97.5% width of $\hat{\tau}_{\mathrm{DiM}} - \tau$ is 1.903. Thus Experiment 1 should be read primarily as a validity and conservatism sanity check for the general finite-sample bound.

Table 2 helps localize, though not fully identify, the source of this conservatism. The exact remaining-pool oracle quantities $V^\star$ and $R^\star$ are noticeably larger than the terminal-assignment diagnostics V_{term} and R_{term}, especially at small n; for example, at $n = 10$, $V^\star = 0.257$ versus $V_{\mathrm{term}} = 0.055$, and $R^\star = 0.657$ versus $R_{\mathrm{term}} = 0.386$.

Because $V^\star$ and $R^\star$ are computed exactly here, this gap is not a Monte Carlo artifact. At the same time, V_{term} and R_{term} are only diagnostic surrogates—not oracle targets and not quantities known to yield valid confidence intervals—so these comparisons should not be interpreted as calibration identities. Rather, they suggest that a substantial part of the conservatism may already enter through the remaining-pool reveal construction, and perhaps more fundamentally through the martingale/Freedman concentration step. Determining whether this looseness is intrinsic to this line of argument, or can be reduced by a sharper finite-population concentration method, remains an interesting direction for future work.

D.2.3 Experiment 2: Finite-sample CIs for OLS–RA

Objective. Having completed basic sanity checks with the DiM estimator, here we assess finite-sample (FS) confidence intervals (CIs) for the OLS regression-adjusted estimator (OLS–RA). Our goals in this experiment are: (i) to ensure at least nominal coverage across various (n, γ); and (ii) to quantify the tightness or conservatism of the FS-CI widths relative to an asymptotic DiM baseline and the empirical 2.5%–97.5% inter-percentile range of $\hat{\tau}_{\text{OLS}} - \tau$.

Setup and Evaluation. We follow the general data-generation protocol in Section D.2.1. We consider $n \in \{25, 50\}$ and $\gamma \in \{0, 0.25, 0.5, 0.75, 1.0, 1.25, 1.5\}$, so that $p = \lceil n^\gamma \rceil$. For the reported run, we use $R = 20$, $N = 500$, and MC budgets $(B_S, B_{\text{pair}}, B_i, B_{\text{cond}}, B_J) = (30, 30, 10, 10, 10)$. The bias parameter $\widehat{B^\star}$ is estimated once per finite-population instance using Algorithm 2, while $(\widehat{V^\star}, \widehat{R^\star})$ are estimated assignment-wise using Algorithm 1.

Empirical Diagnostics for OLS–RA. Recall $\widehat{\zeta}_t(i)$ and the within-i completion variance $\widehat{s}_t^2(i)$ from Section D.1.1. We report three diagnostics:

$$
\begin{aligned}
V_{\text{PQV}} &:= \sum_{t=1}^{n_1} \alpha_t^2 \, \widehat{\mathrm{Var}}_{i \in \mathcal{I}_t}\big(\widehat{\zeta}_t(i)\big), \\
R_{\text{swap}} &:= \max_{1 \le t \le n_1} \alpha_t \max_{i \in \mathcal{I}_t} \max_{1 \le b \le B_{\text{cond}}} \max_{J \in \mathcal{J}_{b,i}} \big|\Delta_{iJ} f\big(\mathcal{S}_{t-1}^{\text{prox}}(i, \mathcal{T}_{b,i})\big)\big|, \\
B_{\text{emp}} &:= \left| \frac{1}{N} \sum_{k=1}^{N} \hat{\tau}^{(k)} - \tau \right|.
\end{aligned}
$$

These quantities play different roles.

- V_{PQV} is the *predictable quadratic variation* (PQV) before the within-i Monte Carlo noise correction. In the current implementation,
$$\widehat{V^\star} = \sum_{t=1}^{n_1} \alpha_t^2 \left[\widehat{\mathrm{Var}}_{i \in \mathcal{I}_t}\big(\widehat{\zeta}_t(i)\big) - \frac{1}{|\mathcal{I}_t|} \sum_{i \in \mathcal{I}_t} \frac{\widehat{s}_t^2(i)}{B_{\text{cond}}} \right]_+ ,$$
where $[z]_+ = \max\{z, 0\}$. Thus V_{PQV} is not a separate oracle target; it is the pre-de-noising version of the variance aggregate built from the same sampled $\widehat{\zeta}_t(i)$ values as $\widehat{V^\star}$. Hence $V_{\text{PQV}} - \widehat{V^\star}$ measures the size of the within-i Monte Carlo noise correction, rather than conservatism relative to the oracle $V^\star$.
- R_{swap} records the sampled raw-swap scale before conditional averaging, whereas $\widehat{R^\star}$ is built from the averaged quantities $\widehat{\zeta}_t(i)$. Thus $R_{\text{swap}} - \widehat{R^\star}$ measures the smoothing induced by Rao–Blackwellization and completion averaging. When $B_J > 0$, R_{swap} is not an exact all-control envelope, but a diagnostic computed from the same sampled objects as $\widehat{R^\star}$.
- B_{emp} provides an empirical reference scale for the bias proxy $\widehat{B^\star}$.

These diagnostics are useful because they are computed from the same sampled objects as $(\widehat{V^\star}, \widehat{R^\star}, \widehat{B^\star})$, so they isolate distinct components of the implemented Monte Carlo pipeline. Unlike the DiM case, we do not report terminal-assignment diagnostics for OLS–RA, because there is no comparably simple closed-form terminal analogue.

Table 3: Coverage and width summaries for OLS–RA at $\delta = 0.05$ (Experiment 2). For each (n, γ), we report FS and Wald coverage, CI width (i.e., 2×radius), and the empirical 2.5%–97.5% width of $\hat{\tau}_{\mathrm{OLS}} - \tau$, obtained by first summarizing each outer replicate over its $N = 500$ assignments and then taking medians across the $R = 20$ outer replicates. For coverage, the parenthetical entries are variances across outer replicates of the replicate-level coverage rates. For CI widths, the parenthetical entries are the medians across outer replicates of the within-replicate width variances. For the empirical inter-percentile widths, the parenthetical entries are variances across outer replicates of the replicate-level IPRs.

n	γ	p	Finite-sample CI (RA)		Asymptotic CI (DiM)		Inter-percentile range
			Coverage	CI width	Coverage	CI width	$q_{0.975} - q_{0.025}$
25	0.00	1	1.000 (0.000)	2.982 (0.212)	0.953 (0.000)	2.413 (0.083)	1.316 (0.036)
25	0.25	3	1.000 (0.000)	4.100 (0.879)	0.960 (0.000)	2.389 (0.104)	1.755 (0.039)
25	0.50	5	1.000 (0.000)	6.380 (14.157)	0.948 (0.000)	2.310 (0.123)	2.453 (0.065)
25	0.75	12	1.000 (0.000)	5.874 (1.993)	0.961 (0.000)	2.327 (0.092)	2.313 (0.052)
25	1.00	25	1.000 (0.000)	4.450 (0.385)	0.958 (0.000)	2.359 (0.111)	1.849 (0.076)
25	1.25	56	1.000 (0.000)	4.922 (0.195)	0.952 (0.000)	2.540 (0.132)	2.103 (0.078)
25	1.50	125	1.000 (0.000)	4.313 (0.112)	0.960 (0.000)	2.240 (0.097)	1.887 (0.058)
50	0.00	1	1.000 (0.000)	2.194 (0.031)	0.966 (0.000)	1.788 (0.030)	0.965 (0.013)
50	0.25	3	1.000 (0.000)	2.242 (0.082)	0.961 (0.000)	1.660 (0.027)	0.970 (0.020)
50	0.50	8	1.000 (0.000)	3.284 (0.535)	0.967 (0.000)	1.707 (0.031)	1.358 (0.042)
50	0.75	19	1.000 (0.000)	4.811 (1.766)	0.967 (0.000)	1.704 (0.029)	1.862 (0.060)
50	1.00	50	1.000 (0.000)	3.399 (0.121)	0.962 (0.000)	1.727 (0.034)	1.419 (0.014)
50	1.25	133	1.000 (0.000)	3.228 (0.053)	0.970 (0.000)	1.683 (0.033)	1.434 (0.049)
50	1.50	354	1.000 (0.000)	3.254 (0.033)	0.964 (0.000)	1.704 (0.032)	1.454 (0.024)

Table 4: Summary of Monte Carlo estimates and empirical diagnostics for OLS–RA (Experiment 2). For each (n, γ), we report medians across the R outer replicates, with variances across outer replicates in parentheses. We do not report terminal-assignment diagnostics here because, unlike the DiM case, OLS–RA has no comparably simple closed-form terminal analogue; instead we report the algorithm-tied diagnostics V_{PQV}, R_{swap}, and B_{emp}.

n	γ	p	Variance (assignment-level)		Range		Bias	
			$\widehat{V^\star}$	V_{PQV}	$\widehat{R^\star}$	R_{swap}	$\widehat{B^\star}$	B_{emp}
25	0.00	1	0.117 (0.001)	0.119 (0.001)	0.311 (0.003)	0.891 (0.020)	0.196 (0.004)	0.021 (0.000)
25	0.25	3	0.175 (0.002)	0.186 (0.002)	0.423 (0.002)	1.913 (0.030)	0.253 (0.024)	0.035 (0.001)
25	0.50	5	0.295 (0.003)	0.380 (0.005)	0.703 (0.006)	7.280 (0.862)	0.381 (0.059)	0.043 (0.003)
25	0.75	12	0.338 (0.006)	0.372 (0.007)	0.674 (0.004)	2.323 (0.065)	0.359 (0.007)	0.062 (0.002)
25	1.00	25	0.239 (0.005)	0.248 (0.005)	0.470 (0.005)	1.096 (0.022)	0.306 (0.008)	0.030 (0.001)
25	1.25	56	0.320 (0.005)	0.324 (0.006)	0.477 (0.008)	0.954 (0.029)	0.326 (0.003)	0.035 (0.000)
25	1.50	125	0.255 (0.004)	0.257 (0.004)	0.444 (0.011)	0.782 (0.017)	0.254 (0.003)	0.024 (0.000)
50	0.00	1	0.062 (0.000)	0.062 (0.000)	0.179 (0.001)	0.390 (0.002)	0.198 (0.002)	0.011 (0.000)
50	0.25	3	0.065 (0.000)	0.066 (0.000)	0.202 (0.001)	0.535 (0.006)	0.198 (0.004)	0.009 (0.000)
50	0.50	8	0.112 (0.001)	0.126 (0.001)	0.326 (0.002)	1.366 (0.031)	0.280 (0.015)	0.026 (0.001)
50	0.75	19	0.186 (0.002)	0.221 (0.003)	0.506 (0.004)	2.434 (0.108)	0.389 (0.011)	0.026 (0.001)
50	1.00	50	0.134 (0.000)	0.139 (0.001)	0.294 (0.001)	0.683 (0.004)	0.332 (0.005)	0.026 (0.000)
50	1.25	133	0.142 (0.002)	0.144 (0.002)	0.267 (0.002)	0.535 (0.007)	0.269 (0.004)	0.011 (0.000)
50	1.50	354	0.149 (0.001)	0.150 (0.001)	0.265 (0.001)	0.505 (0.004)	0.280 (0.004)	0.017 (0.000)

Results. Table 3 reports coverage and widths for OLS–RA FS CIs and DiM Wald CIs, together with the empirical 2.5%–97.5% width of $\hat{\tau}_{\mathrm{OLS}} - \tau$. Table 4 reports the corresponding MC estimates and diagnostics.

The OLS–RA FS intervals attain empirical coverage 1.000 in all reported cells, which are always greater than the nominal coverage 0.95 but are conservative. The FS widths exceed the empirical inter-percentile widths

throughout; in this run, the width/IPR ratio ranges from about 2.24 to 2.60. For example, at $n = 25, \gamma = 0.50$, the FS width is 6.380 whereas the empirical IPR is 2.453, and at $n = 50, \gamma = 0.75$, the corresponding values are 4.811 and 1.862.

The diagnostics are best interpreted structurally rather than as direct lower or upper bounds on the oracle quantities. By construction, V_{PQV} is the pre-de-noising variance aggregate built from the same sampled $\widehat{\zeta}_t(i)$ values used in $\widehat{V^\star}$; therefore the gap $V_{\mathrm{PQV}} - \widehat{V^\star}$ measures the size of the within-i noise correction, not conservatism relative to the true $V^\star$. Likewise, R_{swap} records the raw sampled swap scale before conditional averaging, whereas $\widehat{R^\star}$ is built from the averaged stepwise quantities $\widehat{\zeta}_t(i)$; the gap between them quantifies the smoothing induced by Rao–Blackwellization. The fact that $\widehat{B^\star}$ exceeds B_{emp} in every reported cell shows that the bias proxy is also conservative in these experiments.

Finally, the empirical IPR is largest in moderate-dimensional regimes—near $\gamma = 0.50$ when $n = 25$, and near $\gamma = 0.75$ when $n = 50$. This is consistent with the Q-geometry interpretation in Section 4.2: covariates can reduce residual dispersion, while deletion and insertion sensitivities become largest near the arm-wise interpolation threshold. Larger-budget reruns should reduce MC noise, but are unlikely to change these qualitative conclusions.